Chapter 8

Belief Revision

Pavlos Peppas

8.1 Introduction

Philippa, a Greek nineteen year old student at Patras University, has just discovered that Nikos and Angela are not her true parents; she was adopted when she was six months old from an orphanage in Sao Paulo. The news really shook Philippa. Much of what she used to believe all her life about herself and her family was wrong. After recovering from the initial shock she started putting her thoughts back in order: so that means that Alexandros is not really her cousin, and she did not take her brown eyes from (who she used to believe was) her grandmother, and she no longer needs to worry about developing high blood pressure because of the bad family history from both Nikos' and Angela's side. Moreover, she probably has siblings somewhere in Brazil, and if she really looked into it, she might be entitled to a Brazilian citizenship which could come in handy for that long trip she always wanted to make to Latin America.

This is a typical (although rather dramatic) instance of a belief revision scenario: a rational agent receives new information that makes her change her beliefs. In the principal case where the new information contradicts her initial belief state, the agent needs to withdraw some of the old beliefs before she can accommodate the new information; she also needs to accept the consequences that might result from the interaction of the new information with the (remaining) old beliefs.

The study of the process of belief revision, which gave rise to an exciting research area with the same name[1], can be traced back to the early 1980s. The article that is widely considered to mark the birth of the field is the seminal work of Alchourron, Gardenfors, and Makinson reported in [1]. As a matter of fact, the framework that evolved from [1] – now known as the *AGM paradigm* (or simply *AGM*) after the initials of its three founders – is to this date the dominant framework in Belief Revision.

Of course much has happened since 1985. The formal apparatus developed in [1] has been enhanced and thoroughly studied, new research directions have emerged from

[1] We shall use the capitalized term "Belief Revision" to refer to the *research area*; the same term in lower case letters will be used for the *process* of belief change.

it, connections with neighboring fields have been established, and a lot more is currently under way. This article will journey through the main developments in Belief Revision, pretty much in a historical order, starting with the classical AGM paradigm and following the trail till the present day.

8.2 Preliminaries

Let us first fix some notation and terminology. Alchourron, Gardenfors, and Makinson build their framework working with a formal language L governed by a logic which is identified by its consequence relation $\vdash$. Very little is assumed about L and $\vdash$, making the AGM paradigm quite general. In particular, L is taken to be closed under all Boolean connectives, and $\vdash$ has to satisfy the following properties:

(i) $\vdash \varphi$ for all truth-functional tautologies A (superclassicality).

(ii) If $\vdash (\varphi \to \psi)$ and $\vdash \varphi$, then $\vdash \psi$ (modus ponens).

(iii) $\vdash$ is consistent, i.e. $\nvdash L$.

(iv) $\vdash$ satisfies the deduction theorem, that is, $\{\varphi_1, \varphi_2, \ldots, \varphi_n\} \vdash \psi$ iff $\vdash \varphi_1 \wedge \varphi_2 \wedge \ldots \wedge \varphi_n \to \psi$.

(v) $\vdash$ is compact.

For a set of sentences Γ of L, we denote by $Cn(\Gamma)$ the set of all logical consequences of Γ, i.e. $Cn(\Gamma) = \{\varphi \in L\colon \Gamma \vdash \varphi\}$. A theory K of L is any set of sentences of L closed under $\vdash$, i.e. $K = Cn(K)$. We shall denote the set of all theories of L by $\mathbb{K}_L$. A theory K of L is complete iff for all sentences $\varphi \in L$, $\varphi \in K$ or $\neg\varphi \in K$. We shall denote the set of all consistent complete theories of L by $\mathbb{M}_L$. For a set of sentences Γ of L, $[\Gamma]$ denotes the set of all consistent complete theories of L that contain Γ. Often we shall use the notation $[\varphi]$ for a sentence $\varphi \in L$, as an abbreviation of $[\{\varphi\}]$. For a theory K and a set of sentences Γ of L, we shall denote by $K + \Gamma$ the closure under $\vdash$ of $K \cup \Gamma$, i.e. $K + \Gamma = Cn(K \cup \Gamma)$. For a sentence $\varphi \in L$ we shall often write $K + \varphi$ as an abbreviation of $K + \{\varphi\}$. Finally, the symbols $\top$ and $\bot$ will be used to denote an arbitrary (but fixed) tautology and contradiction of L respectively.

8.3 The AGM Paradigm

In AGM, beliefs are represented as sentences of L and belief sets as theories of L.[2] The process of belief revision is modeled as a function $*$ mapping a theory K and a sentence φ to a new theory $K * \varphi$. Of course certain constraints need to be imposed on $*$ in order for it to capture the notion of *rational belief revision* correctly. A guiding intuition in formulating these constraints has been the *principle of minimal change* according to which a rational agent ought to change her beliefs *as little as possible* in order to (consistently) accommodate the new information. Of course, at first glance it is not clear how one should measure change between belief sets, or even if the notion of minimal change is at all expressible within a purely logical framework.

[2] It should be noted that representing a belief set as a theory, presupposes that agents are *logically omniscient*. In this sense the AGM paradigm is tailored for *ideal reasoners*.

8.3.1 The AGM Postulates for Belief Revision

Despite the apparent difficulties, Gardenfors, [29], succeeded in formulating a set of eight postulates, known as the *AGM postulates for belief revision*,[3] which are now widely regarded to have captured much of what is the essence of rational belief revision:

$(K*1)$ $K*\varphi$ is a theory of L.

$(K*2)$ $\varphi \in K*\varphi$.

$(K*3)$ $K*\varphi \subseteq K+\varphi$.

$(K*4)$ If $\neg\varphi \notin K$ then $K+\varphi \subseteq K*\varphi$.

$(K*5)$ If φ is consistent then $K*\varphi$ is also consistent.

$(K*6)$ If $\vdash \varphi \leftrightarrow \psi$ then $K*\varphi = K*\psi$.

$(K*7)$ $K*(\varphi \wedge \psi) \subseteq (K*\varphi)+\psi$.

$(K*8)$ If $\neg\psi \notin K*\varphi$ then $(K*\varphi)+\psi \subseteq K*(\varphi \wedge \psi)$.

Any function $* : \mathbb{K}_L \times L \mapsto \mathbb{K}_L$ satisfying the AGM postulates for revision $(K*1)$ - $(K*8)$ is called an *AGM revision function* . The first six postulates $(K*1)$ - $(K*6)$ are known as the *basic* AGM postulates (for revision), while $(K*7)$ - $(K*8)$ are called the *supplementary* AGM postulates.

Postulate $(K*1)$ says that the agent, being an ideal reasoner, remains logically omniscient after she revises her beliefs. Postulate $(K*2)$ says that the new information φ should *always* be included in the new belief set. $(K*2)$ places enormous faith on the reliability of φ. The new information is perceived to be so reliable that it prevails over all previous conflicting beliefs, no matter what these beliefs might be.[4] Later in this chapter (Section 8.7) we shall consider ways of relaxing $(K*2)$. Postulates $(K*3)$ and $(K*4)$ viewed together state that whenever the new information φ does not contradict the initial belief set K, there is no reason to remove any of the original beliefs at all; the new belief state $K*\varphi$ will contain the whole of K, the new information φ, and whatever follows from the logical closure of K and φ (and nothing more). Essentially $(K*3)$ and $(K*4)$ express the notion of minimal change in the limiting case where the new information is consistent with the initial beliefs. $(K*5)$ says that the agent should aim for consistency at any cost; the *only* case where it is "acceptable" for the agent to fail is when the new information in itself is inconsistent (in which case, because of $(K*2)$, the agent can't do anything about it). $(K*6)$ is known as the *irrelevance of syntax postulate* . It says that the syntax of the new information has no effect on the revision process; all that matters is its content (i.e. the proposition it represents). Hence, logically equivalent sentences φ and ψ change a theory K in the same way.

[3] Although these postulate where first proposed by Gardenfors alone, they were extensively studied in collaboration with Alchourron and Makinson in [2]; thus their name.

[4] The high priority of φ over previous beliefs may not always be related to its reliability. For example, in the context of the *Ramsey Test* for conditionals, φ is incorporated into a theory K as part of the process of evaluating the acceptability of a counterfactual conditional $\varphi > \psi$ (see [30]).

Finally, postulates $(K * 7)$ and $(K * 8)$ are best understood taken together. They say that for any two sentences φ and ψ, if in revising the initial belief set K by φ one is lucky enough to reach a belief set $K * \varphi$ that is consistent with ψ, then to produce $K * (\varphi \wedge \psi)$ all that one needs to do is to expand $K * \varphi$ with ψ; in symbols $K * (\varphi \wedge \psi) = (K * \varphi) + \psi$. The motivation for $(K * 7)$ and $(K * 8)$ comes again from the principle of minimal change. The rationale is (loosely speaking) as follows: $K * \varphi$ is a minimal change of K to include φ and therefore there is no way to arrive at $K * (\varphi \wedge \psi)$ from K with "less change". In fact, because $K * (\varphi \wedge \psi)$ also includes ψ one might have to make further changes apart from those needed to include φ. If however ψ is consistent with $K * \varphi$, these further changes can be limited to simply adding ψ to $K * \varphi$ and closing under logical implications – no further withdrawals are necessary.

The postulates $(K * 1)$ - $(K * 8)$ are certainly very reasonable. They are simple, elegant, jointly consistent, and they follow quite naturally from the notion of minimal change. Moreover, according to Alchourron, Gardenfors, and Makinson, these postulates are not only sound but, given the limited expressiveness of a purely logical framework, there are also (in some sense) *complete*. Now this is a strong statement, especially since one can show that $(K * 1)$ - $(K * 8)$ do not suffice to uniquely determine the belief set $K * \varphi$ resulting from revising K by φ, given K and φ alone. In other words there is more than one function $*$ that satisfies $(K * 1)$ - $(K * 8)$. Yet the plurality of AGM revision functions should not be seen as a weakness of the postulates but rather as expressing the fact that different people may change their mind in different ways. Hence the AGM postulates simply circumscribe the territory of all different rational ways of revising belief sets.

Nevertheless, one may still be skeptical about whether the territory staked out by $(K * 1)$ - $(K * 8)$ contains nothing more but just rational belief revision functions. Further evidence is needed to support such a strong claim. Such evidence was indeed provided mainly in the form of formal results known as *representation results* connecting the AGM postulates with other models of belief revision. Some of the most important representation results will be discussed later in this chapter.

8.3.2 The AGM Postulates for Belief Contraction

Apart from belief revision, Alchourron, Gardenfors, and Makinson studied another type of belief change called *belief contraction* (or simply *contraction*), which can be described as the process of *rationally* removing from a belief set K a certain belief φ. Contraction typically occurs when an agent loses faith in φ and decides to give it up.[5] Simply taking out φ from K however will not suffice since other sentences that are present in K may reproduce φ through logical closure. Consider for example the theory $K = Cn(\{p \rightarrow q, p, q\})$ and assume that we want to contract K by q. Then, not only do we have to remove q from K, but we also need to give up (at least) one of $p \rightarrow q$ or p, for otherwise q will resurface via logical closure.

