ITERATED BELIEF REVISION:
THEORY & PRACTICE

A thesis submitted to the University of Manchester
for the degree of Master of Science
in the Faculty of Science and Engineering

2004

Jamie Lentin
Department of Mathematics

Contents

List of Figures

Abstract

In many cases we will learn contradictory facts about a situation, yet generally we can resolve these contradictions. If someone told us that a box was full of cherries but we open it to find the box is empty, we would not believe it was both full and empty, rather dismiss the first fact as false. Similarly, belief revision is the study of resolving contradictions in sets of logical propositions, in particular enabling us to add new sentences to the set without introducing inconsistencies.

We will study some main developments in this field; firstly how to represent a system capable of performing belief revision and what properties it should have, particularly when revising multiple times, or iterated revision. Secondly we shall investigate a selection of iterated revision operators based on this theory, and how each compares with our theory and intuitive ideas of how such an operator should work.

Declaration

No portion of the work referred to in this thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institution of learning.

Copyright

Acknowledgements

I would like to thank Jeff Paris, not only for all the help and advice given whilst supervising this work, but also for introducing me to belief revision as part of his course on nonmonotonic logic. Also I would like to thank the engineering and physical sciences research council, for their funding.

Chapter 1

Introduction to Belief Revision

Imagine we had a "store of knowledge" that we can add to. Once we receive new information, we add it into our store. This seems a simple concept, however what if we are trying to add information that is inconsistent with information we already knew?

Example 1.1. *Say we have the following facts.*

1. *I got a present in my stocking last night.*

2. *If I got a present in my stocking, it was delivered by Father Christmas.*

3. *When Father Christmas delivers presents, he uses flying reindeer to do it.*

   But were then told:-

4. *Flying reindeer don't exist.*

If we believed all of these facts at the same time, we'd have to simultaneously believe flying reindeer exist and don't exist. Alternatively to avoid this we could forget everything—but why would the lack of flying reindeer cause us to re-think whether we got a present in our stocking last night?

We know something must be wrong, and at least one of the facts must be incorrect and should be removed, but which? In particular, how can we do this so we get rid of the contradiction, but keep as much knowledge as possible? Belief revision is the study of resolving such issues.

In this work we shall analyse formalisms of belief revision into propositional logic, starting with the very influential Alchourrón, Gärdenfors and Makinson postulates and their relation to rational consequence. We shall then look at several developments that address criticisms of the AGM postulates, particularly in reference to revising multiple times, or iterated revision. Finally we shall present a selection of operators capable of iterated revision, and examine their properties in reference to the postulate sets investigated earlier.

1.1 Preliminaries

Before doing anything, we shall define the language we are working in. Throughout this work, we shall assume a non-empty, finite propositional language L, with the propositional variables being $p, q, r, \ldots$, and SL representing the set of all sentences of L. Since the language is finite, we can define the set of *atoms* At^L, as the set of sentences of the form

$$\bigwedge_{p \in L} \pm p \qquad \text{where } \pm p = p \text{ or } \neg p$$

By the *disjunctive normal form theorem* [9], any sentence θ using any of the standard connectives $\neg, \vee, \wedge, \rightarrow$ is logically equivalent to a sentence of the form

$$\theta \equiv \bigvee_{\alpha \in S_\theta} \alpha \qquad \text{where } S_\theta \subseteq At^L$$

S_θ can be thought of the set of all possible valuations or worlds that satisfy θ, or situations when θ will be true. As proved in [9], we have the following standard properties of S_θ, which we shall use repeatedly.

Theorem 1.2. *For any* $\theta, \phi \in SL$

1. θ *unsatisfiable* $\iff S_\theta = \emptyset$

2. $\theta \models \phi \iff S_\theta \subseteq S_\phi$

3. $S_\theta \cap S_\phi = S_{\theta\wedge\phi}$

4. $S_\theta \cup S_\phi = S_{\theta\vee\phi}$

5. $S_{\neg\theta} = At^L - S_\theta$

1.1.1 Rational Consequence

Along with the classical monotone consequence relation $\models$, we introduce a class of *non-monotonic* rational consequence relations, $\mid\!\sim_{\vec{k}}$, where we interpret $\theta \mid\!\sim_{\vec{k}} \phi$ to mean "if θ then normally ϕ".

Definition 1.3. *Define a k-vector as $\vec{k} = \langle k_1, k_2, k_3, \ldots, k_m \rangle$, such that*

- $k_i \subseteq At^L, \forall i = 1 \ldots m$
- $k_i \cap k_j \neq \emptyset \iff i = j$

For each such $\vec{k}$, we define a binary relation on sentences as follows

Definition 1.4. $\mid\!\sim_{\vec{k}}$ *is a* rational consequence relation, $\mid\!\sim_{\vec{k}} \subseteq SL \times SL$ *iff*

$$\theta \mid\!\sim_{\vec{k}} \phi \iff \begin{array}{l} k_i \cap S_\theta = \emptyset, \forall i = 1, \ldots, m \quad \textit{or} \\ \exists i[k_i \cap S_\theta \neq \emptyset], \textit{ and for the least such } i,\ k_i \cap S_\theta \subseteq S_\phi. \end{array}$$

As a shorthand, also define

$$(\theta)^{\vec{k}} := \begin{cases} \textit{minimum } i \textit{ such that } k_i \cap S_\theta \neq \emptyset & \textit{if } \exists i, k_i \cap S_\theta \neq \emptyset, \\ \infty & \textit{otherwise.} \end{cases}$$

Definition 1.5. $\mid\!\sim_{\vec{k}}$ *is a* consistency-preserving *rational consequence relation, if $\vec{k}$ also satisfies*

$$\bigcup_{i=1}^{m} k_i = At^L$$

Note that since $\bigcup_{i=1}^m k_i = At^L$, $k_i \cap S_\theta = \emptyset, \forall i \Rightarrow S_\theta = \emptyset \Rightarrow \theta$ is inconsistent. So if only considering θ such that θ is consistent, the definition of $\mid\!\sim_{\vec{k}}$ shortens to

$$\theta \mid\!\sim_{\vec{k}} \phi \iff k_i \cap S_\theta \subseteq S_\phi, \textit{ where } i = (\theta)^{\vec{k}}$$

N.B. *When we consider rational consequence relations, we shall* only *need the latter, consistency-preserving rational consequence relations, and unless specifically mentioned, assume that every $\vec{k}$ is such that $\bigcup_{i=1}^{m} k_i = At^L$.*

Although seemingly unrelated, the above definition will soon have a very integral part in our theory revision.

1.2 Knowledge Representation, Revision and Expansion

Before we can reason about such a "store of knowledge" as mentioned at the beginning of this chapter, we need to decide on a representation of such a store. Assuming we can translate the items into a propositional logic language, we shall initially use the following.

Definition 1.6. *A* Knowledge base *K is a set of sentences of L that is deductively closed, i.e.*

$$K \subseteq SL, \qquad K = Cn(K)$$

Where:-

$$Cn(K) = \{\theta : K \models \theta\}$$

A knowledge base K is unsatisfiable if,

$$K \models \theta, \text{ or equivalently, } \theta \in SL \iff K = SL$$

(otherwise K is satisfiable).

Intuitively, a knowledge base is the set of things we would believe at any one moment. It is deductively closed since if we believed p and $p \rightarrow q$, then it seems reasonable to expect us to believe q also, etc. Note that $p \notin K$ does not imply $\neg p \in K$—we don't have to believe p or it's negation, we can be indifferent on the matter also. However, if we believed *both* p and $\neg p$, then $K = SL$ and K is unsatisfiable.

However, even at this stage knowledge bases are not the only way to model our knowledge.

Definition 1.7. *In a finite language situation, an equivalent way of representing current knowledge is a* belief sentence. *Given a knowledge base K, we can generate an equivalent belief sentence ψ, by*

$$\psi := \bigvee \bigcap_{\theta \in K} S_\theta$$

in which case, it is easy to see that $K = Cn(\psi)$ and $\theta \in K \iff \psi \models \theta$.

Belief sentences are easier to handle in some ways since ψ is a single sentence, whereas K is always a non-finite set (since $p \models p \wedge p \models p \wedge p \wedge p \models \ldots$). However, it is easy to see that for each K there are many equivalent sentences ψ, so a belief sentence causes additional complications in checking that equivalent sentences are treated in the same way. Finally we could use a *belief base*, which is purely the set of sentences we have been told, like in example (1.1). However, this will not be relevant to our discussions.

If we weren't worried about keeping K satisfiable, then one obvious way of revising K is by just slinging our new knowledge in. We call this *expansion*, defined by

$$K + \theta = Cn(K \cup \{\theta\})$$

which clearly generates another knowledge base, but there is no guarantee that it is satisfiable.

Example 1.8. $K = Cn(p)$, $K + \neg p = Cn(p, \neg p) = SL$

At this stage the knowledge base—what we believe—contains every sentence possible, which clearly isn't desirable behaviour. It's not reasonable to expect to believe something is both true and false simultaneously.

What we want is a *revision* operator, $K * \theta$, that is guaranteed to produce a satisfiable knowledge base, but naturally this is a lot more complex.

1.3 AGM Postulates

Alchourrón, Gärdenfors and Makinson suggest a series of postulates, or desirable properties, for such an operator [9]. These are, for a knowledge base K and a satisfiable sentence θ,

(*0) $K * \theta$ is satisfiable

(*1) $K * \theta = Cn(K * \theta)$

(*2) $\theta \in K * \theta$

(*3) $\neg\theta \notin K \Rightarrow K * \theta = K + \theta$

(*4) $\theta \equiv \phi \Rightarrow K * \theta = K * \phi$

(*5) For $\theta \wedge \phi$ satisfiable, $\neg\phi \notin K * \theta \Rightarrow (K * \theta) + \phi = K * (\theta \wedge \phi)$

(*0), (*1) specify that the result is another satisfiable knowledge base. (*2) states that θ will always be in our new knowledge base, i.e. the agent always believes what it is told. Of course there are plenty of examples of when we could consider this counter-intuitive—I certainly don't believe all that I'm told—however in the situation that the new sentences are observations, for example, then it seems reasonable. (*3) states that we should do the minimum possible; if there is no reason not to believe something, then we should just add it to what we already know without removing anything from the knowledge base. (*4) states that when given equivalent information, we should end up with the same knowledge base. For example, whether we were told "Henry is in the garden" or "Henry est dans le jardin" shouldn't make any difference to what we believed afterwards[1].

An equivalent way of expressing (*5) is, for $\theta \wedge \phi$ satisfiable, $\neg\phi \notin K * \theta$,

$$\forall\psi, \quad (K * \theta), \phi \models \psi \iff K * (\theta \wedge \phi) \models \psi$$

i.e. if ϕ is a "concrete fact", revising by θ should be the same as revising by $\theta \wedge \phi$—note that in the presence of (*2), the right hand condition is equivalent to $K * (\theta \wedge \phi), \phi \models \psi$

[1]presuming we knew French, of course!

1.3.1 Representations of AGM Revision Operators

For any suggested revision operator, we would currently have to prove each of (*0)–(*5) for our operator. However we already have a representation result from [9].

Theorem 1.9. *$*$ is an AGM-compliant revision operator, i.e. it satisfies (*0)–(*5), iff there exists a consistency preserving rational consequence relation $\mid\!\sim_{\vec{k}}$ such that*

$$\begin{aligned} K &= \{\phi : \ \mid\!\sim_{\vec{k}} \phi\} \quad \textit{or if } k_1 \neq \emptyset,\ K = Cn(\bigvee k_1) \\ K * \theta &= \{\phi : \theta \mid\!\sim_{\vec{k}} \phi\} \end{aligned}$$

Thus not only giving us a framework within which we can base revision operators upon, but also avoids proving that the operator complies with each of (*0)–(*5). Instead, we can just show that an operator is characterised by a $\vec{k}$.

Katsuno and Mendelzon suggest an alternative equivalence, based instead around *faithful orderings*. However, the revision operators that faithful orderings describe use *belief sentences*, as defined in (1.7). Since we have a finite language L, we already know that belief sentences and knowledge bases are equivalent notions, we will now show that the revision operators they define are exactly equivalent to our $*$ operators.

Definition 1.10. *A* total pre-ordering *is a relation $\leq_\psi$ such that:-*

- *$\leq_\psi$ is transitive*
- *$\leq_\psi$ is reflexive*
- *$\leq_\psi$ is total, i.e. $\forall \alpha, \beta \in At^L[\beta \leq_\psi \alpha$ or $\alpha \leq_\psi \beta]$*

As $\leq_\psi$ is a total ordering, we can define $=_\psi$, $<_\psi$ as follows

- *$\alpha =_\psi \beta \iff \alpha \leq_\psi \beta$ and $\beta \leq_\psi \alpha$*
- *$\alpha <_\psi \beta \iff \alpha \leq_\psi \beta$ and $\beta \not\leq_\psi \alpha \iff \beta \not\leq_\psi \alpha$*
  (since $\leq_\psi$ is total, $\beta \not\leq_\psi \alpha \Longrightarrow \alpha \leq_\psi \beta$)

Lemma 1.11. *$<_\psi$ is transitive and irreflexive*

Proof. *Transitive:* Given $\alpha <_\psi \beta, \beta <_\psi \gamma$, assume $\gamma \leq_\psi \alpha$ (i.e. $\alpha \not<_\psi \gamma$).

$$\alpha <_\psi \beta \Rightarrow \alpha \leq_\psi \beta, \text{ with } \gamma \leq_\psi \alpha \Rightarrow \gamma \leq_\psi \beta \Rightarrow \beta \not<_\psi \gamma, \text{ contradiction.}$$

Irreflexive: Assume $\alpha <_\psi \alpha$, then $\alpha \leq_\psi \alpha$ and $\alpha \not\leq_\psi \alpha$, contradiction. □

Definition 1.12. *A* faithful assignment *is a mapping from a sentence ψ to a total pre-order $\leq_\psi$ such that $\forall \alpha, \beta \in At^L$:-*

$$\begin{aligned} \alpha \models \psi, \beta \not\models \psi &\Rightarrow \alpha <_\psi \beta \\ \alpha \models \psi, \beta \models \psi &\Rightarrow \alpha =_\psi \beta \\ \psi \equiv \theta &\Rightarrow \leq_\psi = \leq_\theta \end{aligned}$$

i.e. given ψ, a faithful assignment will order all possible atoms by the "plausibility" of ψ. The atoms in which ψ is true are classified as the smallest by $\leq_\psi$.

Finally, we give their formulation for all AGM revision operators.

Theorem 1.13. *A revision operator $\circ : SL \to SL$ using belief sentences obeys the AGM postulates, i.e. there exists an AGM revision operator $*$ such that*

$$Cn(\psi) * \theta = Cn(\psi \circ \theta)$$

iff there exists a faithful assignment *mapping the belief sentence ψ to a total pre-order $\leq_\psi$ such that*

$$S_{(\psi \circ \theta)} = \{\alpha \in S_\theta : \nexists \beta \in S_\theta [\beta <_\psi \alpha]\}$$

Proof. We shall prove this by showing every operator of this form is equivalent to our already defined set of $*$ operators based on $\vec{k}$.