Like belief revision, belief contraction is formally defined as a function $\dot{-}$ mapping a theory K and a sentence φ to a new theory $K\dot{-}\varphi$. Once again a set of eight postulates was

[5] Another interesting instance of contraction is during argumentation. Consider two agents A and B that argue about a certain issue for which they have opposite views. It is quite likely that *for the sake of argument* the two agents will (temporarily) contract their beliefs to reach some common ground from which they will then starting building their case.

proposed, motivated by the principle of minimal change,[6] to constraint $\dot{-}$ in a way that captures the essence of *rational* belief contraction. These postulates, known as the *AGM postulates for belief contraction*, are the following:

($K\dot{-}1$) $K\dot{-}\varphi$ is a theory.

($K\dot{-}2$) $K\dot{-}\varphi \subseteq K$.

($K\dot{-}3$) If $\varphi \notin K$ then $K\dot{-}\varphi = K$.

($K\dot{-}4$) If $\nvdash \varphi$ then $\varphi \notin K\dot{-}\varphi$.

($K\dot{-}5$) If $\varphi \in K$, then $K \subseteq (K\dot{-}\varphi) + \varphi$.

($K\dot{-}6$) If $\vdash \varphi \leftrightarrow \psi$ then $K\dot{-}\varphi = K\dot{-}\psi$.

($K\dot{-}7$) $(K\dot{-}\varphi) \cap (K\dot{-}\psi) \subseteq K\dot{-}(\varphi \wedge \psi)$.

($K\dot{-}8$) If $\varphi \notin K\dot{-}(\varphi \wedge \psi)$ then $K\dot{-}(\varphi \wedge \psi) \subseteq K\dot{-}\varphi$.

Any function $\dot{-} : \mathbb{K}_L \times L \mapsto \mathbb{K}_L$ that satisfies ($K\dot{-}1$) - ($K\dot{-}8$) is called an *AGM contraction function*. Like the postulates for revision, ($K\dot{-}1$) - ($K\dot{-}8$) split into two groups: the first six postulates ($K\dot{-}1$) - ($K\dot{-}6$) are known as the *basic* AGM postulates for contraction, while ($K\dot{-}7$) - ($K\dot{-}8$) are called the *supplementary* AGM postulates for contraction.

Given the agent's logical omniscience, postulate ($K\dot{-}1$) is self-evident. Also self-evident is ($K\dot{-}2$) since by its very nature, contraction produces a belief set smaller than the original. Postulate ($K\dot{-}3$) says that if φ is not in the initial belief set K to start with, then there is no reason to change anything at all. ($K\dot{-}4$) tells us that the only sentences that are "immutable" are tautologies; all other sentences φ can in principle be removed from the initial beliefs K, and contraction will perform this removal no matter what the cost in epistemic value might be.[7] Postulate ($K\dot{-}5$), known as the *recovery postulate* says that contracting and then expanding by φ will give us back (at least) the initial theory K; in fact, because of ($K\dot{-}2$), we get back precisely K. The motivation behind ($K\dot{-}5$) is again the notion of minimal change: when contracting K by φ we should cut off only the part of K that is "related" to φ and *nothing else*. Hence adding φ back should restore our initial belief set.[8]

Postulate ($K\dot{-}6$), like its belief revision counterpart ($K * 6$), tells us that contraction is not syntax-sensitive: contraction by logically equivalent sentences produces the same result. The last two postulates relate the individual contractions by two sentences φ and ψ, to the contraction by their conjunction $\varphi \wedge \psi$. Firstly notice that to contract K by $\varphi \wedge \psi$ we need to give up either φ or ψ or both. Consider now a belief $\chi \in K$ that survives the contraction by φ, as well as the contraction by ψ (i.e. $\chi \in K\dot{-}\varphi$ and $\chi \in K\dot{-}\psi$). This in a sense means that, within the context of K, χ is not related to neither φ nor ψ and therefore it is also not related to their conjunction $\varphi \wedge \psi$; hence, says ($K\dot{-}7$), by the principle of minimal change χ should not be affected by the contraction of K by $\varphi \wedge \psi$. Finally, for

[6] In this context the principle of minimal change runs as follows: during contraction as little as possible should be given up from the initial belief set K in order to remove φ.

[7] The remarks for postulate ($K * 2$) are also relevant here.

[8] It should be noted though that, despite its intuitive appeal, the recovery postulate is among the most controversial AGM postulates – see [61] for a detailed discussion.

($K\dot{-}8$) assume that $\varphi \notin K\dot{-}(\varphi \wedge \psi)$. Since $K\dot{-}\varphi$ is the minimal change of K to remove φ, it follows that $K\dot{-}(\varphi \wedge \psi)$ can not be larger than $K\dot{-}\varphi$. Postulate ($K\dot{-}8$) in fact makes it smaller or equal to it; in symbols $K\dot{-}(\varphi \wedge \psi) \subseteq K\dot{-}\varphi$.

The AGM postulates for contraction are subject to the same criticism as their counterparts for revision: if completeness is to be claimed, one would need more than just informal arguments about their intuitive appeal. Some hard evidence is necessary.[9]

A first piece of such evidence comes from the relation between AGM revision and contraction functions. That such a connection between the two types of belief change should exist was suggested by Isaac Levi before Alchourron, Gardenfors, and Makinson formulated their postulates. More precisely, Levi argued that one should in principle be able to define revision in terms of contraction as follows: to revise K by φ, first contract K by $\neg\varphi$ (thus removing anything that may contradict the new information) and then expand the resulting theory with φ. This is now known as the *Levi Identity* :

$$K * \varphi = (K\dot{-}\neg\varphi) + \varphi \qquad \text{(Levi Identity)}$$

Alchourron, Gardenfors, and Makinson proved that the functions induced from their postulates do indeed satisfy the Levi Identity:

Theorem 1 *(*Alchourron, Gardenfors, and Makinson [1]*). Let $\dot{-}$ be any function from $\mathbb{K}_L \times L$ to $\mathbb{K}_L$ that satisfies the postulates ($K\dot{-}1$) - ($K\dot{-}8$). Then the function $*$ produced from $\dot{-}$ by means of the Levi Identity, satisfies the postulates ($K * 1$) - ($K * 8$).*[10]

As a matter of fact it turns out that Levi's method of producing revision functions is powerful enough to cover the *entire territory* of AGM revision functions; i.e. for *every* AGM revision function $*$ there is an AGM contraction function $\dot{-}$ that produces $*$ by means of the Levi Identity.

The fact that AGM revision and contraction functions are related so nicely in the way predicted by Levi, is the first piece of formal evidence to provide mutual support for the AGM postulates for contraction and revision.

A process that defines contraction in terms of revision is also available. This is known as the *Harper Identity* :

$$K\dot{-}\varphi = (K * \neg\varphi) \cap K \qquad \text{(Harper Identity)}$$

Like the Levi Identity, the Harper Identity is a sound and complete method for constructing contraction functions; i.e. the function $\dot{-}$ generated from an AGM revision function by means of the Harper Identity satisfies ($K\dot{-}1$) - ($K\dot{-}8$) and conversely, every AGM contraction function can be generated from a revision function by means of the Harper Identity. In fact, by combining the Levi and the Harper Identity one makes a full circle: if we start with an AGM contraction function $\dot{-}$ and use the Levi Identity to produce a revision function $*$, which in turn is then used to produce a contraction function via the Harper Identity, the result we end up with is the same contraction function $\dot{-}$ we started with.

[9] Incidentally, like with the AGM postulates for belief revision, one can show that there exists more than one function $\dot{-}$ satisfying ($K\dot{-}1$) - ($K\dot{-}8$).

[10] The result still holds even if $\dot{-}$ does not satisfy ($K\dot{-}5$) (i.e. the *recovery postulate*).

8.3.3 Selection Functions

Having identified the class of rational revision and contraction functions axiomatically, the next item on the agenda is to develop *constructive models* for these functions. It should be noted that, because of the Levi Identity, any constructive model for contraction functions can immediately be turned into a constructive model for revision functions; the converse is also true thanks to the Harper Identity. In this and the following two sections we shall review the most popular constructions for revision and contraction, starting with *partial meet contractions* – a construction for contraction functions.

Consider a theory K, and let φ be some non-tautological sentence in K that we would like to remove from K. Given that we need to adhere to the principle of minimal change, perhaps the first thing that comes to mind is to identify a maximal subset of K that fails to entail φ and define that to be the contraction of K by φ. Unfortunately, there is, in general, more than one such subset, and it is not at all obvious how to choose between them.[11] Nevertheless, these subsets are a very good starting point. We shall therefore give them a name: any maximal subset of K that fails to entail φ is called a *φ-remainder*.[12] The set of all φ remainders is denoted by $K \perp\!\!\!\perp \varphi$.[13]

As already mentioned, it is not clear how to choose between φ-remainders, since they are all equally good from a purely logical point of view. *Extra-logical factors* need to be taken into consideration to separate the *most plausible* φ-remainders from the rest. In the AGM paradigm, this is accomplished through *selection functions* . Formally, a selection function for a theory K is any function γ that maps a non-empty collection X of subsets of K to a non-empty subset $\gamma(X)$ of X; i.e. $\emptyset \neq \gamma(X) \subseteq X$. Intuitively, a selection function is used to pick up the "best" φ-remainders; i.e. the elements of $\gamma(K \perp\!\!\!\perp \varphi)$ are the most "valuable" (in an epistemological sense) among all φ-remainders.

Clearly, for a fixed theory K, there are many different selection functions, each one with a different set of "best" remainders. Only one of them though corresponds to the extra-logical factors that determine the agent's behavior. Once this function is given, it is possible to uniquely determine the contraction of K by any sentence φ by means of the following condition:

(M-) $\quad K \dot{-} \varphi = \bigcap \gamma(K \perp\!\!\!\perp \varphi)$.

Condition (M-) tells us that in contracting K by φ we should keep only the sentences of K that belong to all maximally plausible φ-remainders. This is a neat and intuitive construction, and it turns out that the functions $\dot{-}$ so produced satisfy many (but not all) of the AGM postulates for contraction.[14] To achieve an exact match between the functions produced from (M-) and the AGM contraction functions, we need to confine the selection functions γ fed to (M-) to those that are *transitively relational*.

A selection function γ is *transitively relational* iff it can be produced from a transitive binary relation $\ll$ in 2^K by means of the following condition:

[11] Consider for example the theory $K = Cn(\{p, q\})$, where p and q are propositional variables, and suppose that we want to contract by $p \wedge q$. There are more that one maximal subsets of K failing to entail $p \wedge q$, one of which contains p but not q, while another contains q but not p.

[12] In other words, a φ-remainder is a subset K' of K such that (i) $K' \nvdash \varphi$, and (ii) for any $K'' \subseteq K$, if $K' \subset K''$ then $K'' \vdash \varphi$.

[13] In the limiting case where φ is a tautology, $K \perp\!\!\!\perp \varphi$ is defined to be $\{K\}$.

[14] In particular, they satisfy the *basic* postulates $(K\dot{-}1)$ - $(K\dot{-}6)$ but fail to satisfy the *supplementary* postulates $(K\dot{-}7)$ and $(K\dot{-}8)$.

(TR) $\gamma(K \perp\!\!\!\perp \varphi) = \{ K' \in K \perp\!\!\!\perp \varphi$: for all $K'' \in K \perp\!\!\!\perp \varphi$, $K'' \ll K' \}$.

Intuitively, $\ll$ is to be understood as an ordering on subsets of K representing comparative epistemological value; i.e. $K'' \ll K'$ iff K' is at least as valuable as K''. Hence, (TR) tells us that a selection function γ is transitively relational if it makes its choices based on an underlying ordering $\ll$; i.e. the "best" remainders picked up by γ are the ones that are most valuable according to $\ll$.

Since this is the first time we encounter an ordering $\ll$ as part of a constructive model for belief change, it is worth noting that such orderings are central to the study of Belief Revision and we shall encounter many of them in the sequel. They come with different names (epistemic entrenchment, system of spheres, ensconcement, etc), they apply at different objects (remainders, sentences, possible worlds, etc) and they may have different intended meanings. In all cases though they are used (either directly or indirectly) to capture the extra-logical factors that come into play during the process of belief revision/contraction.

Any function $\dot{-}$ constructed from a transitive relational selection function by means of (M-), is called a *transitive relational partial meet contraction function*. The following theorem is one of the first major results of the AGM paradigm, and the second piece of formal evidence reported herein in support of the postulates ($K\dot{-}1$) - ($K\dot{-}8$) for contraction (and via the Levi Identity, of the postulates ($K * 1$) - ($K * 8$) for revision):

Theorem 2 (Alchourron, Gardenfors, and Makinson [1]). *Let K be a theory of L and $\dot{-}$ a function from $\mathbb{K}_L \times L$ to $\mathbb{K}_L$. Then $\dot{-}$ is a transitive relational partial meet contraction function iff it satisfies the postulates ($K\dot{-}1$) - ($K\dot{-}8$).*

In other words, when (M-) is fed transitively relational selection functions γ it generates functions $\dot{-}$ that satisfy *all* the AGM postulates for contraction; conversely, *any* AGM contraction function $\dot{-}$ can be constructed from a transitively relational selection function γ by means of (M-).

We conclude this section by considering two special cases of selection functions lying at opposite ends of the selection-functions-spectrum. The first, which we shall denote by γ_F, always selects *all* elements of its argument; i.e. for any X, $\gamma_F(X) = X$. Hence for a fixed theory K the function γ_F for K, picks up *all* φ-remainders for any φ. The contraction function produced from γ_F by means of (M-) is called a *full meet contraction function*. Notice that in the construction of a full meet contraction function $\dot{-}$, γ_F is superfluous since $\dot{-}$ can be produced by means of the following condition:

(F-) $K\dot{-}\varphi = \bigcap K \perp\!\!\!\perp \varphi$.

The distinctive feature of a full meet contraction function is that, among all contraction functions, it always produces the *smallest* theory. In particular, any function $\dot{-} : \mathbb{K}_L \times L \mapsto \mathbb{K}_L$ that satisfies ($K\dot{-}1$) - ($K\dot{-}3$) and ($K\dot{-}5$), is such that $K\dot{-}\varphi$ always includes $\bigcap K \perp\!\!\!\perp \varphi$ for any $\varphi \in L$. As an indication of how severe full meet contraction is, we note that the revision function $*$ produced from it (via the Levi Identity) is such that $K * \varphi = Cn(\varphi)$ for all φ contradicting K; in other words, for any $\neg\varphi \in K$, the (full-meet-produced) revision of K by φ removes *all* previous beliefs (other than the consequences of φ).

At the opposite end of the spectrum are *maxichoice* contraction functions. These are the functions constructed from selection functions γ_M that always pick up only *one* element of their arguments; i.e. for any X, $\gamma_M(X)$ is a *singleton*. Hence, for any sentence φ, when

γ_M is applied to the set $K \perp\!\!\!\perp \varphi$ of all φ-remainders, it selects only one of them as the "best" φ-remainder. It should be noted that such selection functions γ_M are *not* in general transitively relational, and maxichoice contraction functions do not satisfy all the AGM postulates for contraction. A peculiar feature of maxichoice contractions $\dot{-}$ is that they produce (via the Levi Identity) highly "opinionated" revision functions $*$; i.e. whenever the new information φ contradicts the initial belief set K, such functions $*$ always return a *complete* theory $K * \varphi$.

8.3.4 Epistemic Entrenchment

As mentioned earlier, selection functions are essentially a formal way of encoding the extra-logical factors that determine the beliefs that a sentence φ should take away with it when it is rooted out of a theory K.

These extra-logical factors relate to the *epistemic value* that the agent perceives her individual beliefs to have within the context of K. For example a law-like belief ψ such as "all swans are white", is likely to be more important to the agent than the belief χ that "Lucy is a swan". Consequently, if a case arises where the agent needs to choose between giving up ψ or giving up χ (e.g., when contracting with the belief "Lucy is white") the agent will surrender the latter.

Considerations like these led Gardenfors and Makinson [32] to introduce the notion of *epistemic entrenchment* as another means of encoding the extra-logical factors that are relevant to belief contraction. Intuitively, the epistemic entrenchment of a belief ψ is the degree of resistance that ψ exhibits to change: the more entrenched ψ is, the less likely it is to be swept away during contraction by some other belief φ.

Formally, epistemic entrenchment is defined as a preorder $\leq$ on L encoding the relative "retractibility" of individual beliefs; i.e. $\chi \leq \psi$ iff the agent is at least as (or more) reluctant to give up ψ than she is to give up χ. Once again, certain constraints need to be imposed on $\leq$ for it to capture its intended meaning:

(EE1) If $\varphi \leq \psi$ and $\psi \leq \chi$ then $\varphi \leq \chi$.

(EE2) If $\varphi \vdash \psi$ then $\varphi \leq \psi$.

(EE3) $\varphi \leq \varphi \wedge \psi$ or $\psi \leq \varphi \wedge \psi$.

(EE4) When K is consistent, $\varphi \notin K$ iff $\varphi \leq \psi$ for all $\psi \in L$.

(EE5) If $\psi \leq \varphi$ for all $\psi \in L$, then $\vdash \varphi$.

Axiom (EE1) states that $\leq$ is transitive. (EE2) says that the stronger a belief is logically, the less entrenched it is. At first this may seem counter-intuitive. A closer look however will convince us otherwise. Consider two beliefs φ and ψ both of them members of a belief set K, and such that $\varphi \vdash \psi$. Then clearly, if one decides to give up ψ one will also have to remove φ (for otherwise logical closure will bring ψ back). On the other hand it is possible to give up φ and retain ψ. Hence giving up φ produces less epistemic loss than giving up ψ and therefore the former should be preferred whenever a choice exists between the two. Thus axiom (EE2). For axiom (EE3) notice that, again because of logical closure, one can not give up $\varphi \wedge \psi$ without removing at least one of the sentences φ or ψ. Hence either φ or

ψ (or even both) are at least as vulnerable as $\varphi \wedge \psi$ during contraction. We note that from (EE1) - (EE3) it follows that $\leq$ is *total*; i.e. for any two sentences $\varphi, \psi \in L$, $\varphi \leq \psi$ or $\psi \leq \varphi$.

The final two axioms deal with the two ends of this total preorder $\leq$, i.e. with its minimal and its maximal elements. In particular, axiom (EE4) says that in the principal case where K is consistent, all non-beliefs (i.e. all the sentences that are not in K) are minimally entrenched. At the other end of the entrenchment spectrum we have all tautologies, which according to (EE5) are the only maximal elements of $\leq$ and therefore the hardest to remove (in fact, in the AGM paradigm it is impossible to remove them).

Perhaps not surprisingly it turns out that for a fixed belief set K there is more than one preorder $\leq$ that satisfies the axioms (EE1) - (EE5). Once again this is explained by the subjective nature of epistemic entrenchment (different agents may perceive the epistemic importance of a sentence φ differently). However, once the epistemic entrenchment $\leq$ chosen by an agent is given, it should be possible to determine uniquely the result of contracting her belief set K by *any* sentence φ. This is indeed the case; condition (C-) below defines contraction in terms of epistemic entrenchment:[15]

(C-) $\psi \in K\dot{-}\varphi$ iff $\psi \in K$ and either $\varphi < \varphi \vee \psi$ or $\vdash \varphi$.

Gardenfors and Makinson proved the following representation result which essentially shows that for the purpose of belief contraction, an epistemic entrenchment $\leq$ is all the information ones needs to know about extra-logical factors:

Theorem 3 (Gardenfors and Makinson [32]). *Let K be a theory of L. If $\leq$ is a preorder in L that satisfies the axioms (EE1) - (EE5) then the function defined by (C-) is an AGM contraction function. Conversely, if $\dot{-}$ is an AGM contraction function, then there is preorder $\leq$ in L that satisfies the axioms (EE1) - (EE5) as well as condition (C-).*

Theorem 3 is the third piece of formal evidence in support of the postulates (K-1) - (K-8) for contraction.

It should be noted that, thanks to the Levi Identity, (C-) can be reformulated in a way that defines directly a revision function $*$ from an epistemic entrenchment $\leq$:

(E*) $\psi \in K * \varphi$ iff either $(\varphi \rightarrow \neg\psi) < (\varphi \rightarrow \psi)$ or $\vdash \neg\varphi$.

An analog to Theorem 3, connecting epistemic entrenchments with AGM revision functions via (E*), is easily established [84, 76].

8.3.5 System of Spheres

Epistemic entrenchment together with condition (C-) is a *constructive approach* to modeling belief contraction, as opposed to the AGM postulates ($K\dot{-}1$)-($K\dot{-}8$) which model contraction axiomatically. Another constructive approach, this time for belief revision, has been proposed by Grove in [38]. Building on earlier work by Lewis [57], Grove uses a structure called a *system of spheres* to construct revision functions. Like an epistemic entrenchment, a system of sphere is essentially a preorder. However the objects being ordered are no longer sentences but consistent complete theories.