Lemma 1.14. *Given a faithful assignment from a belief sentence ψ to a total pre-order $\leq_\psi$, we can generate a $\vec{k}$, such that*

$$\alpha \leq_\psi \beta \iff \alpha \in k_i \text{ and } \beta \in k_j, \text{ where } i \leq j \quad (1.14\text{i})$$

$$K = Cn(\psi) = Cn(\bigvee k_1) \quad (1.14\text{ii})$$

Proof. Choose a sequence of $\alpha_i \in At^L$ such that,

$$\alpha_1 <_\psi \alpha_2 <_\psi \ldots <_\psi \alpha_n$$

and it cannot be expanded by any further α_i's. This is possible since At^L is finite, there is a finite number of α_i's, and as $<_\psi$ is irreflexive and transitive, no α_i can be repeated. This gives an upper bound on the length of the sequence. Since At^L is non-empty, we can at least have a sequence of length 1.

Define $\vec{k}$ by

$$k_i := \{\beta \in W : \alpha_i =_\psi \beta\}$$

Clearly (1.14i) will hold, so long as every $\alpha \in At^L$ is a member of some k_i. What if one had been missed out?

Claim. $\bigcup_{i=1}^{m} k_i = At^L$. Suppose not, i.e. $\exists\beta\Big[\forall\alpha_i[\alpha_i \neq_\psi \beta]\Big]$

Since $\leq_\psi$ is total, we have

$$\forall\alpha_i[\alpha_i \leq_\psi \beta \text{ or } \beta \leq_\psi \alpha_i]$$

However, by transitivity, only one of the following will hold:-

- $\forall\alpha_i,\ \alpha_i <_\psi \beta$
- $\forall\alpha_i,\ \beta <_\psi \alpha_i$
- $\exists j$ such that $\alpha_j <_\psi \beta <_\psi \alpha_{j+1}$

In all cases, we could then add β into the sequence. But that sequence was maximal—contradiction.

Since $\leq_\psi$ is a faithful assignment, we know that

$$\alpha \in \{\beta \in At^L : \beta \text{ is minimal according to } \leq_\psi\} \iff \alpha \models \psi \iff \alpha \in S_\psi$$

Which, by the above, is equivalent to

$$\alpha \in k_1 \iff \alpha \in S_\psi$$

So $k_1 = S_\psi$, therefore $Cn(\bigvee k_1) = Cn(\bigvee S_\psi) = Cn(\psi)$, and (1.14ii) holds. □

Theorem 1.15. *Any K&M revision operator $\circ$ as in (1.13) is exactly equivalent to a revision operator $*$ defined with $\mid\!\sim_{\vec{k}}$, i.e. for any satisfiable θ,*

$$Cn(\psi \circ \theta) = Cn(\psi) * \theta$$

Proof. By Lemma (1.14), we can generate a $\vec{k}$ that satisfies the conditions for being a revision operator (i.e. $K = Cn(\psi) = Cn(k_1)$). We show that this $\vec{k}$ is what we want.

$$\begin{aligned}
Cn(\psi) * \theta &= \{\phi : \theta \mid\!\sim_{\vec{k}} \phi\} \\
&= \{\phi : \text{least } i \text{ st } k_i \cap S_\theta \neq \emptyset, k_i \cap S_\theta \subseteq S_\phi\} \\
&\quad\quad \text{since } \theta \text{ is satisfiable, we know } \exists i, k_i \cap S_\theta \neq \emptyset. \\
&= \{\phi : \{\alpha \in S_\theta : \alpha \in k_i \text{ and } \nexists \beta \in S_\theta[\beta \in k_j \text{ and } j < i]\} \subseteq S_\phi\} \\
&= \{\phi : \{\alpha \in S_\theta : \nexists \beta \in S_\theta[\beta <_\psi \alpha]\} \subseteq S_\phi\} \\
&= \{\phi : S_{\psi\circ\theta} \subseteq S_\phi\} = \{\phi : \psi \circ \theta \models \phi\} = Cn(\psi \circ \theta)
\end{aligned}$$

□

A simple corollary of this is that any K&M revision operator $\circ$ satisfies (*0)–(*5).

Lemma 1.16. *Given a $\vec{k}$ for a knowledge base K (i.e. such that $K = \{\phi : \mid\!\sim_{\vec{k}} \phi\}$), we can generate a faithful assignment from any sentence ψ such that $K = Cn(\psi)$ to $\leq_\psi$*

Proof. Define $\leq_\psi$ by, $\forall \alpha, \beta \in At^L$,

$$\alpha \leq_\psi \beta \iff \alpha \in k_i \text{ and } \beta \in k_j, \text{ where } i \leq j$$

Since $K = Cn(\bigvee k_1)$, it is easy to see that the first two conditions of being a faithful assignment hold. Lastly, since $\psi \equiv \theta \Rightarrow Cn(\psi) = Cn(\theta)$, then $\leq_\psi = \leq_\theta$, by definition.

□

And in an identical fashion to Theorem (1.15), we can show that the constructed $\leq_\psi$ will give identical results to a revision operator based on $\vec{k}$, thus we can freely go back and forth between the two definitions.

1.4 Failures of AGM

The aim of the AGM postulates is to capture a set of operators that satisfies our intuitive thinking as to how a revision operator should work, and weed out non-intuitive revision operators. However, the next 2 examples show non-intuitive behaviour that is allowed by the AGM postulates.

Example 1.17. *Consider the following language L, which contains the propositional variables*

c — *I will cycle to work*

w — *I will walk to work*

r — *It is raining*

and suppose we had 2 agents with the initial knowledge of

$$K_1 = Cn(\{r, \neg r \rightarrow c\}) \quad \text{and} \quad K_2 = Cn(\{r, \neg r \rightarrow w\})$$

respectively. Since $r \models \neg r \rightarrow c$ *and* $r \models \neg r \rightarrow w$*, it is easy to see that* $K_1 = K_2 = Cn(\{r\})$

However what happens if we try to revise both by $\neg r$*? We'd expect Agent 1 to decide to cycle to work, agent 2 to walk there. However,* (*4) *dictates that the resulting knowledge base or belief sentence should be the same, since* K_1 *and* K_2 *are the same.*

So AGM-compliant postulates seem to be forced to ignore such "hidden knowledge". Also, the only time that AGM rules mention applying multiple revisions is if we combine (*3) and (*5) to give

$$\theta \wedge \phi \text{ satisfiable}, \neg\theta \notin K, \neg\phi \notin K * \theta \Rightarrow (K * \theta) * \phi = K * (\theta \wedge \phi)$$

which, whilst it seems reasonable, only applies in a very specific case. As a result of this, there's very little we can't do when considering iterated revisions. Consider the following example.

Example 1.18. *(adapted from [3]) Consider the following language L, about this new animal we have just seen*

w — *the animal has wings*

f — *the animal can fly*

Suppose we initially believed $K = Cn(\neg w \wedge \neg f)$. *Then, if it suddenly stretched out it's wings, the following is AGM compliant reasoning*

$$K * w = Cn(w \wedge f) \qquad \textit{possible if } \vec{k} = \langle \{\neg w \wedge \neg f\}, \{\neg w \wedge f\}, \{w \wedge f\}, \{w \wedge \neg f\} \rangle.$$

So not only would we now believe it had wings, but also that it can fly. But what if we'd seen it flying first? Using the above, we get $K * f = Cn(\neg w \wedge f)$, *however it is perfectly AGM compliant to then do*

$$\begin{aligned} K * f &= Cn(\neg w \wedge f) && \textit{using same } \vec{k} \textit{ as above} \\ Cn(\neg w \wedge f) * w &= Cn(w) && \textit{if after revising by } f \textit{ we use a new } \vec{k}\textit{, where} \\ & && \vec{k} = \langle \{\neg w \wedge f\}, \{\neg w \wedge \neg f, w \wedge f, w \wedge \neg f\} \rangle. \end{aligned}$$

$$\textit{Therefore } (K * f) * w = Cn(w).$$

So, even though we saw it flying, once we see it has wings we don't conclude it can fly anymore. This doesn't seem like "reasonable" behaviour—we'd expect a result that stated the animal had wings and could fly.

Chapter 2

Beyond AGM: Further Postulates and Formulations

The previous chapter finished with a selection of examples where AGM compliant operators were producing non-intuitive results. In example (1.17), the problem isn't necessarily related to (*4), but rather that we have destroyed the information that agent 1 prefers cycling, and agent 2 prefers walking. As the example shows, this information cannot be stored in the knowledge base, however it can within the revision operators' $\vec{k}$ as "conditional knowledge". For example, if we revised K_1 and K_2, using

$$\vec{k_1} = \langle \{r \wedge \neg p \wedge \neg q\}, \{\neg r \wedge p \wedge \neg q\}, \{\neg r \wedge \neg p \wedge q\}, \dots \rangle$$

$$\vec{k_2} = \langle \{r \wedge \neg p \wedge \neg q\}, \{\neg r \wedge \neg p \wedge q\}, \{\neg r \wedge p \wedge \neg q\}, \dots \rangle$$

as revision operators respectively, then we would get the intended result. Thus we need to consider the revision operators' $\vec{k}$ as part of our current knowledge, not purely as a device to determine the revision operator.

Example (1.18) works since the AGM rules explicitly avoid mentioning multiple, or iterated, revisions. As a result we can get away with anything, since there are no restrictions to the revision operators we can use, so long as they match with our current knowledge base. Since we now know that there is conditional knowledge within the $\vec{k}$ itself, we need further postulates to restrict the choice of $\vec{k}$ when revising iteratively, to preserve this knowledge.

2.1 The Darwiche and Pearl Approach

Darwiche and Pearl, in solution to the above problems, suggest replacing knowledge bases with *epistemic states* [3]. In their paper they define an epistemic state purposefully vague manner, however there are many papers which have re-expressed the definition in a clearer manner, such as [7]. We will take an epistemic state as any structure such that we have two associated functions:-

- A unary function [ ] that takes an epistemic state and returns a knowledge base—a "current-belief" function.
- A binary function that takes an epistemic state and a sentence, and returns a new epistemic state—i.e. a revision operator.

To give a more concrete example of this, we'll express a collection of epistemic states based on our $\vec{k}$ notation.

Definition 2.1. *Define a $\vec{k}$-based epistemic state as a k-vector as in (1.3), along with two associated functions,*

- *A current-belief function [ ] defined by* $[\vec{k}] = \{\phi : \mid\!\sim_{\vec{k}} \phi\}$
- *A revision operator $\odot$ such that* $\vec{k} \odot \theta = \vec{k}'$

So to complete this definition, we will need to give a revision operator $\odot$. Several different suggestions for revision operators for $\vec{k}$-based epistemic states are given in the next chapter, but for now we will focus on on properties such an operator should have.

Before we used K, (or $\{\phi : \mid\!\sim_{\vec{k}} \phi\}$) as our "current knowledge" and the $\vec{k}$ to define the revision operator. Now we are using *all* of the $\vec{k}$ information as part of our "current knowledge". The revision operator is now a function *applied to* $\vec{k}$'s, as opposed to a function *defined by* them.

Since the notation we have used is very similar to our representation of AGM-compliant operators, we can define AGM compliant operators within $\vec{k}$-based epistemic states by the following.

Lemma 2.2. *Given a $\vec{k}$-based epistemic state such that $\odot$ obeys*

$$[\vec{k} \odot \theta] = \{\phi : \theta \mid\!\sim_{\vec{k}} \phi\}$$

and a fixed $K = [\vec{k}]$, then $K * \theta = [\vec{k} \odot \theta]$ is a valid AGM operator

Proof. Using the same $\vec{k}$, all of the conditions in (1.9) follow by definition. □

Given the above, we can give a set of postulates implying the AGM postulates for $\vec{k}$-based epistemic states as follows:-

($\odot$0) $[\vec{k} \odot \theta]$ is satisfiable

($\odot$1) $[\vec{k} \odot \theta] = Cn([\vec{k} \odot \theta])$

($\odot$2) $\theta \in [\vec{k} \odot \theta]$

($\odot$3) $\neg\theta \notin [\vec{k}] \Rightarrow [\vec{k} \odot \theta] = [\vec{k}] + \theta$

($\odot$4) $\theta \equiv \phi \Rightarrow [\vec{k} \odot \theta] = [\vec{k} \odot \phi]$

($\odot$5) For $\theta \wedge \phi$ satisfiable, $\neg\phi \notin [\vec{k} \odot \theta] \Rightarrow [\vec{k} \odot \theta] + \phi = [\vec{k} \odot (\theta \wedge \phi)]$

So an important point to note is that if our operator $\odot$ complies with AGM postulates, then in the single-revision case $[\vec{k} \odot \theta]$, the result is the same regardless how we define our revision operator.

Also in [3], they suggest 4 extra postulates to further restrict the choice of revision operators.

(C1) $\phi \models \theta \Longrightarrow [(\vec{k} \odot \theta) \odot \phi] = [\vec{k} \odot \phi]$

(C2) $\phi \models \neg\theta \Longrightarrow [(\vec{k} \odot \theta) \odot \phi] = [\vec{k} \odot \phi]$

(C3) $\phi \in [\vec{k} \odot \theta] \Longrightarrow \phi \in [(\vec{k} \odot \phi) \odot \theta]$

(C4) $\neg\phi \notin [\vec{k} \odot \theta] \Longrightarrow \neg\phi \notin [(\vec{k} \odot \phi) \odot \theta]$

(C1) and (C2) are very similar in notation, but the interpretation is quite different. (C1) states that if ϕ implies θ anyway, then revising by θ first wouldn't have an effect, since ϕ implies it anyway. For example, if you were told "the pen is on the table" and then "the pen and pencil are on the table", the first statement is rather redundant—we know that from the second anyway. (C2) States that if revising by θ then ϕ is going to cause a contradiction, then the most recent information completely overrides the older information. This can be thought of as an extension of ($\odot$2), that revision always succeeds. For example, if someone told us "Bill Oddie is my father", then "Actually, my father is Tom Jones", since we believe the second sentence by ($\odot$2), there's no way we can believe the first also—we might as well not been told it.

(C3) specifies we can't do what we did in example (1.18), if we revise by θ and can conclude ϕ (as we do in the first case of (1.18)), then we still can if we knew ϕ before θ (what we expect to happen in the second half). For example, if we conclude that someone is Welsh after hearing their Welsh accent, we should certainly think they're Welsh if they told us they were before we decided they had a Welsh accent. (C4) states that if after θ we didn't refute ϕ, there's no reason to if we did originally know ϕ. For example, concluding that someone hates model railways because their name is Steve doesn't seem reasonable. But if someone told us that they loved model railways and that their name was Steve, then it seems an even more bizarre conclusion.

Do the AGM postulates already imply any of these, though? Example (1.18) already gives us an AGM operator that doesn't obey (C3), but we can also give examples of AGM-compliant operators that also break (C1),(C2) and (C4).