[15] At first glance, condition (C-) may seem unnatural. Indeed there is an equivalent and much more intuitive way to relate epistemic entrenchments with contraction functions (see condition (C$\leq$) in [31]). However, condition (C-) is more useful as a construction mechanism for contraction functions.

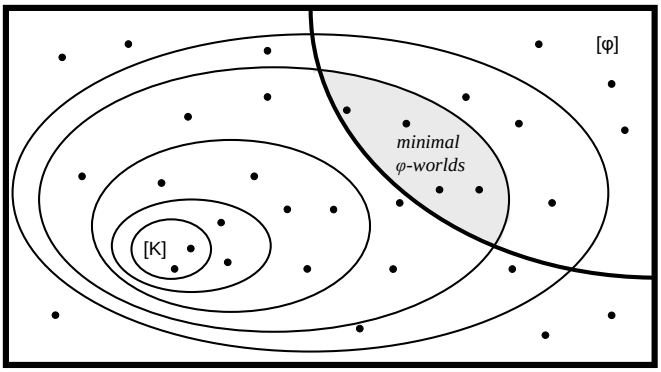

Figure 8.1: A System of Spheres

Given an initial belief set K a *system of spheres centered on* $[K]$ is formally defined as a collection S of subsets of $\mathbb{M}_L$, called *spheres*, satisfying the following conditions (see Figure 8.1):[16]

(S1) S is totally ordered with respect to set inclusion; that is, if $V, U \in S$ then $V \subseteq U$ or $U \subseteq V$.

(S2) The smallest sphere in S is $[K]$; that is, $[K] \in S$, and if $V \in S$ then $[K] \subseteq V$.

(S3) $\mathbb{M}_L \in S$ (and therefore $\mathbb{M}_L$ is the largest sphere in S).

(S4) For every $\varphi \in L$, if there is any sphere in S intersecting $[\varphi]$ then there is also a smallest sphere in S intersecting $[\varphi]$.

Intuitively a system of spheres S centered on $[K]$ represents the *relative plausibility* of consistent complete theories, which in this context play the role of possible worlds: the closer a consistent complete theory is to the center of S, the more plausible it is. Conditions (S1) - (S4) are then read as follows. (S1) says that any two worlds in S are always comparable in terms of plausibility. Condition (S2) tells us that the most plausible worlds are those compatible with the agent's initial belief set K. Condition (S3) says that all worlds appear somewhere in the plausibility spectrum. Finally condition (S4), also known as the *Limit Assumption*, is of a more technical nature. It guarantees that for any consistent sentence φ, if one starts at the outermost sphere $\mathbb{M}_L$ (which clearly contains a φ-world) and gradually progresses towards the center of S, one will eventually meet the *smallest* sphere containing φ-worlds. In other words, the spheres in S containing φ-worlds do not form an infinitely decreasing chain; they always converge to a *limit* which is also *in* S. The smallest sphere in S intersecting $[\varphi]$ is denoted $c(\varphi)$. In the limiting case where φ is inconsistent, $c(\varphi)$ is defined to be equal to $\mathbb{M}_L$.

Suppose now that we want to revise K by a sentence φ. Intuitively, the rational thing to do is to select the most plausible φ-worlds and define through them the new belief set $K * \varphi$:

[16] Recall that $\mathbb{M}_L$ is the set of all consistent complete theories of L, and for a theory K of L, $[K]$ is the set of all consistent complete theories that contain K.

$$\text{(S*)} \qquad K * \varphi = \begin{cases} \bigcap(c(\varphi) \cap [\varphi]) & \text{if } \varphi \text{ is consistent} \\ L & \text{otherwise} \end{cases}$$

Condition (S*) is precisely what Grove proposed as a means of constructing a revision function $*$ from a system of spheres S. Moreover Grove proved that his construction is sound and complete with respect to the AGM postulates for revision:

Theorem 4 (Grove [38]). *Let K be a theory and S a system of spheres centered on $[K]$. Then the revision function $*$ defined via (S*) satisfies the AGM postulates $(K*1)$ - $(K*8)$. Conversely, for any theory K and AGM revision function $*$, there exists a system of spheres S centered on $[K]$ that satisfies (S*).*

Theorem 4 is the fourth and final piece of formal evidence in support of the AGM postulates for revision (and contraction). In a sense, it also marks the end of the "classical era" in Belief Revision.[17] Therefore this is a good point to take a step back and quickly review what has been discussed so far.

Two types of belief change were examined: belief revision and belief contraction. For each of them a set of postulates was proposed to capture the notion of rationality in each case. In formulating the postulates, Alchourron, Gardenfors, and Makinson relied on the principle of minimal change for guidance. Although the two sets of postulates were motivated independently, the connection between revision and contraction predicted by Levi was shown to hold within the AGM paradigm. This result provided the first formal evidence in support of the appropriateness of the AGM postulates.

The second piece of evidence came with the first constructive model proposed by Alchourron, Gardenfors, and Makinson based on selection functions, together with the corresponding representation result matching partial meet contraction functions with the postulates $(K\dot{-}1)$ - $(K\dot{-}8)$.

Later, a second constructive model for contraction functions was introduced by Gardenfors and Makinson, based on the notion of epistemic entrenchment. Formally an epistemic entrenchment is a special preorder on sentences representing the relative resistance of beliefs to change. Gardenfors and Makinson proved that the class of contraction functions produced by epistemic entrenchments coincides precisely with those satisfying the AGM postulates for contraction – yet another piece of strong evidence in support of the AGM postulates.

Grove completed the picture by providing what essentially amounts to a possible world semantics for the AGM postulates $(K*1)$ - $(K*8)$ for revision. His semantics is based on a special preorder on possible worlds called a system of spheres, which is intended to represent the relative plausibility of possible worlds, given the agent's initial belief set. Based on systems of spheres, Grove provided a very natural definition of belief revision. The fact that Grove's intuitive semantics were proven to be sound and complete with respect to the AGM postulates for revision, is perhaps the most compelling formal evidence for the appropriateness of the AGM postulates. Figure 8.2 summaries the first main results of the AGM paradigm.

[17] With one notable exception: Spohn's work [94] on iterated revision which will be discussed in Section 8.6.

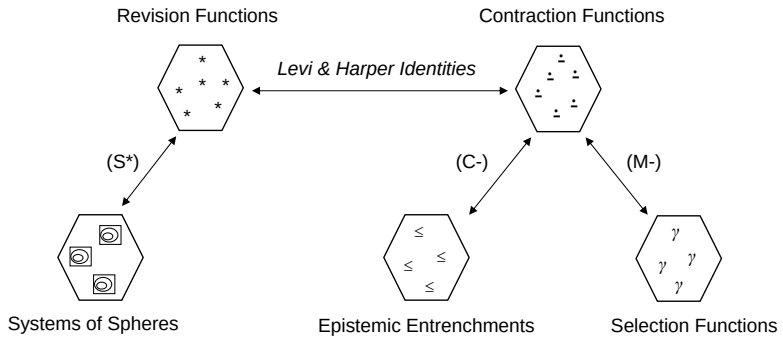

Figure 8.2: The AGM paradigm in the late 1980's

8.4 Belief Base Change

The models and results of the AGM paradigm depicted in Figure 8.2 are so neat, that one almost feels reluctant to change anything at all. Yet these elegant results rest on assumptions that, in a more realistic context, are disputable; moreover some important issues on Belief Revision were left unanswered by the (original) AGM paradigm. Hence researchers in the area took on the task of extending the AGM paradigm while at the same time preserving its elegance and intuitive appeal that has made it so popular and powerful. Considerable efforts have been made to maintain the connections depicted in Figure 8.2 in a more generalized or more pragmatic framework.

Among AGM's founding assumptions, one of the most controversial is the modeling of an agent's beliefs as a *theory*. This is unrealistic for a number of reasons. Firstly, theories are *infinite* objects and as such can not be incorporated directly into a computational framework. Therefore any attempt to use AGM's revision functions in an artificial intelligence application will have to be based on *finite representations* for theories, called *theory bases*. Ideally, a theory base would not only represent (in a finite manner) the sentences of the theory, but also the extra-logical information needed for belief revision.

Computational considerations however are not the only reason that one may choose to move from theories to theory bases. Many authors, [27, 86, 39, 66], make a distinction between the *explicit beliefs* of an agent, i.e. beliefs that the agent accepts in their own right, and beliefs that follow from logical closure. This distinction, goes the argument, plays a crucial role in belief revision since derived beliefs should not be retained if their support in explicit beliefs is gone. To take a concrete example, suppose that Philippa believes that "Picasso was Polish"; call this sentence φ. Due to logical closure, Philippa also holds the derived belief $\varphi \vee \psi$, where ψ can be any sentence expressible in the language, like "Picasso was Australian" or even "There is life on Mars". If later Philippa drops φ, it seems unreasonable to retain $\varphi \vee \psi$, since the latter has no independent standing but owes its presence solely to φ.[18]

Most of the work on belief base revision starts with a theory base B and a preference ordering $<$ on the sentences in B, and provides methods of revising B in accordance with

[18] This is often called the *foundational approach* to belief revision.

$<$. The belief base B is a set of sentences of L, which in principle (but not necessarily) is not closed under logical implication and for all practical purposes it is in fact *finite*. Nebel, [70], distinguishes between approaches that aim to take into account the difference between explicit and derived beliefs on one hand, and approaches that aim to provide a computational model for theory revision on the other. The former give rise to *belief base revision operations*, whereas the latter define *belief base revision schemes*. The main difference between the two is that the output of a belief base revision operation is again a belief base, whereas the output of a belief base revision scheme is a *theory*. This difference is due to the different aims and assumptions of the two groups. Belief base revision operations assume that the primary objects of change are belief bases, not theories.[19] Of course a revision on belief bases can be "lifted" to a revision on theories via logical closure. However this theory revision is simply an epiphenomenon; revision operators act only on the set of explicit beliefs. If one adopts this view, it is clear why the result of a belief base revision operation is again a belief base.

Belief base revision schemes on the other hand have been developed with a different goal in mind: to provide a concise representation of *theory* revision. We have seen that AGM revision functions need the entire theory K and an epistemic entrenchment $\leq$ associated with it to produce the new theory $K * \varphi$ (for any φ). However both K and $\leq$ are infinite objects. Moreover even when K is finite modulo logical equivalence, the amount of information necessary for $\leq$ is (in the worst case) exponential in the size of the finite axiomatization of K. By operating on a belief base B and an associated preference ordering $<$, belief base revision schemes provide a method of producing $K * \varphi$ from succinct representations. In that sense, as noted in [70], belief base revision schemes can be viewed as just another construction model for belief revision alongside epistemic entrenchments, systems of spheres, and selection functions.

In the following we shall review some of the most important belief base revision operations and belief base revision schemes. Our presentation follows the notation and terminology of [70].