Example 2.3. *AGM operator breaking* (C1) *and* (C2)*: Let*

$$\begin{aligned}\vec{k} &= \langle\{\neg p \wedge q\}, \{p \wedge \neg q\}, \{p \wedge q, \neg p \wedge \neg q\}\rangle \\ \vec{k} \odot p \vee q = \vec{k} \odot \neg p = \vec{k'} &= \langle\{\neg p \wedge q\}, \{p \wedge \neg q, p \wedge q, \neg p \wedge \neg q\}\rangle\end{aligned}$$

and since $[\vec{k'}] = [\vec{k} \odot p \vee q] = \{\neg p \wedge q\} = \{\phi : p \vee q \mid\!\sim_{\vec{k}} \phi\}$*, this is AGM compliant.*

We know $p \models p \vee q$*, however,*

$$\begin{aligned}[\vec{k} \odot p] &= \{\phi : p \mid\!\sim_{\vec{k}} \phi\} &&= Cn(\bigvee\{p \wedge \neg q\}) = Cn(p \wedge \neg q) \\ [(\vec{k} \odot p \vee q) \odot p] = [\vec{k'} \odot p] &= \{\phi : p \mid\!\sim_{\vec{k'}} \phi\} &&= Cn(\bigvee\{p \wedge \neg q, p \wedge q\}) = Cn(p)\end{aligned}$$

which breaks (C1).

Equally $p \models \neg(\neg p)$ *and* $[\vec{k'}] = \{\phi : \neg p \mid\!\sim_{\vec{k}} \phi\}$, *however,*

$$\begin{aligned} [\vec{k} \odot p] &= \{\phi : p \mid\!\sim_{\vec{k}} \phi\} &&= Cn(\bigvee\{p \wedge \neg q\}) = Cn(p \wedge \neg q) \\ [(\vec{k} \odot \neg p) \odot p] = [\vec{k'} \odot p] &= \{\phi : p \mid\!\sim_{\vec{k'}} \phi\} &&= Cn(\bigvee\{p \wedge \neg q, p \wedge q\}) = Cn(p) \end{aligned}$$

which breaks (C2).

AGM operator breaking (C4)*: Let*

$$\begin{aligned} \vec{k} &= \langle\{\neg p \wedge q\}, \{p \wedge \neg q, p \wedge q\}, \{\neg p \wedge \neg q\}\rangle \\ \vec{k} \odot q = \vec{k'} &= \langle\{\neg p \wedge q\}, \{p \wedge \neg q\}, \{p \wedge q, \neg p \wedge \neg q\}\rangle \end{aligned}$$

and since $[\vec{k'}] = [\vec{k} \odot q] = \{\neg p \wedge q\} = \{\phi : q \mid\!\sim_{\vec{k}} \phi\}$, *this is AGM compliant. However,*

$$\begin{aligned} [\vec{k} \odot p] &= Cn(\bigvee\{p \wedge \neg q, p \wedge q\}) &&\Rightarrow \neg q \notin [\vec{k} \odot p] \\ [(\vec{k} \odot q) \odot p] = [\vec{k'} \odot p] &= Cn(\bigvee\{p \wedge \neg q\}) &&\Rightarrow \neg q \in [(\vec{k} \odot q) \odot p] \end{aligned}$$

So since we have an example of AGM-compliant operators that break each rule, (C1)–(C4) can't be derived from AGM postulates. However, they are not completely contradictory either, in the next chapter we shall produce some examples of operators that satisfy all these postulates.

2.2 The Lehmann Approach

An alternative framework is suggested by Lehmann in [8]. Although it was envisaged before Darwiche and Pearl's epistemic states, it can also be thought of as an epistemic state. However, the postulates Lehmann provides have different properties to that of Darwiche and Pearl.

Definition 2.4. *Let* σ, τ, *... represent finite sequences of satisfiable sentences joined by '$\cdot$'. So, for example*

$$\sigma = \theta \cdot \phi \cdot \psi \qquad \sigma \cdot \mu = \theta \cdot \phi \cdot \psi \cdot \mu$$

and define Λ *as the empty sequence*

Definition 2.5. *Let a sequence-based epistemic state be a finite sequence* σ*, where the revision operator is just '$\cdot$', i.e. revision merely causes us to add items to the sequence, and* $[\sigma]$ *be the knowledge base resulting from revising by all the sentences in the sequence.*

Note that since we have defined our sequences as an epistemic state, we can translate the rules from the previous section purely by interchanging symbols. We will also take 2 sequences to be equivalent if

$$\sigma_1 \equiv \sigma_2 \iff [\sigma_1 \cdot \tau] = [\sigma_2 \cdot \tau], \forall \tau$$

i.e. if the result of revising the 2 sequences is the same, now and after any additional revisions.

Lehmann forms a set of postulates similar to the AGM postulates. For finite sequences σ, τ and satisfiable sentences θ, ϕ

(I0) $[\sigma]$ is satisfiable

(I1) $[\sigma] = Cn([\sigma])$

(I2) $\theta \in [\sigma \cdot \theta]$

(I3) $\phi \in [\sigma \cdot \theta] \implies \theta \rightarrow \phi \in [\sigma]$

(I4) $\theta \in [\sigma] \implies [\sigma \cdot \tau] = [\sigma \cdot \theta \cdot \tau]$

(I5) $\phi \models \theta \implies [\sigma \cdot \theta \cdot \phi \cdot \tau] = [\sigma \cdot \phi \cdot \tau]$

(I6) $\neg\phi \notin [\sigma \cdot \theta] \implies [\sigma \cdot \theta \cdot \phi \cdot \tau] = [\sigma \cdot \theta \wedge \phi \cdot \tau]$

(I7) $[\sigma \cdot \neg\theta \cdot \theta] \subseteq [\sigma] + \theta$

(I0), (I1), (I2), are clearly equivalent to postulates ($\odot$0), ($\odot$1), ($\odot$2) in our new language and (I5) is a strengthening of (C1). The remaining AGM postulates can be derived, as shown in the following result due to Lehmann [8].

Theorem 2.6. *Given a revision operator* [ ] *on sequences that satisfies* (I0)–(I7)*,*

- $\theta \equiv \phi \Longrightarrow [\sigma \cdot \theta] = [\sigma \cdot \phi]$ — *i.e* ($\odot$4)

  Proof.

  $$\theta \models \phi \Rightarrow \phi \in [\sigma \cdot \theta] \text{ by (I2),(I1) } \Rightarrow [\sigma \cdot \theta \cdot \tau] = [\sigma \cdot \theta \cdot \phi \cdot \tau] \text{ by (I4)}$$

  But $\phi \models \theta \Rightarrow [\sigma \cdot \theta \cdot \phi \cdot \tau] = [\sigma \cdot \phi \cdot \tau]$ by (I5), so take $\tau = \Lambda$. □

- $\neg\theta \notin [\sigma] \Longrightarrow [\sigma \cdot \theta] = [\sigma] + \theta$ — *i.e.* ($\odot$3)

  Proof. Given such θ, it's enough to prove $[\sigma \cdot \theta] \subseteq [\sigma] + \theta$ and $[\sigma \cdot \theta] \supseteq [\sigma] + \theta$

  Claim. $[\sigma \cdot \theta] \subseteq [\sigma] + \theta$

  Take ϕ such that $\phi \in [\sigma \cdot \theta]$ By (I3), we get $\theta \rightarrow \phi \in [\sigma]$.

  Since the result of expansion (i.e. +) is deductively closed, $\phi \in [\sigma] + \theta$

  Claim. $[\sigma \cdot \theta] \supseteq [\sigma] + \theta$: Since $\theta \in [\sigma \cdot \theta]$ by (I2), enough to show $[\sigma \cdot \theta] \supseteq [\sigma]$.

  Given $\phi \in [\sigma]$,

  $\neg\theta \notin [\sigma] = [\sigma \cdot \phi]$, by (I4). Thus, by (I6), $[\sigma \cdot \phi \cdot \theta] = [\sigma \cdot \phi \wedge \theta]$

  By (I2), $\phi \wedge \theta \in [\sigma \cdot \phi \wedge \theta] \Longrightarrow \phi \in [\sigma \cdot \phi \cdot \theta]$

  But by (I4), $[\sigma \cdot \phi \cdot \theta] = [\sigma \cdot \theta]$. Therefore $\phi \in [\sigma \cdot \theta]$ □

Lehmann also gives the following representation theorem for such revision operators, which we shall not prove.

Theorem 2.7. *Any Lehmann operator* [ ] *on sequences satisfies* (I0)–(I7) *iff we have a* $\vec{k}$ *such that*

$$[\sigma] = Cn(\bigvee Leh(\sigma)_1)$$

where $Leh()$ *maps finite sequences into* $2^{At^L} \times \mathbb{N}$, *and is defined by*

$$Leh(\Lambda) = (k_1, 1)$$

$$Leh(\sigma \cdot \theta) = \begin{cases} (Leh(\sigma)_1 \cap S_\theta, Leh(\sigma)_2) & \textit{if } Leh(\sigma)_1 \cap S_\theta \neq \emptyset \\ ((\bigcup_{i=1}^{n} k_i) \cap S_\theta, n) & \textit{otherwise} \end{cases}$$

$$\textit{Where } n = max\{Leh(\sigma)_2 + 1, (\theta)^{\vec{k}}\}$$

To indicate that we are using a specific $\vec{k}$, *we can write* $[\ ]^{\vec{k}}$.

Intuitively, the $Leh()$ function finds the largest subset of k_1 that satisfies all sentences so far. However if at some point in the recursion this becomes the emptyset we start again, considering a large enough k_i to resolve this.

Using this theorem, we can also prove that the Lehmann Postulates imply (C3) and (C4).

Theorem 2.8. *Any Lehmann operator satisfying* (I0)–(I7) *also satisfies*

$$\phi \in [\sigma \cdot \theta] \Rightarrow \phi \in [\sigma \cdot \phi \cdot \theta] \text{ — } i.e. \text{ (C3)}$$

$$\neg\phi \notin [\sigma \cdot \theta] \Rightarrow \neg\phi \notin [\sigma \cdot \phi \cdot \theta] \text{ — } i.e. \text{ (C4)}$$

Proof of (C3). Any such operator [ ] will be equivalent to the form in (2.7). Assume $\phi \in [\sigma \cdot \theta] = Cn(\bigvee Leh(\sigma \cdot \theta)_1)$.
Assume $S_\theta \cap Leh(\sigma)_1 \neq \emptyset$, so $S_\phi \supseteq Leh(\sigma \cdot \theta)_1$. In this case,

$$\begin{aligned} S_\phi \supseteq Leh(\sigma \cdot \theta)_1 &= Leh(\sigma)_1 \cap S_\theta \\ &= Leh(\sigma)_1 \cap S_\phi \cap S_\theta \\ &= Leh(\sigma \cdot \phi)_1 \cap S_\theta \qquad \text{since } S_\phi \supseteq S_\theta \cap Leh(\sigma)_1 \neq \emptyset \\ &\qquad\qquad\qquad\qquad\qquad \Rightarrow S_\phi \cap Leh(\sigma)_1 \neq \emptyset \\ &= Leh(\sigma \cdot \phi \cdot \theta)_1 \end{aligned}$$

On the other hand, assume $S_\theta \cap Leh(\sigma)_1 = \emptyset \Rightarrow S_\phi \supseteq (\bigcup_{i=1}^n k_i) \cap S_\theta$, $n = Leh(\sigma \cdot \theta)_2$.

If $S_\phi \cap Leh(\sigma)_1 = \emptyset$, $Leh(\sigma \cdot \phi)_1 = (\bigcup_{i=1}^m k_i) \cap S_\phi$, where $m = Leh(\sigma \cdot \phi)_2$. Then either $Leh(\sigma \cdot \phi \cdot \theta)_1 = (\bigcup_{i=1}^m k_i) \cap S_\phi \cap S_\theta \subseteq S_\phi$ and done, or

$$(\bigcup_{i=1}^{m} k_i) \cap S_\phi \cap S_\theta = \emptyset \qquad \text{and} \qquad Leh(\sigma \cdot \phi \cdot \theta)_1 = (\bigcup_{i=1}^{\max\{m+1,n\}} k_i) \cap S_\theta$$

However if $m \geq n$,

$$\begin{aligned} Leh(\sigma \cdot \phi)_1 &= (\textstyle\bigcup_{i=1}^m k_i) \cap S_\phi \\ &\supseteq (\textstyle\bigcup_{i=1}^n k_i) \cap S_\phi \\ &\supseteq (\textstyle\bigcup_{i=1}^n k_i) \cap S_\theta \quad \text{since } S_\phi \supseteq (\textstyle\bigcup_{i=1}^n k_i) \cap S_\theta \end{aligned}$$

but $(\bigcup_{i=1}^m k_i) \cap S_\phi \cap S_\theta = \emptyset$, contradiction. So $Leh(\sigma \cdot \phi \cdot \theta)_1 = (\bigcup_{i=1}^n k_i) \cap S_\theta \subseteq S_\phi$.

Otherwise, if $S_\phi \cap Leh(\sigma)_1 \neq \emptyset$, then $S_\theta \cap S_\phi \cap Leh(\sigma)_1 = S_\theta \cap Leh(\sigma \cdot \phi)_1 = \emptyset$. and $Leh(\sigma)_2 = Leh(\sigma \cdot \phi)_2$, so $Leh(\sigma \cdot \theta) = Leh(\sigma \cdot \phi \cdot \theta)$. □

Proof of (C4). Identical, but instead of checking that $S_\phi \supseteq Leh(\sigma \cdot \phi \cdot \theta)_1$, check that $\exists \alpha \in S_\phi, \alpha \in Leh(\sigma \cdot \phi \cdot \theta)_1$ □

Although it initially looks quite different, the $Leh()$ function is identical to rational consequence in certain situations.

Corollary 2.9. *In the case of a single-sentence sequence, any Lehmann operator* $[\ ]^{\vec{k}}$ *is AGM-compliant, i.e* $*$ *defined by* $K * \theta = [\theta]^{\vec{k}}$, $K = \{\phi : \mid\!\sim_{\vec{k}} \phi\}$ *complies with AGM postulates*

Proof. Consider $Leh(\theta)$. If $k_1 \cap S_\theta \neq \emptyset$, then $Leh(\theta) = (k_1 \cap S_\theta, 1)$. Otherwise,

$$Leh(\theta) = ((\bigcup_{i=1}^{n} k_i) \cap S_\theta, n)), \qquad \text{where } n = max\{2, (\theta)^{\vec{k}}\} = (\theta)^{\vec{k}}, \text{ since } (\theta)^{\vec{k}} > 1$$

$$\implies Leh(\theta) = ((\bigcup_{i=1}^{(\theta)^{\vec{k}}} k_i) \cap S_\theta, n))$$

In both cases, $Leh(\theta)_1 = k_i \cap S_\theta$, for least i such that $k_i \cap S_\theta \neq \emptyset$. So

$$\begin{aligned} K * \theta &= [\theta]^{\vec{k}} = Cn(\bigvee Leh(\theta)_1) \\ &= Cn(\bigvee k_i \cap S_\theta) \qquad \text{for least } i \text{ such that } k_i \cap S_\theta \neq \emptyset \\ &= \{\phi : \theta \mid\!\sim_{\vec{k}} \phi\} \end{aligned}$$ □

A simple observation is that as the number of revisions in σ grows, so will the size of n. Because of this, any Lehmann operator will tend toward *trivial revision*, defined by

$$Leh(\sigma \cdot \theta)_1 = \begin{cases} Leh(\sigma)_1 \cap S_\theta & \text{if } Leh(\sigma)_1 \cap S_\theta \neq \emptyset \\ At^L \cap S_\theta = S_\theta & \text{otherwise} \end{cases}$$

so when contradicting information is found in a sequence, all previous information is discarded. Whilst this seems very undesirable, it's worth noting we only reach this point after receiving enough *inconsistent* sentences, for example the sequence $p \cdot \neg p \cdot \ldots \cdot p \cdot \neg p$. In this case, trivial revision (i.e. purely considering the latest revision as our current knowledge) seems reasonable.

2.3 Differences between D. & P. and Lehmann

A lot of the Lehmann Postulates are stronger than Darwiche and Pearl or AGM. For example, (C1) states that if you revise by a weaker condition then a stronger condition, the first revision is redundant. However, (I5) also states that the weaker condition is redundant after further revisions also (due to the τ appended to the end of the sequence). (I5) in Darwiche and Pearl notation is

$$\phi \models \theta \Longrightarrow (\vec{k} \odot \theta) \odot \phi \equiv \vec{k} \odot \phi$$

Note that in (C1), the knowledge bases were equal after the revision. In this instance, the epistemic states are equivalent, i.e. the knowledge bases will be equal after any further revisions. Lehmann also comments that (I7) is the strongest version of (C2) that was consistent with his rules.

Consider the following case of (I5), using $q \wedge p \models p$

$$[p \wedge q \cdot \neg q] = [p \cdot p \wedge q \cdot \neg q]$$

and if $\neg q \notin [p]$, we can use (I6) to gain

$$[p \wedge q \cdot \neg q] = [p \cdot q \cdot \neg q]$$

So either the right-hand side no longer deduces p after $\neg q$, or the left-hand side doesn't completely discount $p \wedge q$, and still deduces p. Intuitively, this is unlikely to be derivable in the Darwiche and Pearl rules, since we cannot use (C1) with an additional revision after applying the rule (or any of the postulates, in fact). However, we do have (C2), of which three instances are

$$[p \wedge q \cdot \neg q] = [\neg q], \qquad [p \cdot p \wedge q \cdot \neg q] = [p \cdot \neg q], \qquad [p \cdot q \cdot \neg q] = [p \cdot \neg q].$$

So for a Darwiche and Pearl operator, there is clearly going to be a difference between revising by $p \wedge q$ and revising by p then q. On the other hand, Lehmann states that there should be no difference at all. Which is more sensible behaviour is not a clear cut decision, though.

Example 2.10. ***Lehmann Behaviour:*** *Suppose we were talking to a group of people about a party last night.*

Jane tells us....

1. *"Peter came to the party..." — p*
2. *"...and Quincy came too" — q, or combining the 2 sentences, $p \wedge q$*

However, the host later tells us

3. *"Nah, Quincy wasn't here." — $\neg q$*

Whether we view this sequence of sentences as $p \cdot q \cdot \neg q$ or $p \wedge q \cdot \neg q$ is just a matter of how we are translating the English into propositional logic—it shouldn't make a difference as to the end result. Who we believe was at the party shouldn't change simply as a result of how we choose to interpret the situation into logic.

Darwiche and Pearl Behaviour: *Alternatively, imagine each sentence came from a different source.*

Alfred tells us....

1. *"Peter came to the party" — p*

Jane tells us....

2. *"Quincy came to the party" — q*

However, the host later tells us

3. *"Nah, Quincy wasn't here." — $\neg q$*

In this case, clearly we'd dismiss Jane's knowledge, presuming she was got far too drunk and started imagining things, but we'd certainly not discount Peter being there because of it. However if we applied the same reasoning to the Lehmann case above (i.e. $p \wedge q \cdot \neg q$), something must be clouding Jane's judgement, since we know Quincy wasn't there. Equally why should we believe her that Peter was there?

The above example shows two different arguments for each style of revision. Is it the case we could combine the two, though? Formalising the above, we can prove that the 2 sets of postulates are incompatible.

Theorem 2.11. *There is no revision operator that satisfies the Lehmann equivalent of* (C2),

$$\phi \models \neg\theta \implies [\sigma \cdot \theta \cdot \phi] = [\sigma \cdot \phi]$$

and (I0)–(I7).

Proof. Assume we have such an operator, [ ]. Since $\phi \wedge \theta \models \theta$,

$$[\sigma \cdot \theta \cdot \phi \wedge \theta \cdot \neg\phi] = [\sigma \cdot \phi \wedge \theta \cdot \neg\phi] \qquad \text{by (I5)}$$

But since $\neg\phi \models \neg(\phi \wedge \theta)$

$$[\sigma \cdot \phi \wedge \theta \cdot \neg\phi] = [\sigma \cdot \neg\phi] \text{ and } [\sigma \cdot \theta \cdot \phi \wedge \theta \cdot \neg\phi] = [\sigma \cdot \theta \cdot \neg\phi] \qquad \text{by (C2)}$$

So, combining the above, $[\sigma \cdot \theta \cdot \neg\phi] = [\sigma \cdot \neg\phi]$. Since θ and ϕ were arbitrary consistent sentences, any revision sequence $[\sigma \cdot \theta] = [\theta]$

Take $\phi \wedge \psi$ consistent. Then, by (I2), (I1), $\phi, \psi \in [\phi \wedge \psi]$. Since $[\sigma \cdot \theta] = [\theta]$, using (I4),

$$[\phi \wedge \psi] = [\phi \wedge \psi \cdot \phi] = [\phi] = [\psi]$$

But taking θ such that $\theta \wedge \psi$ consistent, $\theta \wedge \phi$ inconsistent,

$$[\theta] = [\psi] = [\phi], \text{ contradiction.}$$ □

However, at the beginning of this section we noted that Lehmann provided more powerful postulates by claiming equality after any additional revisions τ. We can generate a Darwiche and Pearl compatible version of the Lehmann Postulates by removing the τ from (I4)–(I6) and replacing (I7) as follows,

(I4') $\theta \in [\sigma] \implies [\sigma] = [\sigma \cdot \theta]$

(I5') $\phi \models \theta \implies [\sigma \cdot \theta \cdot \phi] = [\sigma \cdot \phi]$

(I6') $\neg\phi \notin [\sigma \cdot \theta] \implies [\sigma \cdot \theta \cdot \phi] = [\sigma \cdot \theta \wedge \phi]$

(I7′) $\phi \models \neg\theta \implies [\sigma \cdot \theta \cdot \phi] = [\sigma \cdot \phi]$

These conditions don't break any of our proofs that a Lehmann operator satisfies the AGM postulates. Also, they imply (C1) and (C2), since they are directly equivalent to (I5′) and (I7′) respectively. In the next chapter, we shall provide an operator that satisfies these conditions, as well as (C3) and (C4), showing that any contradictions have been resolved with this modification.

2.4 Update vs. Revision

As noticed by Katsuno & Mendelzon in [6], what is currently considered as revision should be divided up into 2 cases; *revision*, where new facts are learnt about a static world and *update*, where we are receiving new information about a changing world. Consider the following example adapted from [6].

Example 2.12. *Suppose we have a table with either a book or a magazine on it, but not both—so $L = \{b, m\}$, $K = Cn((b \wedge \neg m) \vee (\neg b \wedge m))$. We send a robot in to put the book on the table, i.e. we want to revise K by b.*

Since $\neg b \notin K$, by (*3)*, $K * b = K + b = Cn(\{(b \wedge \neg m) \vee (\neg b \wedge m), b\}) = Cn(b \wedge \neg m)$ — but why would we now conclude the magazine is on the floor?*

Alternatively, view the same situation as above, but instead of sending a robot in to do the dirty work, we walk in ourselves and see the book on the table. It now seems perfectly reasonable to conclude that the magazine isn't on the table.

The example above isn't necessarily disputing (*3). In the latter part, when we are learning new information about a static world, AGM revision gives a reasonable result. However when the world is dynamic, i.e. when the robot walks in and changes it, AGM revision doesn't give an intuitive answer. Katsuno & Mendelzon suggest using update operators $\dot{*}$ instead. They give a heavily altered version of the AGM postulates for update operators which, for satisfiable θ (as usual), are

(U0) K is satisfiable $\Rightarrow K \dot{*} \theta$ is satisfiable

(U1) $K \dot{*} \theta = Cn(K \dot{*} \theta)$ — same as (*1)

(U2) $\theta \in K \dot{*} \theta$ — same as (*2)

(U3) $\theta \in K \Rightarrow K \dot{*} \theta = K$

(U4) $\theta \equiv \phi \Rightarrow K \dot{*} \theta = K \dot{*} \phi$ — same as (*4)

(U5) $(K \dot{*} \theta) + \phi \subseteq K \dot{*} (\theta \wedge \phi)$

(U6) $\phi \in K \dot{*} \theta, \theta \in K \dot{*} \phi \Longrightarrow K \dot{*} \theta = K \dot{*} \phi$

(U7) For K complete, $(K \dot{*} \theta) \cap (K \dot{*} \phi) \subseteq K \dot{*} (\theta \vee \phi)$
(K is *complete* iff $\forall \theta \in SL$, $\theta \in K$ or $\neg\theta \in K$)

(U8) $(K_1 \cap K_2) \dot{*} \theta = (K_1 \dot{*} \theta) \cap (K_2 \dot{*} \theta)$

Firstly note that (U3) is a weakened version of (*3), so the restriction giving problems in example (2.12) is eased. However, a direct corollary of (U3) is

$$K \text{ is unsatisfiable } \Rightarrow K \dot{*} \theta \text{ is unsatisfiable}$$

i.e. once we have an unsatisfiable knowledge base, any update operator cannot resolve the situation. As a result, (U0) is weakened accordingly. Whilst this sounds very undesirable, the only time we can get K unsatisfiable is if we start off with an unsatisfiable K, due to (U0). (U5) is also a weakening of the original AGM postulate. (U6) states that if we will know one fact if the other happens and vice versa, it doesn't matter which we find out about, our end knowledge is the same. For example, if we see the sun has set, we know it's night-time. If we know it's night-time, then the sun would have set. It doesn't matter which we find out, since our total knowledge on the matter would be the same. (U7) can be seen in the following example. I know there's going to be a stain on the carpet both if I spill coffee on the carpet and if I spill red wine on the carpet—which I actually spill doesn't matter, red wine or coffee, there's going to be a stain on the carpet. (U8) implies that each atom in which K is true is considered separately. Take example (2.12), and consider the following thought process

1. We know that the situation in the room is either $b \wedge \neg m$ or $\neg b \wedge m$

2. If the Robot had found that the situation was $b \wedge \neg m$, there would be no work for it to be done, so it would return, leaving the room untouched.

3. If the Robot had found that the situation was $\neg b \wedge m$, then it would put the book on the table, resulting in $b \wedge m$.

4. So after the update, we know that the situation can be either $b \wedge \neg m$ or $b \wedge m$ — i.e. b.

Katsuno & Mendelzon also provide a representation theorem for their update postulates in [6], as below

Definition 2.13. *Let $S_K = \bigcap_{\theta \in K} S_\theta$, i.e extend our definition of S_θ to knowledge bases.*

Theorem 2.14. *Any update operator $\dot{*}$ obeys* (U0)–(U8) *iff there exists a $\vec{k_\alpha}$ for each $\alpha \in S_K$ such that*

$$K \dot{*} \theta = \bigcap_{\alpha \in S_K} \{\phi : \theta \mid\!\sim_{\vec{k_\alpha}} \phi\} \qquad \text{where } k_{\alpha,1} = \{\alpha\}$$

Or, using our definition of AGM revision from (1.9),

$$K \dot{*} \theta = \bigcap_{\alpha \in S_K} Cn(\alpha) * \theta$$

The effect of the postulate (U8) is clearly visible in the representation theorem, when compared to (1.9). An update operation consists of a separate AGM revision for each possible situation (i.e. atom). For example, in the thought process above we can show that our revisions at steps 2 and 3 are in fact AGM compliant revisions.

Step 2: $b \mid\!\sim_{\vec{k}} b \wedge \neg m$ where $\vec{k} = \langle \{b \wedge \neg m\}, \{\neg b \wedge m\}, \{\neg b \wedge \neg m, b \wedge m\} \rangle$

Step 3: $b \mid\!\sim_{\vec{k}} b \wedge m$ where $\vec{k} = \langle \{\neg b \wedge m\}, \{b \wedge m\}, \{\neg b \wedge \neg m, b \wedge \neg m\} \rangle$

By starting with an example with subtly different wording we have ended up with a completely different semantics, highlighting the importance of clear wording for such examples. Our new representation seems reasonable, however we are still using knowledge bases and have (*4). Problems like example (1.17) are still apparent in this formulation and like AGM, does not consider iteration. If we used an epistemic state formulation like our $\vec{k}$-based epistemic states though, the formulation would clearly be more complex—simply applying the same theory would result in having an epistemic state of $\vec{k}$'s for each possible atom, as opposed to just one.

2.5 Conclusions

Although the AGM postulates have been very influential in shaping further works, as they stood in the previous chapter they are not the panacea of belief revision. Especially when considering the case of iterated revision, which is largely unmentioned. Their rationale for this was one of simplicity—start off with investigating the results of one revision, and then use these results to better equip yourself when dealing with the iterated case.

Largely this has worked. All of the postulate-sets presented above use the AGM postulates as a basis of their reasoning, even if the framework has changed significantly since the knowledge bases used in the first chapter. Amongst the postulate sets presented though, there is no clear winner. For each of the three postulate sets, we have given an example of when one postulate-set gives a more intuitive result than another. As in many cases, it seems that coping with the vagueness of natural language is a tough issue.

Chapter 3

Iterated Revision Operators

We have presented several sets of postulates and representation theorems, but yet to give concrete examples of revision operators. We shall now look at a selection of operators capable of considering iterated revision in an intuitive manner, comparing to postulates from the last chapter.

A recurring theme in the last chapters is that an ordering of knowledge is required to perform iteration, in our case in the form of a $\vec{k}$. This makes intuitive sense—when faced with an example like (1.1), it seems natural to "weigh-up" the facts to decide which you have a firmer belief in, keep firmly regarded facts and disregard other contradictory information. We shall classify operators by how they choose to do this ordering.