8.4.1 Belief Base Change Operations

An obvious approach to define belief base operations is to follow the lead of the partial meet construction in Section 8.3.3.

In particular, let B be a belief base and φ a sentence in L. The definition of a φ-remainder can be applied to B without any changes despite the fact that B is (in principle) not closed under logical implication. The same is true for selection functions over subsets of B. Hence condition (M-) can be used verbatim for constructing belief base contraction functions. We repeat (M-) below for convenience, this time with a B subscript to indicate that it is no longer restricted to theories (for the same reason K is replaced by B):

$(\text{M-})_B \qquad B \dot{-} \varphi = \bigcap \gamma(B \perp\!\!\!\perp \varphi).$

Any function $\dot{-}$ constructed from a selection function γ by means of $(\text{M-})_B$ is called a *partial meet belief base contraction function* . Hansson, [41], characterized partial meet belief base contraction in terms of the following postulates:[20]

[19] Apart from the degenerate case where the two are identical.

[20] Earlier publications by Hansson also report on similar results. However [41] gives a more detailed and uniform presentation of his work on belief base contraction.

($B\dot{-}1$) If $\nvdash \varphi$ then $\varphi \notin Cn(B\dot{-}\varphi)$.

($B\dot{-}2$) $B\dot{-}\varphi \subseteq B$.

($B\dot{-}3$) If it holds for all subsets B' of B that $\varphi \in B'$ iff $\psi \in B'$, then $B\dot{-}\varphi = B\dot{-}\psi$.

($B\dot{-}4$) If $\psi \in B$ and $\psi \notin B\dot{-}\varphi$, then there is a set B' such that $B\dot{-}\varphi \subseteq B' \subseteq B$ and that $\varphi \notin Cn(B')$ but $\varphi \in Cn(B' \cup \{\psi\})$.

Theorem 5 (Hansson [41]). *A function $\dot{-}$ from $2^L \times L$ to 2^L is a partial meet belief base contraction function iff it satisfies ($B\dot{-}1$) - ($B\dot{-}4$).*

Some remarks are due regarding partial meet belief base contractions and their associated representation result. Firstly, the extra-logical information needed to produce $\dot{-}$ is not encoded as an ordering on the sentences of B (as it is typically the case with belief base revision schemes), but by a selection function γ on the subsets of B. Secondly, γ is not necessarily *relational*; i.e. γ is not necessarily defined in terms of a binary relation $\ll$. If such an assumption is made, further properties of the produced belief base contraction function $\dot{-}$ can be derived (see [41]). Finally, although this construction generates belief base contractions, it can also be used to produce belief base revisions by means of the following variant of the Levi Identity:

(BL) $B * \varphi = (B\dot{-}\neg\varphi) \cup \{\varphi\}$.

An alternative to partial meet belief base contraction is *kernel contraction* introduced by Hansson in [39] and originating from the work of Alchourron and Makinson on *safe contraction* [2].

Let B be a belief base and φ a sentence of L. A *φ-kernel* of B is a minimal subset of B that entails φ; i.e. B' is a φ-kernel of B iff $B' \subseteq B$, $B' \vdash \varphi$, and no proper subset of B' entails φ. We shall denote the set of all φ-kernels by $B \parallel \varphi$. An *incision function* σ for B is a function that maps a set X of subsets of B to a subset of $\bigcup X$ such that for all $T \in X$, $\sigma(X) \cap T \neq \emptyset$; i.e. σ picks up a subset of $\bigcup X$ that cuts across all elements of X.

Given an incision function σ for a belief base B, one can construct a contraction function $\dot{-}$ as follows:

($B\dot{-}$) $B\dot{-}\varphi = B - \sigma(B \parallel \varphi)$.

A function $\dot{-}$ constructed from an incision function σ by means of (K$\dot{-}$) is called a *kernel contraction function*. Once again Hansson, [39, 41], has provided an axiomatic characterization of kernel contractions. To this end, consider the postulate ($B\dot{-}5$) below:

($B\dot{-}5$) If $\psi \in B$ and $\psi \notin B\dot{-}\varphi$, then there is a set B' such that $B' \subseteq B$ and that $\varphi \notin Cn(B')$ but $\varphi \in Cn(B' \cup \{\psi\})$.

Clearly ($B\dot{-}5$) is a weaker version of ($B\dot{-}4$). It turns out that this weakening of ($B\dot{-}4$) is all that is needed for a precise characterization of kernel contractions:

Theorem 6 (Hansson [39, 41]). *A function $\dot{-}$ from $2^L \times L$ to 2^L is a kernel contraction function iff it satisfies ($B\dot{-}1$) - ($B\dot{-}3$) and ($B\dot{-}5$).*

It follows immediately from Theorems 5 and 6, that every partial meet belief base contraction is also a kernel contraction. The converse is not true. It is however possible to devise restrictions on incision functions such that the induced subclass of kernel contractions, called *smooth* kernel contractions, coincides with the family of partial meet belief base contraction (see [39]).

Once again, the belief base variant of the Levi Identity (BL) can be used to produce belief base revisions from kernel contractions.

8.4.2 Belief Base Change Schemes

Turning to belief base revision schemes, we need to keep in mind that, while these schemes operate on (prioritized) belief bases, their outcome are theories.

Perhaps the simplest such scheme is the analog of full meet contractions for belief bases shown below. As usual in condition (F-)$_B$ below B denotes a belief base and φ a sentence of L.

(F-)$_B$ $\quad B\dot{-}\varphi = \bigcap_{B' \in (B \perp\!\!\!\perp \varphi)} Cn(B' \cup \{\varphi \rightarrow \psi : \psi \in B\})$.

Let us call the functions produced from (F-)$_B$ *base-generated full meet contraction functions*. In a sense (F-)$_B$ can be viewed as a special case of (M-)$_B$; namely the case where the selection function γ picks up *all* φ-remainders. There are however two important differences between the two conditions. Firstly a new term, $\{\varphi \rightarrow \psi : \psi \in B\}$, has been added to (F-)$_B$ whose purpose is to secure the validity of the recovery postulate ($K\dot{-}6$) (see Theorems 7, 8 below). Secondly, in (F-)$_B$ the φ-remainders (together with the new term) are first *closed under logical implication* before intersected. Hence $B\dot{-}\varphi$, as defined by (F-)$_B$, is always a *theory*, which, furthermore, is not necessarily expressible as the logical closure of some subset of the initial belief base B.

As mentioned earlier, (F-)$_B$ can be viewed as the construction of a contraction function $\dot{-}$ mapping $Cn(B)$ and φ to the theory $B\dot{-}\varphi$. As such, one can assess base-generated full meet contraction functions against the AGM postulates. To this end, consider the following two new postulates from [86]:

($\dot{-}$8r) $\quad K\dot{-}(\varphi \wedge \psi) \subseteq Cn(K\dot{-}\varphi \cup K\dot{-}\psi)$

($\dot{-}$8c) $\quad$ If $\psi \in K\dot{-}(\varphi \wedge \psi)$ then $K\dot{-}(\varphi \wedge \psi) \subseteq K\dot{-}\varphi$.

It turns out that in the presence of ($K\dot{-}1$) - ($K\dot{-}7$), the above two postulates follow from ($K\dot{-}8$). Rott and del Val independently proved the following characterization of base-generated full meet contraction functions:

Theorem 7 (Rott [86], del Val [15]). *A function $\dot{-}$ from $\mathbb{K}_L \times L$ to $\mathbb{K}_L$ is a base-generated full meet contraction function iff it satisfies ($K\dot{-}1$) - ($K\dot{-}7$), ($\dot{-}8r$) and ($\dot{-}8c$).*

Notice that, by selecting all φ-remainders, condition (F-)$_B$ treats all sentences in the belief base B as equal. If however the belief base B is *prioritized*, a refined version of (F-)$_B$ is needed; one which among all φ-remainders selects only those whose sentences have the highest priority. In particular, assume that the belief base B is partitioned into n priority classes $B_1, B_2, \cdots, B_n$, listed in increasing order of importance (i.e. for $i < j$, the sentences in B_i are less important than the sentences in B_j). Given such a prioritization of B one can define an ordering on subsets of B as follows:

(B$\ll$) For any $T, E \subseteq B$, $T \ll E$ iff there is an $1 \leq i \leq n$ such that $T \cap B_i \subset E \cap B_i$ and for all $i < j \leq n$, $T \cap B_j = E \cap B_j$.

We can now refine (F-)$_B$ to select only the best φ-remainders with respect to $\ll$. In particular, let $max(B \perp\!\!\!\perp \varphi)$ denote the set of maximal φ-remainders with respect to $\ll$; i.e. $max(B \perp\!\!\!\perp \varphi) = \{B' \in (B \perp\!\!\!\perp \varphi)$: for all $E \subseteq B$, if $B' \ll E$ then $E \vdash \varphi\}$. The prioritized version of (F-)$_B$ is condition (P-)$_B$ below:

(P-)$_B$ $\quad B\dot{-}\varphi = \bigcap_{B' \in max(B \perp\!\!\!\perp \varphi)} Cn(B' \cup \{\varphi \rightarrow \psi : \psi \in B\})$.

All functions induced from (P-)$_B$ are called *base-generated partial meet contraction functions* . Clearly, all base-generated full meet contraction functions are also partial meet (simply set the partition of B to be a singleton; i.e. only containing B itself). Perhaps surprisingly, the converse is also true:

Theorem 8 (Rott [86], del Val [15]). *A function $\dot{-}$ from $\mathbb{K}_L \times L$ to $\mathbb{K}_L$ is a base-generated partial meet contraction function iff it satisfies ($K\dot{-}1$) - ($K\dot{-}7$), ($\dot{-}8r$) and ($\dot{-}8c$).*

It is important not to misread the equivalence between full meet and partial meet base-generated contraction functions implied by Theorem 8. In particular, Theorem 8 *does not* imply that for a given prioritized belief base B the function induced by (P-)$_B$ is the same as the function induced by (F-)$_B$. What Theorem 8 does entail is that for any prioritized belief base B there exists a non-prioritized belief base B', which is logically equivalent to B (i.e. $Cn(B) = Cn(B')$), and such that the contraction function produced from B and (P-)$_B$ is the same as the contraction function produced from B' and (F-)$_B$.

Clearly partial meet base-generated contraction functions can be used to construct revision functions via the Levi Identity; predictably, these functions are called *partial meet base-generated revision functions* . Among them, Nebel [69, 70] identified a special class, called *linear belief base revision functions*, with some interesting properties. Formally a linear belief base revision function is defined as a partial meet base-generated revision function that is produced from a *totally ordered* prioritized belief base; that is, the priority classes $B_1, B_2, \cdots, B_n$ of the initial belief base B are all *singletons*. It turns out that linear belief base revision functions coincide precisely with the original AGM revision functions:

Theorem 9 (Nebel [69], del Val [15]). *A function $*$ from $\mathbb{K}_L \times L$ to $\mathbb{K}_L$ is a linear belief base revision function iff it satisfies ($K*1$) - ($K*8$).*

The last belief base change scheme that we shall consider in this section is based on the notion of *ensconcement* [98, 100, 99]. Intuitively, an ensconcement is an ordering $\preccurlyeq$ on a belief base B that can be "blown up" to a full epistemic entrenchment $\leq$ related to $Cn(B)$. We can also think of it in another way. Consider a theory K and an epistemic entrenchment $\leq$ related to K that defines (via (E*)) a revision function $*$.[21] If we want to generate $*$ from a belief base B of K, we also need some sort of "base" for $\leq$. That is precisely what an ensconcement is: a (typically) concise representation of an epistemic entrenchment.

Formally, an ensconcement ordering $\preccurlyeq$ on a belief base B is a total preorder on B satisfying the following conditions:

[21] To be more precise, $\leq$ gives only a *partial* definition of $*$; namely only its restriction to K. For a complete specification of $*$ we would need a whole family of epistemic entrenchments, one for every theory of L. This abuse in terminology occurs quite frequently in this chapter.

($\leqslant$1) For all non-tautological sentences φ in B, $\{\psi \in B : \varphi \prec \psi\} \nvdash \varphi$.

($\leqslant$2) For all $\varphi \in B$, φ is a tautology iff $\psi \leqslant \varphi$ for all $\psi \in B$.

Clearly an ensconcement ordering $\leqslant$ satisfies the following *priority consistency condition* [85]:

(PCC) For all $\varphi \in B$, if B' is a nonempty subset of B that entails φ then there is a $\psi \in B'$ such that $\psi \leqslant \varphi$.

Rott, [85], has shown that (PCC) is a *necessary and sufficient* condition for the extension of any total preorder $\leqslant$ to an epistemic entrenchment $\leq$ on $Cn(B)$. Hence ensconcement orderings are always extensible to epistemic entrenchments; Williams in [100, 99], provided an explicit construction of such an extension.

In particular, Williams starts by defining the notion of a *cut* of an ensconcement ordering $\leqslant$ with respect to a sentence φ as follows:

(Cut) $$cut(\varphi) = \begin{cases} \{\psi \in B : \{\chi \in B : \psi \leqslant \chi\} \nvdash \varphi\} & \text{if } \nvdash \varphi \\ \emptyset & \text{otherwise} \end{cases}$$

Using the notion of a a cut, Williams then proceeds to generate a binary relation $\leq$ over the entire language from a given ensconcement ordering $\leqslant$, by means of the following condition:

(EN1) For any $\varphi, \psi \in L$, $\varphi \leq \psi$ iff $cut(\psi) \subseteq cut(\varphi)$.

It turns out that the binary relation $\leq$ so defined is indeed an epistemic entrenchment:

Theorem 10 (Williams [98, 99, 100]). *Let B be a belief base and $\leqslant$ an ensconcement ordering on B. The binary relation $\leq$ generated from $\leqslant$ by means of (EN1) is an epistemic entrenchment related to $Cn(B)$ (i.e. it satisfies the axioms (EE1) - (EE5)).*

From Theorem 10 and (E*) (Section 8.3.4) it follows immediately that the function $*$ defined by condition (EN2) below is an AGM revision function (i.e. it satisfies the postulates $(K*1)$-$(K*8)$).

(EN2) $\psi \in Cn(B) * \varphi$ iff $cut(\varphi \to \psi) \subset cut(\varphi \to \neg\psi)$ or $\vdash \neg\varphi$.

In fact, it turns out that the converse is also true; i.e. any AGM revision function can be constructed from some ensconcement ordering by means of (EN2). Hence the belief base change scheme produced from ensconcement orderings and (EN2) is as expressive as any of the constructive models discussed in Section 8.3, with the additional bonus of being generated from finite structures (in principle). This however is not the only virtue of ensconcement orderings; combined with condition (EN3) below, they produce a very attractive belief base change operator $\circledast$:

(EN3) $B \circledast \varphi = cut(\neg\varphi) \cup \{\varphi\}$.

Notice that, as expected from belief base change operators, the outcome of $\circledast$ is (typically) not a theory but rather a *theory base*. What makes $\circledast$ such an attractive belief base change operator is that, when lifted to the theory level via logical closure, it generates AGM revision functions.

More precisely, let B be a belief base, $\preccurlyeq$ an ensconcement ordering on B, and $\circledast$ the belief base change operator produced from B and $\preccurlyeq$ via (EN3). The function $*$ defined as $Cn(B) * \varphi = Cn(B \circledast \varphi)$ is called an *ensconcement-generated revision function*.[22]

Theorem 11 (Williams [99, 100]). *The class of ensconcement-generated revision functions coincides with the class of AGM revision functions.*

We conclude this section by noting that, in principle, the computational complexity of (propositional) belief revision is NP-hard (typically at the second level of the polynomial hierarchy). For an excellent survey on computational complexity results for belief revision, see [70].

8.5 Multiple Belief Change

From belief base revision we will now move to the other end of the spectrum and examine the body of work in *multiple* belief change. Here, not only is the initial belief set K infinite (since it is closed under logical implication), but it can also be revised by an *infinite* set of sentences. The process of rationally revising K by a (possibly infinite) set of sentences Γ is called *multiple revision*. Similarly, rationally contracting K by a (possibly infinite) Γ is called *multiple contraction*.

Extending the AGM paradigm to include multiple revision and contraction is not as straightforward as it may first appear. Subtleties introduced by infinity need to be treated with care if the connections within the AGM paradigm between postulates and constructive models are to be preserved.

8.5.1 Multiple Revision

As it might be expected, multiple revision is modeled as a function $\oplus$ mapping a theory K and a (possibly infinite) set of sentences Γ, to a new theory $K \oplus \Gamma$. To contrast multiple revision functions with the revision functions discussed so far (whose input are sentences), we shall often call the latter *sentence* revision functions.

Lindstrom, [59], proposed the following generalization of the AGM postulates for multiple revision:[23]

($K \oplus 1$) $K \oplus \Gamma$ is a theory of L.

($K \oplus 2$) $\Gamma \subseteq K \oplus \Gamma$.

($K \oplus 3$) $K \oplus \Gamma \subseteq K + \Gamma$.

[22] It turns out that this ensconcement-generated revision function $*$ has yet another interesting property. It can be constructed from $\preccurlyeq$ following another route: $\preccurlyeq$ gives rise to an epistemic entrenchment $\leq$ by means of (EN1), which in turn produces a revision function by means of (E*), which turns out to be identical with $*$.

[23] There are in fact some subtle differences between the definition of $\oplus$ presented herein and the one given by Lindstrom in [59], which however are only superficial; the essence remains the same.

($K \oplus 4$) If $K \cup \Gamma$ is consistent then $K + \Gamma \subseteq K \oplus \Gamma$.

($K \oplus 5$) If Γ is consistent then $K \oplus \Gamma$ is also consistent.

($K \oplus 6$) If $Cn(\Gamma) = Cn(\Delta)$ then $K \oplus \Gamma = K \oplus \Delta$.

($K \oplus 7$) $K \oplus (\Gamma \cup \Delta) \subseteq (K \oplus \Gamma) + \Delta$.

($K \oplus 8$) If $(K \oplus \Gamma) \cup \Delta$ is consistent then $(K \oplus \Gamma) + \Delta \subseteq K \oplus (\Gamma \cup \Delta)$.

It is not hard to verify that ($K \oplus 1$) - ($K \oplus 8$) are indeed generalizations of ($K * 1$) - ($K * 8$), in the sense that multiple revision collapses to sentence revision whenever the input set Γ is a singleton. To put it more formally, if $\oplus$ satisfies ($K \oplus 1$) - ($K \oplus 8$) then the function $* : \mathbb{K}_L \times L \mapsto \mathbb{K}_L$ defined as $K * \varphi = K \oplus \{\varphi\}$, satisfies the AGM postulates ($K * 1$) - ($K * 8$).

In [77, 80] it was also shown that multiple revision can be constructed from systems of spheres, pretty much in the same way that its sentence counterpart is constructed. More precisely, let K be a theory and S a system of spheres centered on $[K]$. From S a multiple revision function $\oplus$ can be produced as follows:

(S$\oplus$) $$K \oplus \Gamma = \begin{cases} \cap(c(\Gamma) \cap [\Gamma]) & \text{if } [\Gamma] \neq \emptyset \\ L & \text{otherwise} \end{cases}$$

Condition (S$\oplus$) is a straightforward generalization of (S*) and has the same intuitive interpretation: to revise a theory K by a (consistent) set of sentences Γ, pick the most plausible Γ-worlds and define $K \oplus \Gamma$ to be the theory corresponding to those worlds.

Yet, not every system of spheres is good enough to produce a multiple revision function. Two additional constraints, named (SM) and (SD), are needed that are presented below. First however one more definition: we shall say that a set V of consistent complete theories is *elementary* iff $V = [\cap V]$.[24]

(SM) For every nonempty consistent set of sentences Γ, there exists a smallest sphere in S intersecting $[\Gamma]$.

(SD) For every nonempty $\Gamma \subseteq L$, if there is a smallest sphere $c(\Gamma)$ in S intersecting $[\Gamma]$, then $c(\Gamma) \cap [\Gamma]$ is *elementary*.

A system of spheres S which on top of (S1) - (S4) satisfies (SM) and (SD) is called *well-ranked* .

The motivation for condition (SM) should be clear. Like (S4) (to which (SM) is a generalization) condition (SM) guarantees the existence of minimal Γ-worlds (for any consistent Γ), through which the revised theory $K \oplus \Gamma$ is produced.

What may not be clear is the need for condition (SD). It can be shown that conditions (S1) - (S4) do not suffice to guarantee that all spheres in an arbitrary system of spheres are elementary.[25] Condition (SD) requires that at the very least, whenever a non-elementary sphere V *minimally* intersects $[\Gamma]$, the set $V \cap [\Gamma]$ is elementary.

[24] In classical Model Theory, the term "elementary" refers to a class of *models* rather than a set of consistent complete theories (see [11]). Yet, since in this context consistent complete theories play the role of possible worlds, this slight abuse of terminology can be tolerated.

[25] If that was the case, (SD) would had been vacuous since the intersection of any two elementary sets of consistent complete theories is also elementary.