3.1 Temporal Operators

Many of the postulates so far, notably (*2), (C1), (C2) and (I2) suggest a heavy preference to newer information over older information. In particular, we always incorporate the newest information into our knowledge, regardless of how much it contradicts information previously in the knowledge base. This idea can be expanded upon to completely define a revision operation–let new information be more believable than old.

3.1.1 σ-Liberation

Richard Booth and others suggest a simple operator based on this thinking in [2] and [1]. Although originally designed for removing sentences from a knowledge base, we will present it as an operator for iterated revision using the Lehmann formulation.

Definition 3.1. *Define σ-Liberation as a revision operator $[\ ]^\Gamma$, where:-*

$$\begin{aligned}
[\sigma]^\Gamma &= Cn(\Gamma(\sigma, \emptyset)) \\
\Gamma(\sigma \cdot \theta, A) &= \begin{cases} \Gamma(\sigma, A) & \text{if } A, \theta \models \neg\theta \text{ i.e. } A \cup \{\theta\} \text{ is unsatisfiable,} \\ \Gamma(\sigma, A \cup \{\theta\}) & \text{otherwise.} \end{cases} \\
\Gamma(\Lambda, A) &= A
\end{aligned}$$

So our function Γ traverses through the list backwards, adding sentences to the list A, as long as the new sentence is consistent with the contents of A.

Example 3.2. *Suppose we had a sequence of revisions $\sigma = q \rightarrow r \cdot \neg q \cdot p \wedge q \cdot p$. Then*

$$\begin{aligned}
[\sigma]^\Gamma &= Cn(\Gamma(q \rightarrow r \cdot \neg q \cdot p \wedge q \cdot p, \emptyset)) \\
&= Cn(\Gamma(q \rightarrow r \cdot \neg q \cdot p \wedge q, \{p\})) && \text{since } p \not\models \neg p \\
&= Cn(\Gamma(q \rightarrow r \cdot \neg q, \{p \wedge q, p\})) && \text{since } p, p \wedge q \not\models \neg(p \wedge q) \\
&= Cn(\Gamma(q \rightarrow r, \{p \wedge q, p\})) && \text{since } p, p \wedge q, \neg q \models q \equiv \neg(\neg q) \\
&= Cn(\Gamma(\Lambda, \{q \rightarrow r, p \wedge q, p\})) && \text{since } p, p \wedge q, q \rightarrow r \not\models \neg(q \rightarrow r) \\
&= Cn(\{q \rightarrow r, p \wedge q, p\}) = Cn(p \wedge q \wedge r)
\end{aligned}$$

So every sentence is kept, apart from $\neg q$. Notice that σ-Liberation has no problem with "hidden knowledge" such as $q \rightarrow r$, which we showed to be a problem for AGM revision operators in example (1.17). This is the rationale for the name, in fact.

Although it is simplest to define this operator on sequence-based epistemic states, it is not in fact compliant with (I5). Consider the following instance

$$[p \cdot q \wedge p \cdot \neg q] = [q \wedge p \cdot \neg q]$$

$$\Gamma(p \cdot q \wedge p \cdot \neg q, \emptyset) = \Gamma(p \cdot q \wedge p, \{\neg q\}) = \Gamma(p, \{\neg q\}) = \Gamma(\Lambda, \{p, \neg q\}) = \{p, \neg q\}$$

However,

$$\Gamma(q \wedge p \cdot \neg q, \emptyset) = \Gamma(q \wedge p, \{\neg q\}) = \Gamma(\Lambda, \{\neg q\}) = \{\neg q\}$$

From our work in section 2.3, this suggests that is actually more likely to obey Darwiche and Pearl postulates, which is what we shall prove below.

Lemma 3.3.

$$A \subseteq \Gamma(\sigma, A) \tag{3.3i}$$

$$\Gamma(\sigma, \emptyset) \not\models \neg\theta \Rightarrow \Gamma(\sigma, \emptyset) \cup \{\theta\} = \Gamma(\sigma, \theta) \tag{3.3ii}$$

Proof of (3.3i). Show by induction on $|\sigma|$ that $A \subseteq \Gamma(\sigma, A)$, for any satisfiable A.

(Λ): $\Gamma(\Lambda, A) = A$, so trivial.

$(\sigma) \Rightarrow (\sigma \cdot \theta)$: Consider $\Gamma(\sigma \cdot \theta, A)$. If $A, \theta \models \neg\theta$, then $\Gamma(\sigma \cdot \theta, A) = \Gamma(\sigma, A) \supseteq A$ by I.H. otherwise $A, \theta \not\models \neg\theta$ i.e. $A \cup \{\theta\}$ is satisfiable, then

$$\Gamma(\sigma \cdot \theta, A) = \Gamma(\sigma, A \cup \{\theta\}) \supseteq A \cup \{\theta\} \supseteq A \qquad \text{by I.H.}$$

Proof of (3.3ii). Assume $\Gamma(\sigma, \emptyset) \not\models \neg\theta$ and take $\sigma = \phi_1 \cdot \ldots \cdot \phi_m$. Any recursion step with $i > 0$ is going to be of the form

$$\begin{aligned}
\Gamma(\phi_1 \cdot \ldots \cdot \phi_i, A_i) &= \begin{cases} \Gamma(\phi_1 \cdot \ldots \cdot \phi_{i-1}, A_i) & \text{if } A_i, \phi_i \models \neg\phi_i \\ \Gamma(\phi_1 \cdot \ldots \cdot \phi_{i-1}, A_i \cup \{\phi_i\}) & \text{otherwise} \end{cases} \\
&= \Gamma(\phi_1 \cdot \ldots \cdot \phi_{i-1}, A_{i-1}) \\
&\qquad \text{where } A_{i-1} = \begin{cases} A_i & \text{if } A_i, \phi_i \models \neg\phi_i \\ A_i \cup \{\phi_i\} & \text{otherwise} \end{cases}
\end{aligned}$$

for unique A_i. Note that $\Gamma(\sigma, \emptyset) = A_0 \supseteq A_1 \supseteq \ldots \supseteq A_m = \emptyset$. Since $\Gamma(\sigma, \emptyset) \not\models \neg\theta$, $A_i \not\models \neg\theta$, $\forall A_i$.

Claim. $\forall A_i$, $i = 1, \ldots, m$ $A_i, \phi_i \models \neg\phi_i \iff A_i, \theta, \phi_i \models \neg\phi_i$

$$A_i, \phi_i \models \neg\phi_i \Rightarrow \theta, A_i, \phi_i \models \neg\phi_i \qquad \text{by monotonicity}$$

$$\begin{aligned}
A_i, \phi_i \not\models \neg\phi_i &\Rightarrow A_{i-1} = A_i \cup \{\phi_i\} && \text{by definition} \\
&\Rightarrow A_{i-1} \cup \{\theta\} \text{ is satisfiable} && \text{since } A_{i-1} \subseteq \Gamma(\sigma, \emptyset) \\
&\Rightarrow \theta, A_i, \phi_i \not\models \neg\phi_i
\end{aligned}$$

Consider the same recursion, but starting with $\Gamma(\sigma, \{\theta\})$.

$$\Gamma(\phi_1 \cdot \ldots \cdot \phi_i, A_i \cup \{\theta\}) = \begin{cases} \Gamma(\phi_1 \cdot \ldots \cdot \phi_{i-1}, A_i \cup \{\theta\}) & \text{if } A_i, \theta, \phi_i \models \neg\phi_i \\ \Gamma(\phi_1 \cdot \ldots \cdot \phi_{i-1}, A_i \cup \{\theta, \phi_i\}) & \text{otherwise} \end{cases}$$

But $A_i, \phi_i \models \neg\phi_i \iff A_i, \theta, \phi_i \models \neg\phi_i$, so

$$= \Gamma(\phi_1 \cdot \ldots \cdot \phi_{i-1}, A_{i-1} \cup \{\theta\})$$

Finally, we show by recursion that $\Gamma(\phi_1 \cdot \ldots \cdot \phi_i, A_i \cup \{\theta\}) = \Gamma(\phi_1 \cdot \ldots \cdot \phi_i, A_i) \cup \{\theta\}$, $i = 0, \ldots, m$

$(i = 0)$: $\Gamma(\Lambda, A_0) \cup \{\theta\} = A_0 \cup \{\theta\} = \Gamma(\Lambda, A_0 \cup \{\theta\})$

$(i = i+1)$: $\Gamma(\phi_1 \cdot \ldots \cdot \phi_{i+1}, A_{i+1} \cup \{\theta\}) = \Gamma(\phi_1 \cdot \ldots \cdot \phi_i, A_i \cup \{\theta\}) = \Gamma(\phi_1 \cdot \ldots \cdot \phi_i, A_i) \cup \{\theta\}$
$= \Gamma(\phi_1 \cdot \ldots \cdot \phi_{i+1}, A_{i+1}) \cup \{\theta\}$, by inductive hypothesis. □

Theorem 3.4. *The σ-liberation operator $[\sigma]^\Gamma$ obeys* (I0)–(I3)*, as well as* (I4')–(I6')

Proof. (I0): Show by induction on $|\sigma|$ that $\Gamma(\sigma, A)$ is satisfiable, for any satisfiable A.

(Λ): $\Gamma(\Lambda, A) = A$, so trivial.

$(\sigma)\Rightarrow(\sigma \cdot \theta)$: $\Gamma(\sigma \cdot \theta, A) = \begin{cases} \Gamma(\sigma, A) & \\ \Gamma(\sigma, A \cup \{\theta\}) & \text{if } A, \theta \not\models \neg\theta, \text{ i.e. } A \cup \{\theta\} \text{ is satisfiable.} \end{cases}$

In both cases, we are done by the inductive hypothesis.

So since $A = \emptyset$ is satisfiable, $\Gamma(\sigma, \emptyset)$ is satisfiable, so $[\sigma]^\Gamma$ is satisfiable.

(I1): Trivial by definition.

(I2): $[\sigma \cdot \theta]^\Gamma = Cn(\Gamma(\sigma \cdot \theta, \emptyset)) = Cn(\Gamma(\sigma, \{\theta\}))$, since $\theta \not\models \neg\theta$ (i.e. θ is satisfiable).

By (3.3i), $\{\theta\} \subseteq \Gamma(\sigma, \{\theta\})$, so $\theta \in \Gamma(\sigma \cdot \theta, \emptyset)$.

(I3): Assume $\phi \in [\sigma \cdot \theta]^\Gamma$, so $\Gamma(\sigma, \{\theta\}) \models \phi$.

If $\Gamma(\sigma, \emptyset) \models \neg\theta \Rightarrow \Gamma(\sigma, \emptyset) \models \theta \rightarrow \phi, \forall\phi \Rightarrow \theta \rightarrow \phi \in [\sigma]$ Otherwise, $\Gamma(\sigma, \emptyset) \not\models \neg\theta \Rightarrow Cn(\Gamma(\sigma, \emptyset), \theta) = Cn(\Gamma(\sigma, \{\theta\}))$ by (3.3ii), so

$$\Gamma(\sigma, \{\theta\}) \models \phi \Rightarrow \Gamma(\sigma, \emptyset), \theta \models \phi \Rightarrow \Gamma(\sigma, \emptyset) \models \theta \rightarrow \phi$$

(I4'): $\theta \in [\sigma]^\Gamma \Rightarrow \Gamma(\sigma, \emptyset) \models \theta \Rightarrow \Gamma(\sigma, \emptyset) \not\models \neg\theta$, so by (3.3ii),

$$Cn(\Gamma(\sigma \cdot \theta, \emptyset)) = Cn(\Gamma(\sigma, \{\theta\})) = Cn(\Gamma(\sigma, \emptyset) \cup \{\theta\}) = Cn(\Gamma(\sigma, \emptyset))$$

$$\text{i.e. } [\sigma \cdot \theta]^\Gamma = [\sigma]^\Gamma$$

(I5'): Assume $\phi \models \theta$. Then

$$\begin{aligned}
\Gamma(\sigma \cdot \theta \cdot \phi, \emptyset) &= \Gamma(\sigma \cdot \theta, \phi) \\
&= \Gamma(\sigma, \{\theta, \phi\}) \\
&= \Gamma(\sigma, \{\theta \wedge \phi\}) && \text{as } \Gamma, \theta, \phi \models \psi \iff \Gamma, \theta \wedge \phi \models \psi \\
&= \Gamma(\sigma, \{\phi\}) = \Gamma(\sigma \cdot \phi, \emptyset) && \text{as } \phi \equiv \theta \wedge \phi
\end{aligned}$$

$$\text{Therefore, } [\sigma \cdot \theta \cdot \phi]^\Gamma = [\sigma \cdot \phi]^\Gamma$$

(I6'): Assume $\neg\phi \notin [\sigma \cdot \theta] \Rightarrow \Gamma(\sigma, \{\theta\}) \not\models \neg\phi$. Then

$$\begin{aligned}
\Gamma(\sigma \cdot \theta \cdot \phi, \emptyset) &= \Gamma(\sigma \cdot \theta, \{\phi\}) \\
& \quad \text{if } \theta, \phi \models \neg\phi \Rightarrow \theta \models \neg\phi \Rightarrow \Gamma(\sigma, \{\theta\}) \models \neg\phi, \text{ contradiction} \\
&= \Gamma(\sigma, \{\theta, \phi\}) \\
&= \Gamma(\sigma, \{\theta \wedge \phi\}) = \Gamma(\sigma \cdot \theta \wedge \phi, \emptyset)
\end{aligned}$$

$$\text{Therefore, } [\sigma \cdot \theta \cdot \phi]^\Gamma = [\sigma \cdot \theta \wedge \phi]^\Gamma$$

(I7'): $\Gamma(\sigma \cdot \theta \cdot \phi, \emptyset) = \Gamma(\sigma \cdot \theta, \{\phi\}) = \Gamma(\sigma, \phi)$, since $\phi, \theta \models \neg\theta$. So $[\sigma \cdot \theta \cdot \phi]^\Gamma = [\sigma \cdot \phi]^\Gamma$. □

Which, by previous results, is enough to show ($\odot$0)–($\odot$5), as well as (C1) and (C2). The remaining postulates are easy to prove.

Theorem 3.5. *The σ-liberation operator $[\sigma]^\Gamma$ obeys*

- $\phi \in [\sigma \cdot \theta]^\Gamma \implies \phi \in [\sigma \cdot \phi \cdot \theta]^\Gamma$ *— i.e.* (C3)
- $\neg\phi \notin [\sigma \cdot \theta]^\Gamma \implies \neg\phi \notin [\sigma \cdot \phi \cdot \theta]^\Gamma$ *— i.e.* (C4)

Proof. (C3): Firstly, $\phi \in [\sigma \cdot \theta]^\Gamma \Rightarrow \Gamma(\sigma, \{\theta\}) \models \phi \Rightarrow \theta \not\models \neg\phi$, since $\theta \in \Gamma(\sigma, \{\theta\})$

$$\Gamma(\sigma \cdot \phi, \{\theta\}) = \Gamma(\sigma, \{\phi, \theta\}) \qquad \text{since } \theta, \phi \not\models \neg\phi$$

By (3.3i), $\{\phi, \theta\} \subseteq \Gamma(\sigma, \{\phi, \theta\}) \Rightarrow \phi \in \Gamma(\sigma, \{\phi, \theta\}) \Rightarrow \phi \in [\sigma \cdot \phi \cdot \theta]^\Gamma$

(C4) is identical, since $\neg\phi \notin [\sigma \cdot \theta]^\Gamma \Rightarrow \Gamma(\sigma, \{\theta\}) \not\models \neg\phi \Rightarrow \theta \not\models \neg\phi$ □

So we have a simple revision operator which fully complies with a set of postulates, why not just stop here? Although it obeys our postulates, it's still an overly-simple model of everyday behaviour. For instance, you may see several offers every day to "Win a free holiday! £1000! A brand new car! No purchase necessary!", but I just dismiss all of them—after all, mother always told me you cannot get something for nothing. Despite the fact that all the adverts are received since I last spoke to my mother, it is my Mother's knowledge I rate highest. Also, in the introduction it was stated that we wanted to get a consistent knowledge base, whilst removing as little information as possible. Consider the following sequence

$$\sigma = p \wedge q \cdot p \wedge r \cdot p \wedge s \cdot p \wedge t \cdot p \wedge u \cdot \neg p \cdot v, \qquad \forall v \in SL, v \neq p, q, r, s, t, u$$

the set of sentences would be consistent simply if $\neg p$ was removed. However σ-Liberation would remove all sentences $p \wedge \delta$ and keep $\neg p$.

3.2 Quantitative Operators

Another method of ranking is to use an explicit value—"I accept θ with degree of plausibility/trustworthiness n". This seems like a reasonable notion in cases when, for example, you have several information sources, you could assign values to each of the information sources in terms of trustworthiness. An observation you make yourself would get "trustworthiness" 100, something your mother told you "trustworthiness" 70, something your mate told you down the pub "trustworthiness" 20 and something written on a promotional flyer just come through your letter box "trustworthiness" 3.

3.2.1 Spohn Conditionalization

This is the motivation for the technique Spohn uses in various papers, for example [10]. Although he expresses it in terms of ranking functions κ, we shall again use a $\vec{k}$-based epistemic state.

Definition 3.6. *Given a $\vec{k}$ such that $k_1 \neq \emptyset$, we say that a* consistent *sentence θ is believed with* degree of firmness *μ, iff*

- $\mid\!\sim_{\vec{k}} \theta$ *and μ minimal such that $k_{(\mu+1)} \cap S_{\neg\theta} \neq \emptyset$,*
- *otherwise* $\not\mid\!\sim_{\vec{k}} \theta$ *and μ minimal such that $k_{(1-\mu)} \cap S_\theta \neq \emptyset$.*

Intuitively, the higher μ is, the stronger we believe in θ. A high positive or negative μ implies we strongly agree or disagree with θ, respectively—not dissimilar to a lecturer evaluation questionnaire.

N.B. *In both the below definitions, if any k_i is undefined, $k_i = \emptyset$. Most importantly, $k_i = \emptyset, i \leq 0$.*

Definition 3.7 (Spohn Conditionalization). *Define a collection of revision operators on $\vec{k}$-based epistemic states, $\vec{k} \oplus_\lambda \theta = \vec{k'}$, using the following algorithm.*

1. *Split $\vec{k}$ into 2 halves, $\vec{k_\theta}$ and $\vec{k_{\neg\theta}}$, so that $\alpha \models \theta\ \forall\alpha \in k_{\theta,i}$, $\alpha \not\models \theta\ \forall\alpha \in k_{\neg\theta,i}$*
2. *$k'_{\theta,i} := k_{\theta,i+\zeta}$, where $\zeta = (\top)^{\vec{k_\theta}} - 1$, or minimal such that $k_{\theta,\zeta+1} \neq \emptyset$—i.e. remove any initial empty $k_{\theta,i}$'s.*
3. *$k'_{\neg\theta,i+\lambda} := k_{\neg\theta,i+\eta}$, where $\eta = (\top)^{\vec{k_{\neg\theta}}} - 1$, or minimal such that $k_{\neg\theta,\eta+1} \neq \emptyset$—i.e. make sure there are λ initial empty $k_{\neg\theta,i}$'s.*
4. *$k'_i := k'_{\neg\theta,i} \cup k'_{\theta,i}$—i.e. combine the 2 new halves.*

Intuitively, the revision adjusts our vector so that θ has *degree of firmness* λ, shifting everything else in a linear fashion around this constraint. For example, figure (3.1) shows each step of a revision of $\vec{k} \oplus_3 \theta$.

For each step, we have a diagram for the relevant $\vec{k}$. The hatched sections represent when a k_i has atoms that are in S_θ or $S_{\neg\theta}$ (in steps 2 and 3, the hatched sections are the non-empty k_i's). After revision, all the "θ atoms" have been moved up, so that the smallest is now in k'_1, whereas all the "$\neg\theta$ atoms" have been moved down, so that there is 3 empty k_i's before them.

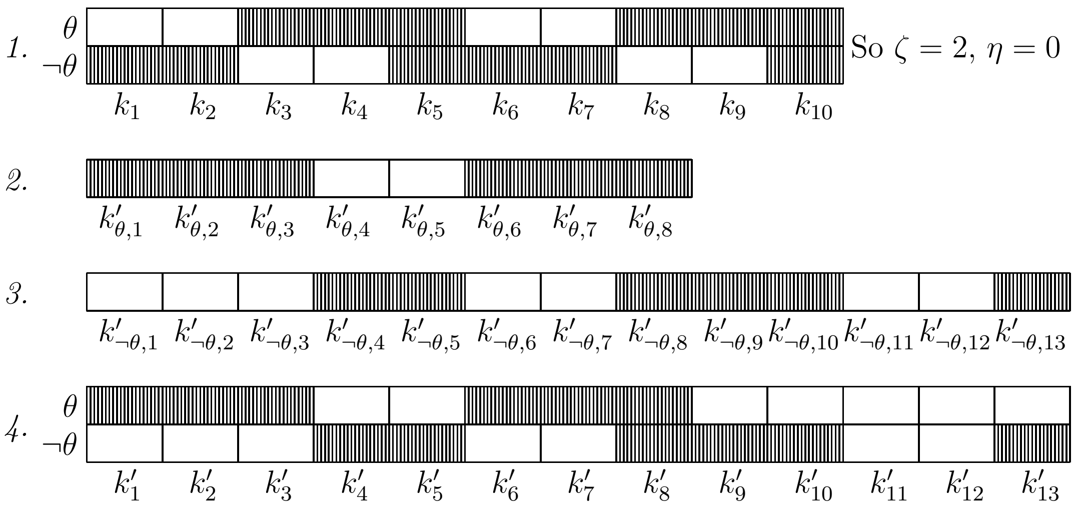

Figure 3.1: $\vec{k} \oplus_3 \theta$

In the extreme cases, if we revise by a tautology, i.e. $\vec{k} \oplus_\lambda \top$, then in the third step, $\bigcup k_{\perp,i} = \emptyset$. So the net effect of the revision is to remove any initial empty k_i's in $\vec{k}$, however if $k_1 \neq \emptyset$, then it has no affect at all.

Alternatively, if we revise by a contradiction, i.e. $\vec{k} \oplus_\lambda \perp$, then $\vec{k_{\neg\perp}} = \vec{k}$ and $\bigcup k_{\perp,i} = \emptyset$. So the revision ensures that the first λ worlds of $\vec{k'}$ are empty. So each has well-defined results, however are not entirely useful—but then, neither is trying to revise by a contradiction or a tautology.

Postulate Conformance

Firstly, the conditionalization operators are richer than the operators which Darwiche and Pearl describe in their postulates; we have a whole range of revision operators to consider, as opposed to just one. Particularly when considering iterated revision, there is a choice of applying the postulate to a fixed λ, or allowing variation through the iteration.

Notice that, although the algorithm will result in a valid $\vec{k}$, it is not guaranteed that all the k_i's in the result will be non-empty. For example, in figure (3.1), k'_{11} and k'_{12} are empty. Since we are trying to specify the "degree of firmness" for a θ, the absolute position of worlds is important, not just their relative position.

Theorem 3.8. *Fix $\lambda > 0$. A Spohn conditionalization operator $\oplus_\lambda$, meets postulates*

($\odot$0)–($\odot$5)

Proof. ($\odot$0)–($\odot$5): Let $\vec{k}' = \vec{k} \oplus_\lambda \theta$, θ consistent.

$$
\begin{aligned}
[\vec{k} \oplus_\lambda \theta] &= [\vec{k}'] \\
&= Cn(\bigvee(k_{\neg\theta,1+\eta-\lambda} \cup k_{\theta,1+\zeta})) && \text{where } \eta, \zeta \text{ min. st } k_{\neg\theta,\eta+1} \neq \emptyset, k_{\theta,\zeta+1} \neq \emptyset \\
&= Cn(\bigvee(\emptyset \cup k_{\theta,1+\zeta})) && \text{since } 1+\eta-\lambda \leq \eta, \\
&= Cn(\bigvee(k_{\theta,i})) && \text{where } i \text{ minimal st. } k_{\theta,i} \neq \emptyset, \\
&= Cn(\bigvee(k_i \cap S_\theta)) && \text{where } i \text{ minimal st. } k_i \cap S_\theta \neq \emptyset, \\
&= \{\phi : \theta \mid\!\sim_{\vec{k}} \phi\}
\end{aligned}
$$

Therefore $\oplus_\lambda$ obeys ($\odot$0)–($\odot$5), by lemma (2.2). □

Corollary 3.9. *For a single-revision case $[\vec{k} \oplus_\lambda \theta]$, the choice of $\lambda > 0$ is irrelevant.*

Proof. From the proof of above, $[\vec{k} \oplus_\lambda \theta] = \{\phi : \theta \mid\!\sim_{\vec{k}} \phi\}$, which doesn't mention λ. □

Note the stipulation that $\lambda > 0$. Clearly if $\lambda = 0$, then we may end up with θ and $\neg\theta$ atoms in k_1, therefore $\theta \notin [\vec{k} \oplus_0 \theta]$ and ($\odot$2) is broken. However, this condition is enough get $\oplus_\lambda$ operators that comply with (C1) and (C2).

Theorem 3.10. *Fix $\lambda, \mu > 0$. Spohn conditionalization operators obey* (C1) *and* (C2)

Proof of (C1). Consider $\vec{k'} = \vec{k} \oplus_\lambda \theta$, and use same notation as in definition of $\oplus$.

$$[(\vec{k} \oplus_\lambda \theta) \oplus_\mu \phi] = Cn(\bigvee(k'_i \cap S_\phi)) \qquad \text{where i min st } k'_i \cap S_\phi \neq \emptyset$$

since $\phi \models \theta, S_\phi \subseteq S_\theta \Longrightarrow S_\phi \subseteq \bigcup_{j=1}^{m} k'_{\theta,j}$ therefore

$$= Cn(\bigvee(k'_{\theta,i} \cap S_\phi)) \qquad \text{where i min st } k'_{\theta,i} \cap S_\phi \neq \emptyset$$

since $k'_{\theta,i} := k_{\theta,i+\zeta}$ for fixed ζ, the position of any world relative to another is unchanged, so

$$= Cn(\bigvee(k_{\theta,j} \cap S_\phi)) \qquad \text{where j min st } k_{\theta,j} \cap S_\phi \neq \emptyset$$

again, since $S_\phi \subseteq S_\theta, k_{\neg\theta,i} \cap S_\theta = \emptyset, \forall i$

$$= Cn(\bigvee(k_j \cap S_\phi)) \qquad \text{where j min st } k_j \cap S_\phi \neq \emptyset$$
$$= [\vec{k} \oplus_\mu \phi]$$

□

Proof of (C2). Identical to (C1), but using $k_{\neg\theta}$ in place of k_θ.

However to gain (C3) and (C4) for Spohn conditionalization, we have to require that in revising by ϕ, we don't decrease the degree of firmness in ϕ. Although possible by setting $\lambda < (\neg\phi)^{\vec{k}}$, it isn't very intuitive anyway. If we already believed ϕ (to a degree), if someone told us ϕ again, it shouldn't result in our belief in ϕ decreasing, regardless how untrustworthy they were.

Theorem 3.11. *With $\lambda > 0$, If $\lambda \geq (\neg\phi)^{\vec{k}}$, then*

$$\phi \in [\vec{k} \oplus_\lambda \theta] \Rightarrow \phi \in [(\vec{k} \oplus_\lambda \phi) \oplus_\lambda \theta] \qquad \text{— i.e. (C3)}$$

$$\neg\phi \notin [\vec{k} \oplus_\lambda \theta] \Rightarrow \neg\phi \notin [(\vec{k} \oplus_\lambda \phi) \oplus_\lambda \theta] \qquad \text{— i.e. (C4)}$$

Proof of (C3). Let $\vec{k'} = \vec{k} \oplus_\lambda \phi$, and use $\vec{k_\phi}$, $\vec{k_{\neg\phi}}$ as in definition of $\vec{k} \oplus_\lambda \phi$.

$$\begin{array}{rcll}
\phi \in [\vec{k} \oplus_\lambda \theta] & \Rightarrow & \theta \mid\!\sim_{\vec{k}} \phi & \text{by (3.8)} \\
& \Rightarrow & k_i \cap S_\theta \subseteq S_\phi & \text{where } i \text{ min st } k_i \cap S_\theta \neq \emptyset \\
k_i \cap S_\theta & = & k_{\phi,i} \cap S_\theta & \text{where } i \text{ min st } k_{\phi,i} \cap S_\theta \neq \emptyset
\end{array}$$

Since $k'_{\phi,i} := k_{\phi,i+\zeta}$, for fixed ζ, the position of any world relative to another is unchanged, so

$$\begin{array}{rcll}
k_i \cap S_\theta & = & k'_{\phi,j} \cap S_\theta & \text{where } j \text{ min st } k'_{\phi,j} \cap S_\theta \neq \emptyset, j \leq i
\end{array}$$

Since $\lambda \geq (\neg\phi)^{\vec{k}}$, then $k'_{\neg\phi,p} = k_{\neg\phi,q}$, $p \geq q$. So $k'_{\neg\phi,j}$ must be empty, giving

$$\begin{array}{rcll}
k_i \cap S_\theta & = & (k'_{\neg\phi,j} \cup k'_{\phi,j}) \cap S_\theta & \text{where } j \text{ min st } k'_{\phi,j} \cap S_\theta \neq \emptyset, j \leq i \\
& = & k'_j \cap S_\theta & \text{where } j \text{ min st } k'_j \cap S_\theta \neq \emptyset, j \leq i
\end{array}$$

$$\Rightarrow \theta \mid\!\sim_{\vec{k'}} \phi \Rightarrow \phi \in [(\vec{k} \oplus_\lambda \phi) \oplus_\lambda \theta]$$

□

Proof of (C4). Identical, but instead of checking that $k'_j \cap S_\theta \subseteq S_\phi$, check that $\exists \alpha \in S_\phi, \alpha \in k'_j \cap S_\theta$

Although the Spohn conditionalization operators satisfy our postulates, intuitively the quantification doesn't quite match with everyday behaviour. Whilst phraseology such as we accept x "to a degree" is part of common language, we would never specify what that degree was, certainly not numerically. If we were implementing a system based on Spohn conditionalization though, this would not necessarily be a problem—we could devise a scheme by experimentation and comparing the results against what we intuitively expect. However, such a system would not strictly be reasoning independently.