Condition (SD) is a technical one necessitated by the possibility of an infinite input Γ (see [80] for details). Fortunately however (SD) and (SM) are the only additional conditions necessary to elevate the connection between revision functions and systems of spheres to the infinite case:

Theorem 12 (Peppas [77, 80]). *Let K be a theory of L. If S is a well ranked system of spheres centered on $[K]$ then the function $\oplus$ induced from S by means of (S$\oplus$) satisfies the postulates (K$\oplus$1) - (K$\oplus$8). Conversely, for any function $\oplus : \mathbb{K}_L \times 2^L \mapsto \mathbb{K}_L$ that satisfies the postulates (K$\oplus$1) - (K$\oplus$8), there exists a well ranked system of spheres S centered on $[K]$ such that (S$\oplus$) holds for all $\Gamma \subseteq L$.*

Apart from the above systems-of-spheres construction for multiple revision, Zhang and Foo, [108], also lifted the epistemic-entrenchment-based construction to the infinite case.

More precisely, Zhang and Foo start by introducing a variant of an epistemic entrenchment called a *nicely ordered partition*. Loosely speaking, a nicely ordered partition is equivalent to an *inductive epistemic entrenchment* ; i.e. an epistemic entrenchment $\leq$ which (in addition to satisfying (EE1) - (EE5)) is such that every nonempty set of sentences Γ has a minimal element with respect to $\leq$.[26] From a nicely ordered partition one can construct a multiple revision function, pretty much in the same way that a sentence revision function is produced from an epistemic entrenchment.[27] Zhang and Foo prove that the family of multiple revision functions so constructed is *almost* the same as the class of functions satisfying the postulates $(K \oplus 1)$ - $(K \oplus 8)$; for an *exact* match between the two an extra postulate is needed (see [108] for details).

We conclude this section with a result about the possibility of reducing multiple revision to sentence revision.

We have already seen that when Γ is a singleton, the multiple revision of K by Γ is the same as the sentence revision of K by the sentence in Γ. Similarly, one can easily show that when Γ is finite, $K \oplus \Gamma = K * \wedge\Gamma$, where $\wedge\Gamma$ is defined as the conjunction of all sentences in Γ. But what happens when Γ is infinite? Is there still some way of reducing multiple revision to sentence revision? Consider the following condition:

$(K \oplus F)$ $\quad K \oplus \Gamma = \bigcap\{((K * \wedge\Delta) + \Gamma) : \Delta \text{ is a finite subset of } \Gamma\}$.

According to condition $(K \oplus F)$, to reduce multiple revision to sentence revisions when the input Γ is infinite, one should proceed as follows: firstly, the initial theory K is revised by every finite conjunction $\wedge\Delta$ of sentences in Γ, then each such revised theory $K * \wedge\Delta$ is expanded by Γ, and finally all expanded theories $(K * \wedge\Delta) + \Gamma$ are intersected.

Let us call a multiple revision function $\oplus$ that can be reduced to sentence revision by means of $(K \oplus F)$ *sentence-reducible* at K. In [80], a precise characterization of sentence reducible functions was given in terms of the systems of spheres that correspond to them. More precisely, consider the condition (SF) below regarding a system of sphere S centered on $[K]$:

[26] A sentence φ is minimal in Γ with respect to $\leq$ iff $\varphi \in \Gamma$ and for all $\psi \in \Gamma$, if $\psi \leq \varphi$ then $\varphi \leq \psi$. Notice that inductiveness is weaker than the better known property of *well-orderedness*; the latter requires the minimal element in each nonempty set Γ to be *unique*.

[27] A synonym for multiple revision is *infinitary* revision. In fact this is the term used by Zhang and Foo in [108].

(SF) For every $H \subseteq S$, $\bigcup H$ is elementary.

According to (SF), for any collection H of spheres in S, the set resulting from the union of the spheres in H is elementary.[28] The theorem below shows that the multiple revision functions produced by well-ranked systems of spheres satisfying (SF), are precisely those that are sentence reducible at K:

Theorem 13 (Peppas [80]). *Let K be a theory of L and $\oplus$ a multiple revision function. The function $\oplus$ is sentence-reducible at K iff there exists a well-ranked system of spheres S centered on $[K]$ that induces $\oplus$ by means of ($S\oplus$) and that satisfies (SF).*

8.5.2 Multiple Contraction

Unlike multiple revision where the AGM postulates have an obvious generalization, in multiple contraction things are not as clear. The reason is that there are (at least) three different ways of interpreting multiple contraction, giving rise to three different operators called *package contraction*, *choice contraction*, and *set contraction*. The first two are due to Fuhrmann and Hansson, [28], while the third has been introduced and analyzed by Zhang and Foo [106, 108].

Given a theory K and a (possibly infinite) set of sentences Γ, package contraction removes all (non-tautological) sentences in Γ from K. Choice contraction on the other hand is more liberal; it only requires that *some* (but not necessarily all) of the sentences in Γ are removed from K. Fuhrmann and Hansson, [28], have proposed natural generalizations of the AGM postulates for both package contraction and choice contraction. They also obtained preliminary representation results relating their postulates with constructive methods for package and choice contraction. These results however are limited to generalizations of the *basic* AGM postulates; they do not include (generalizations of) the supplementary ones. A few years later, Zhang and Foo [106, 108] obtained such fully-fledged representation results for set contraction.

Set contraction is slightly different in spirit from both package and choice contraction. Given a theory K and a set of sentences Γ, the goal with set contraction is not to remove part or the whole of Γ from K, but rather *to make K consistent with Γ*. At first sight this may seem like an entirely new game, but in fact it is not. For example, it can be shown that for any consistent sentence φ, the set contraction of K by $\{\varphi\}$ is the same as the sentence contraction of K by $\neg\varphi$.[29]

Zhang and Foo define set contraction as a function $\ominus$ mapping a theory K and a set of sentences Γ to the theory $K \ominus \Gamma$, that satisfies the following postulates:

($K \ominus 1$) $K \ominus \Gamma$ is a theory of L.

($K \ominus 2$) $K \ominus \Gamma \subseteq K$.

($K \ominus 3$) If $K \cup \Gamma$ is consistent then $K \ominus \Gamma = K$.

[28] It is not hard to see that (SF) entails (SD). Simply notice that from (SF) it follows that all spheres in S are elementary, and consequently, the intersection of $[\Gamma]$ (for any set of sentences Γ) with *any* sphere of S is also elementary.

[29] Similarly to sentence revision, in this section we shall use the term *sentence contraction* to refer to the original contraction functions whose inputs are sentences rather than sets of sentences..

($K \ominus 4$) If Γ is consistent then $\Gamma \cup (K \ominus \Gamma)$ is consistent.

($K \ominus 5$) If $\varphi \in K$ and $\Gamma \vdash \neg\varphi$ then $K \subseteq (K \ominus \Gamma) + \varphi$.

($K \ominus 6$) If $Cn(\Gamma) = Cn(\Delta)$ then $K \ominus \Gamma = K \ominus \Delta$.

($K \ominus 7$) If $\Gamma \subseteq \Delta$ then $K \ominus \Delta \subseteq (K \ominus \Gamma) + \Delta$.

($K \ominus 8$) If $\Gamma \subseteq \Delta$ and $\Delta \cup (K \ominus \Gamma)$ is consistent, then $K \ominus \Gamma \subseteq K \ominus \Delta$.

Given the different aims of sentence and set contraction, it should not be surprising that ($K \ominus 1$)-($K \ominus 8$) are not a straightforward generalization of ($K \dot{-} 1$)-($K \dot{-} 8$). For the same reason the generalized version of Levi and Harper Identities presented below, [108], are slightly different from what might be expected:

$$K \oplus \Gamma = (K \ominus \Gamma) + \Gamma \qquad \text{(Generalized Levi Identity)}$$

$$K \ominus \Gamma = (K \oplus \Gamma) \cap K \quad \text{(Generalized Harper Identity)}$$

Zhang provides support for the set contraction postulates by lifting the validity of the Harper and Levi Identities to the infinite case:

Theorem 14 (Zhang [106]). *Let $\ominus$ be a set contraction function satisfying the postulates ($K \ominus 1$) - ($K \ominus 8$). Then the function $\oplus$ produced from $\ominus$ by means of the Generalized Levi Identity, satisfies the postulates ($K \oplus 1$) - ($K \oplus 8$).*

Theorem 15 (Zhang [106]). *Let $\oplus$ be a multiple revision function that satisfies the postulates ($K \oplus 1$) - ($K \oplus 8$). Then the function $\ominus$ produced from $\oplus$ by means of the Generalized Harper Identity, satisfies the postulates ($K \ominus 1$) - ($K \ominus 8$).*

Apart from the above results, Zhang and Foo reproduced for set contraction Gardenfors' and Makinson's epistemic-entrenchment construction. More precisely, Zhang and Foo presented a construction for set contractions based on nicely ordered partitions, and proved that the family of set contractions so defined is *almost* the same (in fact is a proper subset of) the family of functions satisfying the postulates ($K \ominus 1$) - ($K \ominus 8$); once again an exact match can be obtained if an extra postulate is added to ($K \ominus 1$) - ($K \ominus 8$).

8.6 Iterated Revision

We shall now turn to one of the main shortcomings of the early AGM paradigm: its lack of any guidelines for *iterated revision*.

Consider a theory K coupled with a structure encoding extra-logical information relevant to belief change, say a system of spheres S centered on $[K]$. Suppose that we now receive new information φ, such that $\varphi \notin K$, thus leading us to the new theory $K * \varphi$. Notice that $K * \varphi$ is *fully determined* by K, φ, and S, and moreover, as Grove has shown, the transition from the old to the new belief set satisfies the AGM postulates. But what if at this point we receive further evidence ψ, which, to make the case interesting, is inconsistent with $K * \varphi$ (but not self-contradictory; i.e. $\nvdash \neg\psi$). Can we produce $K * \varphi * \psi$ from what we already know (i.e. K, S, φ, and ψ)? The answer, perhaps surprisingly, is *no*. The reason

is that at $K * \varphi$ we no longer have the additional structure necessary for belief revision; i.e. we don't know what the "appropriate" system of spheres for $K * \varphi$ is, and without that there is very little we can infer about $K * \varphi * \psi$.[30] But why not simply keep the original system of spheres S? For one thing, this would violate condition (S2) which requires that the minimal worlds (i.e. the worlds in the smallest sphere) are $(K * \varphi)$-worlds (and not K-worlds as they are in S). We need a new system of spheres S' centered on $[K * \varphi]$ that is in some sense the *rational* offspring of S and φ. Unfortunately the AGM postulates give us no clue about how to produce S'. The AGM paradigm focuses only on one-step belief change; iterated belief change was left unattended.

8.6.1 Iterated Revision with Enriched Epistemic Input

Spohn, [94], was one of the first to address the problem of iterated belief revision, and the elegance of his solution has influenced most of the proposals that followed. This elegance however comes with a price; to produce the new preference structure from the old one, Spohn requires as input not only the new information φ, but also the *degree of firmness* by which the agent accepts the new information. Let us take a closer look at Spohn's solution (to simplify discussion, in this section we shall consider only revision by *consistent* sentences on *consistent* theories).