3.2.2 Relative Conditionalization

A subtly different way of using this notation is to take the revision as "increase our belief in θ by λ". This is defined below.

Definition 3.12 (Relative Conditionalization). *Define a collection of revision* *operators on $\vec{k}$-based epistemic states, $\vec{k} \ominus_\lambda \theta = \vec{k'}$, as follows*

$$k'_i = k_{\theta,\eta+i+\lambda} \cup k_{\neg\theta,\eta+i}$$

Where η minimal such that $k_{\theta,\eta+\lambda} \cup k_{\neg\theta,\eta} \neq \emptyset$, k_θ defined as before.

To show the difference between the operators, figure (3.2) represents a series of revisions $\vec{k'} = \vec{k} \ominus_\lambda \theta$, $\lambda = 1, 2, 3$.

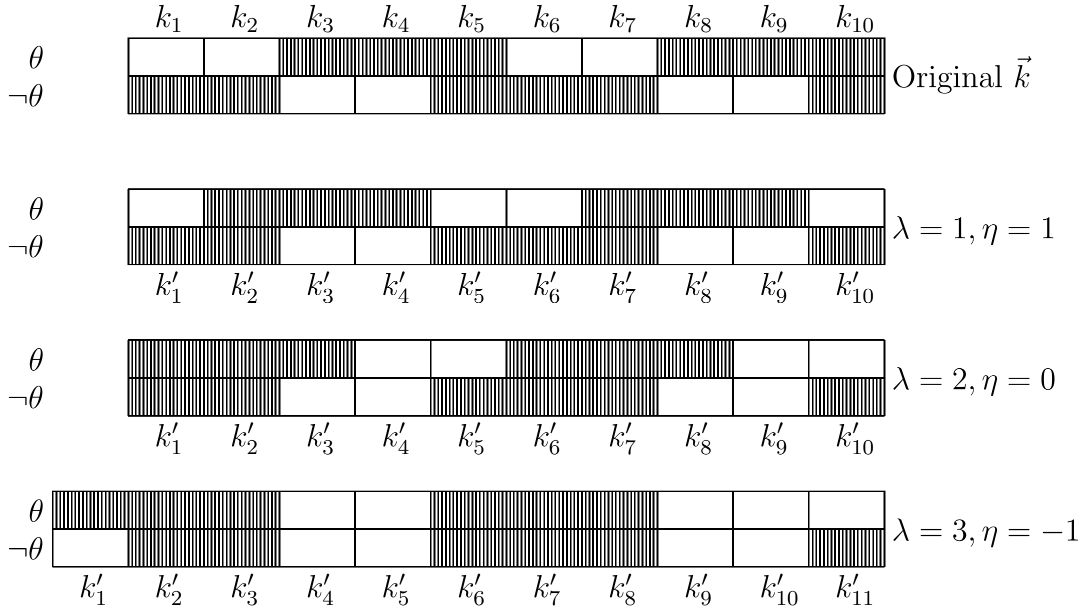

Figure 3.2: $\vec{k'} = \vec{k} \ominus_\lambda \theta$, $\lambda = 1, 2, 3$

θ atoms move up by λ up until the last revision, where since there is no k_0 to move to, the net effect is that the $\neg\theta$ worlds are moved down by $\lambda - (\theta)^{\vec{k}}$.

Clearly the relative conditionalization operators are not compliant with any complete set of postulates so far, since in the example in figure (3.2), $\theta \notin [\vec{k} \ominus_1 \theta]$, which breaks postulates ($\odot$2), (I2) and (U1), so the class of operators does not fit into any of our postulate sets, however this doesn't mean the notion is completely preposterous.

Example 3.13. *Consider the following sequence of revisions $\vec{k} \ominus_1 \theta \ominus_1 \theta \ominus_1 \theta$ equivalent to the revisions in figure (3.2) and where θ means "The moon is made of cheese".*

The first person who tells you that it's made of cheese, you're likely to dismiss them as being stupid, but another person says he heard it being announced this morning *as well. After these 2 revisions you are now unsure, and go to check the BBC news website, and again, it says the moon has just been discovered to be made out of cheese. After this third revision, you now believe the fact.*

3.3 Comparative Operators

Another way to rank our belief in sentences is to say "I believe θ as much as I believe ϕ". For example, "I'll believe that when pigs fly". Unlike using arbitary ordinals, this does resemble reasoning we might use on an everyday basis. We could also use it in the "trustworthiness" example at the beginning of the last section, by believing everything we hear down the pub as much as some reference sentence ϕ, for example.

3.3.1 Fermé and Rott's Revision by Comparison

Fermé and Rott describes a method of reasoning in this manner, where we revise θ to be "at least as true as ψ" [4]. Again, we shall express his operator in terms of a $\vec{k}$-based epistemic state.

Definition 3.14 (Revision by Comparison). *Given a $\vec{k}$, define a collection of revision operators $\otimes_\psi$, $\psi \in SL$, ψ not a tautology, as*

$$\vec{k} \otimes_\psi \theta = \vec{k}''',$$

the result of the revision being that θ has the same degree of firmness as ψ, where:-

$$a = (\neg\psi)^{\vec{k}}$$

$$k'_i = \begin{cases} k_i \cap S_\theta & \text{if } i < a \\ k_i & \text{otherwise} \end{cases}$$

$$k''_i = \begin{cases} k'_i \cup \bigcup_{j=1}^{a}(k_i \cap S_{\neg\theta}) & \text{if } i = a \\ k'_i & \text{otherwise} \end{cases}$$

$$k'''_i = k''_{b+i-1} \qquad \text{where } b \text{ is minimal such that } k''_b \neq \emptyset$$

So $\vec{k'}$ has all the $S_{\neg\theta}$ "below" $\neg\psi$ removed. $\vec{k''}$ re-inserts them at the minimal k''_i such that $\neg\psi$ holds. Finally, $\vec{k'''}$ removes any initial empty k_i's that we may have formed.

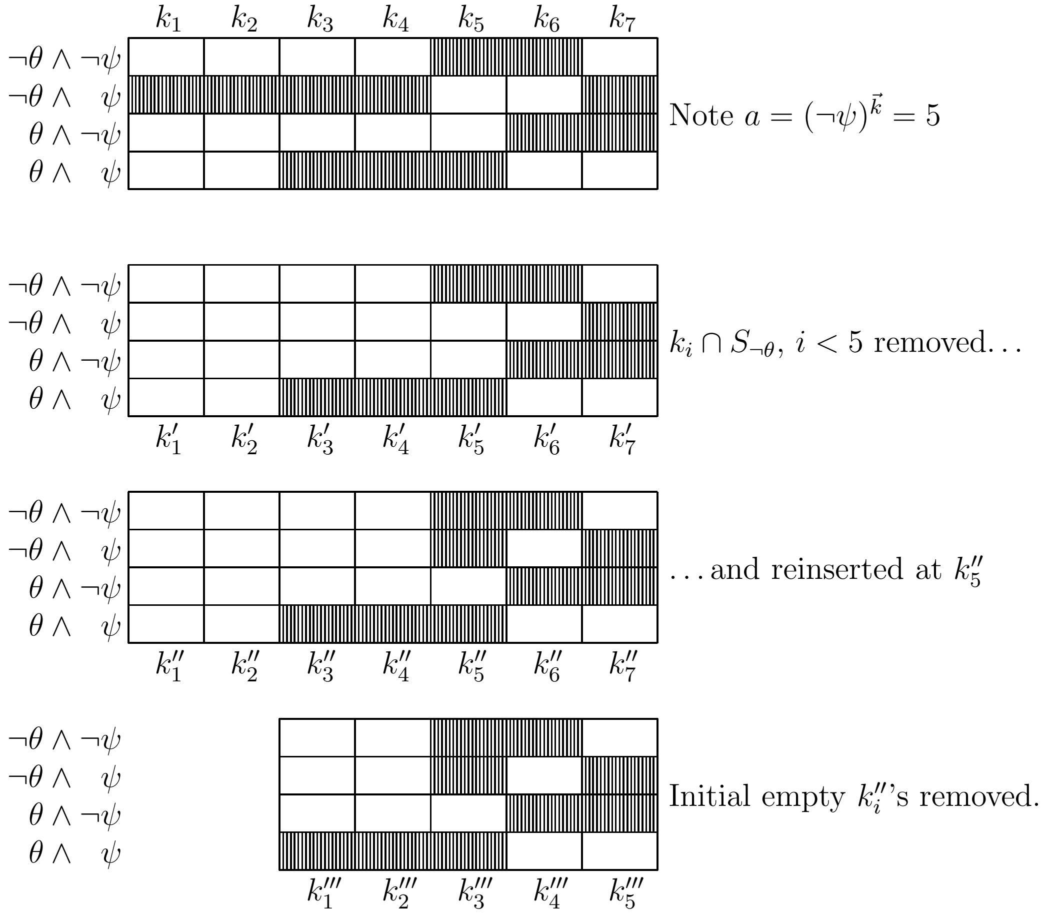

Figure 3.3: $\vec{k'''} = \vec{k} \otimes_\psi \theta$, where $a = 5$

Figure (3.3) shows a revision of $\vec{k} \otimes_\psi \theta$, step by step. The end result is atoms in $k_1, \ldots, k_4 \cap S_{\neg\theta}$ have been compacted down to k'''_3, i.e. the first k'''_i that contains an atom $\in S_{\neg\psi}$. The first thing to notice is that in the process, information has been lost—the preferences between the worlds $S_{\neg\theta\wedge\psi} \cap k_{1\ldots4}$ isn't present in the new $\vec{k'''}$. Given a series of revisions, using the revision by comparison will tend toward having no preference at all. Note this implies that we can't form a reasonable set of beliefs from "total ignorance" (i.e. $k_1 = At^L$), since the result of any revision by comparison on such a $\vec{k}$ will be $\vec{k}$ again.

Whilst the above is the intended result, it's obviously it's not going to work in all cases either. We are only getting $\theta \in [\vec{k} \otimes_\psi \theta]$ in the above case since $(\theta)^{\vec{k}} < (\neg\psi)^{\vec{k}}$. Consider instead figure (3.4).

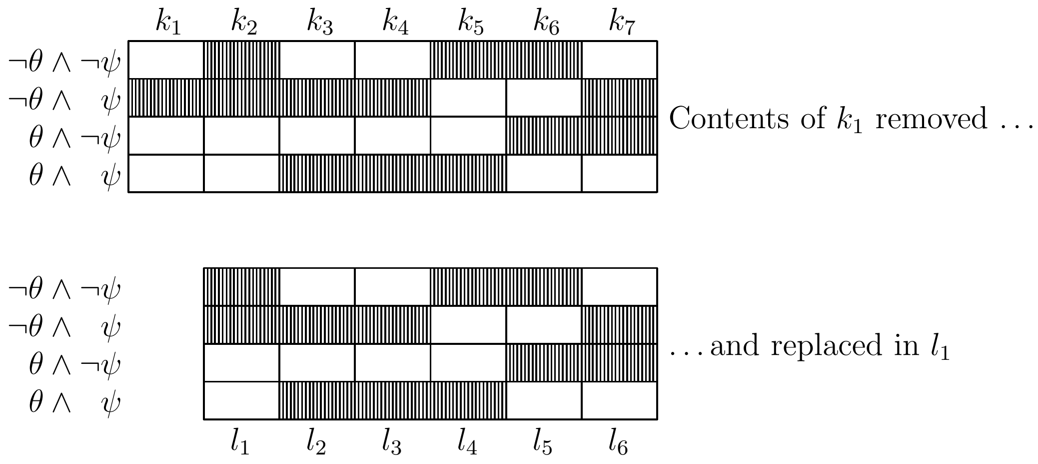

Figure 3.4: $\vec{l} = \vec{k} \otimes_\psi \theta$ when $(\neg\psi)^{\vec{k}} \leq (\theta)^{\vec{k}}$

In the successful case of figure (3.3), $[\vec{k} \otimes_\psi \theta] = \theta \wedge \psi$. In figure (3.4), $[\vec{k} \otimes_\psi \theta] = \neg\theta$— not only do we believe the inverse of θ, we also no longer believe ψ holds. Our revision has failed. We've done what the algorithm intends to do, that is change our degree of firmness in $\neg\theta$ to match $\neg\psi$, however, this was not a strong enough condition to cause us to believe θ as a result.

Postulate compatibility

Theorem 3.15. *Any revision by comparison operator, $\otimes_\psi$, obeys* ($\odot$0)–($\odot$5) *just when $(\theta)^{\vec{k}} < (\neg\psi)^{\vec{k}}$, i.e. ψ restricted to successful cases.*

Proof. If $(\theta)^{\vec{k}} \not< (\neg\psi)^{\vec{k}}$, figure (3.4) gives us a counter example to ($\odot$2), so assume we have $\vec{k'''} = \vec{k} \otimes_\psi \theta$ such that $(\theta)^{\vec{k}} < (\neg\psi)^{\vec{k}}$.

Using the same notation as in Definition (3.14), we have $k''_j = k_j \cap S_\theta, \forall j < (\neg\psi)^{\vec{k}}$. Hence,

$$k''_j \subseteq S_\theta, \qquad \forall j < (\neg\psi)^{\vec{k}}$$

$$\{\phi : \theta \mid\!\sim_{\vec{k}} \phi\} = Cn(k_i \cap S_\theta) = Cn(k''_i) \qquad \text{where } i \text{ is minimal st. } k_i \cap S_\theta \neq \emptyset$$

But, by the above, i is also minimal such that $k''_i \neq \emptyset$, therefore

$$Cn(k''_i) = Cn(k'''_1) = [\vec{k'''}].$$

□

So $\otimes$ only obeys AGM postulates so long as we only consider "successful" revisions. This condition is enough to give us some of the Darwiche & Pearl conditions also.

Lemma 3.16. *For a revision by comparison operator* $\vec{k'''} = \vec{k} \otimes_\psi \theta$*,*

$$\alpha \in S_\theta, \alpha \in k''_i \Rightarrow \alpha \in k_i \qquad \textit{and} \qquad \alpha \in S_{\neg\theta}, \alpha \in k''_j, \alpha \in k_i \Rightarrow j \geq i$$

Proof. By examining the algorithm, it is trivial to see that

$$\alpha \in S_\theta \text{ or } (\alpha)^{\vec{k}} \geq (\neg\psi)^{\vec{k}}, \alpha \in k_i \Rightarrow \alpha \in k''_i$$

The only case when an atom is moved is when $\alpha \in S_{\neg\theta}$ and $(\alpha)^{\vec{k}} < (\neg\psi)^{\vec{k}}$, in which case

$$\alpha \in k_i \Rightarrow \alpha \in k''_j \qquad \text{where } j = (\neg\psi)^{\vec{k}}, \text{ therefore } j \geq i = (\alpha)^{\vec{k}}. \qquad \square$$

Theorem 3.17. *For any AGM revision by comparison operator, i.e.* $\otimes_\psi$ *such that any revision of* $\vec{k} \otimes_\psi \theta, \phi$ *is successful,*

$$\phi \models \theta \Rightarrow [(\vec{k} \otimes_\psi \theta) \otimes_\psi \phi] = [\vec{k} \otimes_\psi \phi] \qquad \textit{— i.e.}\ \mathsf{(C1)}$$

Proof. Let $\vec{k'''} = \vec{k} \otimes_\psi \theta$, and $\vec{k'}$, $\vec{k''}$ be the intermediary steps in this revision.

$$\begin{aligned} [(\vec{k} \otimes_\psi \theta) \otimes_\psi \phi] &= \{\varphi : \phi \mid\!\sim_{\vec{k'''}} \varphi\} \quad \text{By (3.15)} \\ &= Cn(\bigvee k'''_i \cap S_\phi) \quad i \text{ minimal such that } k'''_i \cap S_\phi \neq \emptyset \end{aligned}$$

Since $k'''_i = k_b + i - 1''$, for fixed b, the position of any world relative to another is unchanged, so

$$= Cn(\bigvee k''_i \cap S_\phi) \qquad i \text{ minimal such that } k''_i \cap S_\phi \neq \emptyset$$

However since $\phi \models \theta$, any $\alpha \in S_\phi, \alpha \in S_\theta$, so $k''_i \cap S_\phi = k_i \cap S_\phi$ by (3.16), i.e. the position of any world in S_ϕ remains unaltered between k''_i and k_i.

$$= Cn(\bigvee k_i \cap S_\phi) \qquad i \text{ minimal such that } k_i \cap S_\phi \neq \emptyset$$

$$= \{\varphi : \phi \mid\!\sim_{\vec{k}} \varphi\} = [\vec{k} \otimes_\psi \phi] \qquad \square$$

Note that throughout the proof, the only time we use the value of ϕ is so that revision by ϕ is successful, i.e. $(\phi)^{\vec{k}} < (\neg\psi)^{\vec{k}}$. Because of this, the postulate still holds if we vary ψ in the revisions used in (C1).

However, unless we impose severe restrictions, (C2) cannot hold—as noted before, the ordering between $\neg\theta$ worlds is lost as part of the revision, a byproduct of this is in general, (C2) will fail.

Theorem 3.18. *For any AGM revision by comparison operator, i.e.* $\otimes_\psi$ *such that a revision of* $\vec{k} \otimes_\psi \phi$ *is successful,*

$$\phi \in [\vec{k} \otimes_\psi \theta] \Rightarrow \phi \in [(\vec{k} \otimes_\psi \phi) \otimes_\psi \theta] \qquad \text{— i.e. (C3)}$$

$$\neg\phi \notin [\vec{k} \otimes_\psi \theta] \Rightarrow \neg\phi \notin [(\vec{k} \otimes_\psi \phi) \otimes_\psi \theta] \qquad \text{— i.e. (C4)}$$

Proof of (C3). Let $\vec{k'''} = \vec{k} \otimes_\psi \phi$, and $\vec{k''}$ and $\vec{k'}$ be the intermediary steps in this revision.

$$\begin{aligned} \phi \in [\vec{k} \otimes_\varphi \theta] \;&\Rightarrow\; \theta \mid\!\sim_{\vec{k}} \phi && \text{by (3.15)} \\ &\Rightarrow\; k_i \cap S_\theta \subseteq S_\phi && \text{where i min s.t. } k_i \cap S_\theta \neq \emptyset \end{aligned}$$

But using (3.16), $k_i \cap S_\theta = k''_i \cap S_\theta$ (since they are a subset of S_ϕ). But is i still minimal?

Again, using (3.16), any j such that $k_j \cap S_\theta \subseteq S_{\neg\phi}$, $k''_l \cap S_\theta = k_j \cap S_\theta, l \geq j$. So since none of these were minimal in $\vec{k}$, they certainly aren't in $\vec{k''}$. So,

$$= k''_i \cap S_\theta \subseteq S_\phi \qquad \text{where } i \text{ min s.t. } k''_i \cap S_\theta \neq \emptyset$$

Since $k'''_i = k''_{b+i-1}$, for fixed b, the position of any world relative to another is unchanged, so

$$= k'''_j \cap S_\theta \subseteq S_\phi \qquad \text{where } j \text{ min s.t. } k'''_j \cap S_\theta \neq \emptyset$$

$$\Rightarrow \theta \mid\!\sim_{\vec{k'''}} \phi \Rightarrow \phi \in [(\vec{k} \otimes_\psi \phi) \otimes_\varphi \theta] \qquad \text{by (3.15)}$$

□

Proof of (C4). Identical, but instead of checking that $k'''_j \cap S_\theta \subseteq S_\phi$, check that $\exists\alpha \in S_\phi, \alpha \in k'''_j \cap S_\theta$ □

Again as in the proof of (C1), notice that we don't require a fixed value for ψ, but merely that the revision is successful.

Postulate-wise, revision by comparison hasn't done as well as Spohn conditionalization. We only get AGM-compliant revision in the relatively specific "successful" case, and only (C2) in trivial cases (i.e. when the revision has no effect). Worst of all, after enough revisions, $\vec{k}$ tends to $k_1 = At^L$.

3.3.2 Comparative Conditionalization

Instead, what happens when we try and use the Spohn conditionalization as a comparative operator to get similar effects?

Definition 3.19. *Define* comparative conditionalization *as follows*

$$\vec{k} \oplus_\psi \theta = \vec{k} \oplus_\lambda \theta \vee \neg\psi \qquad \textit{where } \lambda = (\neg\psi)^{\vec{k}} - (\theta \vee \neg\psi)^{\vec{k}}$$

So, instead of combining $\neg\theta \wedge \psi$-atoms, they are shifted back. Because of this, $\vec{k}$ will no longer tend to $k_1 = At^L$ as they did in revision by comparison. As an example, consider figure (3.5), which performs the same revision as figure(3.3) but using comparative conditionalization.

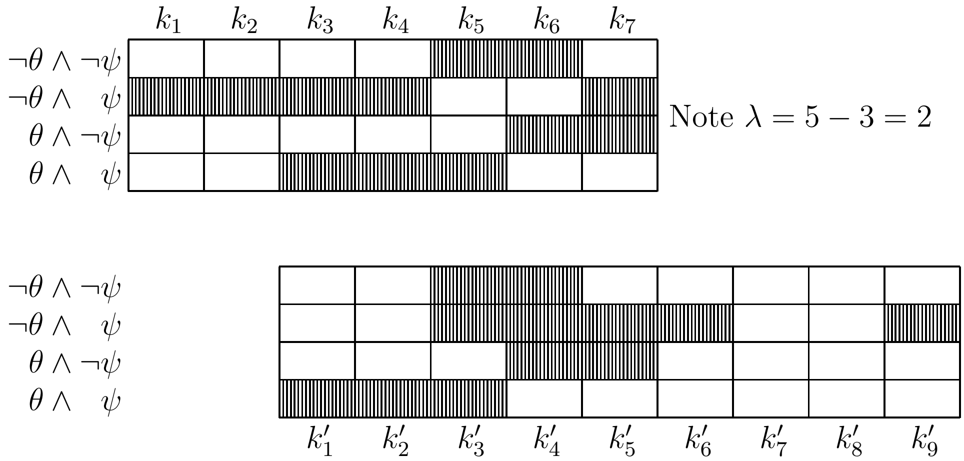

Figure 3.5: $\vec{k'} = \vec{k} \oplus_\psi \theta$, where $a = 5$

By (3.8), we have that $\oplus_\lambda$ is an AGM revision operator iff $\lambda > 0$. However,

$$\begin{aligned} \lambda \leq 0 \quad &\Longleftrightarrow \quad (\neg\psi)^{\vec{k}} - (\theta \vee \neg\psi)^{\vec{k}} \leq 0 \\ &\Longleftrightarrow \quad (\neg\psi)^{\vec{k}} - min\{(\theta)^{\vec{k}}, (\neg\psi)^{\vec{k}}\} \leq 0 \iff (\neg\psi)^{\vec{k}} \leq (\theta)^{\vec{k}} \end{aligned}$$

which is the exact condition for a revision by comparison operator to be an AGM operator, so we have not escaped this requirement for "successful revision" at all.

From (3.10) we have (C2) since, unlike revision by comparison, we are preserving the ordering between $\neg\theta \wedge \psi$ atoms. However, note that revision by comparison preserves ordering between k_6 and k_7 atoms, whereas in comparative conditionalization, they are split. Depending on the scenario, preserving these orderings might be a more desirable property.

3.4 Conclusions

Again, like the postulate sets, none of the iterated revision operators solve all of our problems, although we can demonstrate each operator's usefulness in several cases. Spohn conditionalisation is certainly the most versatile, we have given several variations showing how it can be applied to different scenarios, however as noted it's reliance on arbitary ordinals is at odds with our aim of reproducing intuitive thinking.

One of the main things that is now apparent is that the facts alone are not necessarily enough. Only σ-liberation requires nothing more than the input sentences to perform revision, but we deemed it over-simplistic. How much you believe someone is never purely based on what they tell you, but a collection of information; their relationship to you, the tone of their voice, even things like appearance—"Don't trust him, his eyes are too close together".

Chapter 4

Conclusions and Future Work

First, we examined belief revision using knowledge bases and the AGM postulate system, noting it's relationship with rational consequence. Also we demonstrated the equivalence between knowledge bases with rational consequence operators and belief sentences with faithful orderings.

After examining some failures of the AGM system, we gave 3 refinements on the idea—$\vec{k}$ and sequence based epistemic state formulations, which were designed to cope better with iterated revisions than knowledge bases did. We introduced Lehmann's and Darwiche and Pearls' postulate systems for these epistemic states, examining the differences between the two. We also gave a non-iterative formulation for an update operator, that gave better results when considering change in a dynamic world, as opposed to new information about a static world.

Finally, we used our epistemic-state formulations to present a series of revision operators capable of iterated revision, evaluating their usefulness in terms of intuitive properties and postulates. However, several intuitively reasonable cases were noted that were not possible at the moment.

4.1 Non-naïve Revision

Many avenues of investigation are currently restricted by the fact that regardless of which postulate set we are using, we are forced to accept the new information as true. For example, we noted earlier that Spohn's relative conditionalization seemed a reasonable idea, but because it's not guaranteed that $\theta \in [\vec{k} \ominus_\lambda \theta]$, it doesn't obey any of the postulate sets. However, example (3.13) gave us an intuitively sensible situation where such reasoning is applied.

Also, in section 3.1.1, we noted that when we had the situation

$$\sigma = p \wedge q \cdot p \wedge r \cdot p \wedge s \cdot p \wedge t \cdot p \wedge u \cdot \neg p \cdot v, \qquad \forall v \in SL, v \neq p, q, r, s, t, u$$

in terms of removing the minimal amount of knowledge, it would make more sense to remove $\neg p$ than every sentence $p \wedge \delta$. However we are prevented from using such reasoning in general, since

$$\neg p \in [p \wedge q \cdot p \wedge r \cdot p \wedge s \cdot p \wedge t \cdot p \wedge u \cdot \neg p] \qquad \text{by (I2)}$$

and given a v such that $\neg v \notin [p \wedge q \cdot p \wedge r \cdot p \wedge s \cdot p \wedge t \cdot p \wedge u \cdot \neg p]$,

$$\begin{aligned}&[p \wedge q \cdot p \wedge r \cdot p \wedge s \cdot p \wedge t \cdot p \wedge u \cdot \neg p \cdot v] \\ &\quad = [p \wedge q \cdot p \wedge r \cdot p \wedge s \cdot p \wedge t \cdot p \wedge u \cdot \neg p] + v \supseteq Cn(\neg p \wedge v) \qquad \text{by } (\odot 3), \text{(I2)},\end{aligned}$$

So the only case when we could keep the sentences $p \wedge \delta$ is when $\neg v \in [\neg p]$, by (C2)/(I7'). In the example above, it seemed reasonable with an arbitary independent v though.

However, we cannot simply remove the requirement that we accept the newest sentence without repercussions. Just considering knowledge bases, if we removed (*2) and (*5), "stubborn" operators such as

$$K * \theta = \begin{cases} K + \theta & \text{if } K + \theta \text{ is satisfiable,} \\ K & \text{otherwise.} \end{cases}$$

become possible, where knowledge is only accepted if it doesn't contradict what we already know—clearly such revision operators are undesirable.

An alternative approach would be to put a "wrapper around" AGM-compliant functions, similar to the representation of update operators. For instance,

$$[\sigma]^{\Gamma'} = [Ex(\sigma)]^{\Gamma} = Cn(\Gamma(Ex(\sigma), \emptyset))$$

where $Ex()$ is a function that takes a sequence, returning the sequence with "exceptional" or "unbelievable" sentences removed. Instead of modifying Darwiche and Pearl or Lehmann Postulates, a separate set of conditions for such a function could be investigated.

Bibliography

[1] R. Booth, On the Logic of Iterated Non-prioritised Revision, *http://isys.informatik.uni-leipzig.de/mitarb/booth/hagenws.pdf*

[2] R. Booth, S. Chopra, A. Ghose, T. Meyer, Belief Liberation (and Retraction), *Proceedings of the Ninth Conference on Theoretical Aspects of Rationality and Knowledge*, 20 (2003) 159–172

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

[4] E. Fermé and H. Rott, Revision by comparison, *Artificial Intelligence* 157 (2004) 5–47

[5] H. Katsuno and A. O. Mendelzon, Propositional knowledge base revision and minimal change, *Artificial Intelligence* 52 (1991) 263–294

[6] H. Katsuno and A. O. Mendelzon, On the difference between updating a knowledge base and revising it, *Cambridge Tracts in Theoretical Computer Science* 29 (1992) 183–203

[7] S. Konieczny, R. Pérez, A framework for iterated revision, *Journal of Applied Non-classical Logics* 10 pt. 3–4 (2000) 339–367

[8] D. Lehmann, Belief revision, revised, *Proc. Fourteenth International Joint Conference on Articial Intelligence, IJCAI'95* (1995) 1534–1540

[9] J. Paris, Non-monotonic Logic. *http://www.maths.man.ac.uk/DeptWeb/Homepages/jbp/mt4181.ps*

[10] W. Spohn, Ranking finctions, AGM style, in Spinning Ideas, Electronic Essays Dedicated to Peter Gärdenfors on His Fiftieth Brithday, 1999 *http://www.lucs.lu.se/spinning/categories/dynamics/Spohn/Spohn.pdf*