To start with, Spohn uses a richer structure than a system of spheres to represent the preference information related to a belief set K. He calls this structure an *ordinal conditional function* (OCF). Formally, an OCF κ is a function from the set $\mathbb{M}_L$ of possible worlds to the class of ordinals such that at least one world is assigned the ordinal 0. Intuitively, κ assigns a plausibility grading to possible worlds: the larger $\kappa(r)$ is for some world r, the less plausible r is.[31] This plausibility grading can easily be extended to sentences: for any consistent sentence φ, we define $\kappa(\varphi)$ to be the κ-value of the most plausible φ-world; in symbols, $\kappa(\varphi) = \min(\{\kappa(r) : r \in [\varphi]\})$.

Clearly, the most plausible worlds of all are those whose κ-value is zero. These worlds define the belief set that κ is related to. In particular, we shall say that the belief set K is related to the OCF κ iff $K = \bigcap\{r \in \mathbb{M}_L\colon \kappa(r) = 0\}$. Given a theory K and an OCF κ related to it, Spohn can produce the revision of K by any sentence φ, *as well as* the new ordinal conditional function related to $K * \varphi$. The catch is, as mentioned earlier, that apart from φ, its degree of firmness d is also needed as input. The new OCF produced from κ and the pair $\langle \varphi, d\rangle$ is denoted $\kappa * \langle \varphi, d\rangle$ and it is defined as follows:[32]

$$\text{(CON)} \qquad \kappa * \langle \varphi, d\rangle(r) = \begin{cases} \kappa(r) - \kappa(\varphi) & \text{if } r \in [\varphi] \\ \\ \kappa(r) - \kappa(\neg\varphi) + d & \text{otherwise} \end{cases}$$

Essentially condition (CON) works as follows. Starting with κ, all φ-worlds are shifted "downwards" against all $\neg\varphi$-worlds until the most plausible of them hit the bottom of the rank; moreover, all $\neg\varphi$-worlds are shifted "upwards" until the most plausible of them are

[30] In fact, all we can deduce is that $K * \varphi * \psi$ is a theory containing ψ.

[31] In this sense an ordinal conditional function κ is quite similar to a system of spheres S: both are formal devices for ranking possible worlds in terms of plausibility. However κ not only tells us which of any two worlds is more plausible; it also tells us by *how much* is one world more plausible than the other.

[32] The left subtraction of two ordinals α, β such that $\alpha \geq \beta$, is defined as the unique ordinal γ such that $\alpha = \beta + \gamma$.

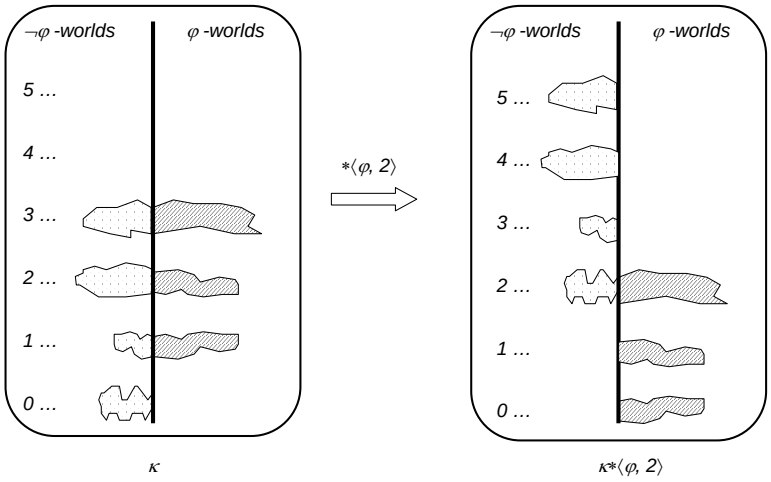

Figure 8.3: Spohn's Conditionalization

at distance d from the bottom (see Figure 8.3). Spohn calls this process *conditionalization* (more precisely, the $\langle\varphi, d\rangle$-*conditionalization* of κ) and argues that is the right process for revising OCFs.

Conditionalization is indeed intuitively appealing and has many nice formal properties, including compliance with the AGM postulates[33] (see [94, 31, 101]). Moreover notice that the restriction of κ to $[\varphi]$ and to $[\neg\varphi]$ remains unchanged during conditionalization, hence in this sense the principle of minimal change is observed not only for transitions between belief sets, but also for their associated OCFs.

There are however other ways of interpreting minimal change in the context of iterated revision. Williams in [101] proposes the process of *adjustment* as an alternative to conditionalization. Given an OCF κ, Williams defines the $\langle\varphi, d\rangle$-*adjustment of* κ, which we denote by $\kappa \circ \langle\varphi, d\rangle$, as follows:

$$\text{(ADJ)} \qquad \kappa \circ \langle\varphi, d\rangle(r) = \begin{cases} 0 & \text{if } r \in [\varphi],\ d > 0, \text{ and } \kappa(r) = \kappa(\varphi) \\ d & \text{if } r \in [\neg\varphi], \text{ and } \kappa(r) = \kappa(\neg\varphi) \text{ or } \kappa(r) \leq d \\ \kappa(r) & \text{otherwise} \end{cases}$$

Adjustment minimizes changes to the grades of possible worlds in *absolute* terms. To see this, notice that in the principal case where $\kappa(\varphi) > 0$ and $d > 0$,[34] the only φ-worlds that change grades are the most plausible ones (wrt κ), whose grade becomes zero. Moreover, the only $\neg\varphi$-worlds that change grades are those with grades smaller that d, or, if no such world exists, the minimal $\neg\varphi$-worlds whose grade becomes d. Like conditionalization, adjustment satisfies all AGM postulates for revision.

[33] That is, given an OCF κ and any $d > 0$, the function $*$ defined as $K * \varphi = \bigcap\{r \in \mathbb{M}_L : \kappa * \langle\varphi, d\rangle(r) = 0\}$ satisfies the AGM postulates $(K*1)$ - $(K*8)$.

[34] This is the case where the new information φ contradicts the original belief set (since $\kappa(\neg\varphi) > 0$, the agent originally believes $\neg\varphi$).

The entire apparatus of OCFs and their dynamics (conditionalization or adjustment) can be reproduced using sentences rather than possible worlds as building blocks. To this end, Williams, [101], defined the notion of *ordinal epistemic entrenchments functions* (OEF) as a special mapping from sentences to ordinals, intended to encode the resistance of sentences to change: the higher the ordinal assigned to sentence, the higher the resistance of the sentence. As the name suggests, an OEF is an enriched version of an epistemic entrenchment (in the same way that an OCF is an enriched version of a system of spheres). Williams formulated the counterparts of conditionalization and adjustment for OEF and proved their equivalence with the corresponding operation on OCFs.

In [67], Nayak took this line of work one step further. Using the original epistemic entrenchment model to encode sentences resistance to change, he considers the general problem of epistemic entrenchment dynamics. The novelty in Nayak's approach is that the epistemic input is no longer a simple sentence as in AGM, or even a sentence coupled with a degree of firmness as in OCF dynamics, but rather another epistemic entrenchment; i.e. an initial epistemic entrenchment $\leq$ is revised by another epistemic entrenchment $\leq'$, producing a new epistemic entrenchment $\leq * \leq'$. Notice that because of (EE4) (see Section 8.3.4), an epistemic entrenchment uniquely determines the belief set it relates to; we shall call this set the *content* of an epistemic entrenchment. Hence epistemic entrenchment revision should be interpreted as follows. The initial epistemic entrenchment $\leq$ represents both the original belief set K (defined as its content) as well as the preference structure related to K. The input $\leq'$ represents prioritized evidence: the content K' of $\leq'$ describes the new information, while the ordering on K' is related (but not identical) to the relative strength of acceptance of the sentences in K'. Finally, $\leq * \leq'$ encodes both the posterior belief set as well as the preference structure associated with it.

Nayak proposes a particular construction for epistemic entrenchment dynamics and shows that the induced operator satisfies (a generalized version of) the AGM postulates for revision. Compared to Williams' OEFs dynamics, Nayak's work is closer to the AGM tradition (both use epistemic entrenchments to represent belief states and plausibility is represented in relative rather than absolute terms). On the other hand however, when it comes to the modeling the epistemic input, Nayak departs even further than Williams from the AGM paradigm; an epistemic entrenchment (used by Nayak) is a much more complex structure than a weighted sentence (used by Williams), which in turn is richer than a simple sentence (used in the original AGM paradigm).

8.6.2 Iterated Revision with Simple Epistemic Input

This raises the question of whether a solution to iterated revision can be produced using only the apparatus of the original AGM framework; that is, using epistemic entrenchments (or systems of spheres or selection functions) to model belief states, and simple sentences to model epistemic input.

One of the most influential proposals to this end is the work of Darwiche and Pearl ("DP" for short), [14]. The first important feature of this work is that, contrary to the original approach of Alchourron, Gardenfors and Makinson (but similarly to Spohn [94], Williams [101], and Nayak [67]), revision functions operate on *belief states*, not on *belief sets*. In the present context a belief state (also referred to as an *epistemic state*) is defined as a belief set coupled with a structure that encodes relative plausibility (*e.g.*, an epistemic entrenchment, a system of spheres, etc.). Clearly a belief state is a richer model that a

belief set. Hence it could well be the case that two belief states agree on their belief content (i.e. their belief sets), but behave differently under revision because of differences in their preference structures. For ease of presentation, and although this is not required by Darwiche and Pearl, in the rest of this section we shall identify belief states with systems of spheres; note that given a system of spheres S we can easily retrieve its belief content – simply notice that $c(\top)$ is the smallest sphere of S and therefore $\cap c(\top)$ is the belief set associated with S.[35] We shall often abuse notation and write for a sentence φ that $\varphi \in S$ instead of $\varphi \in \cap c(\top)$.

With these conventions, $*$ becomes a function that maps a system of spheres S and a sentence φ, to a new system of spheres $S * \varphi$. Darwiche and Pearl reformulated the AGM postulates accordingly to reflect the shift from belief sets to belief states. They also proposed the following four additional postulates to regulate iterated revisions:[36]

(DP1) If $\varphi \vdash \chi$ then $(S * \chi) * \varphi = S * \varphi$.

(DP2) If $\varphi \vdash \neg\chi$ then $(S * \chi) * \varphi = S * \varphi$.

(DP3) If $\chi \in S * \varphi$ then $\chi \in (S * \chi) * \varphi$.

(DP4) If $\neg\chi \notin S * \varphi$ then $\neg\chi \notin (S * \chi) * \varphi$.

Postulate (DP1) says that if the subsequent evidence φ is logically stronger than the initial evidence χ then φ overrides whatever changes χ may have made. (DP2) says that if two contradictory pieces of evidence arrive sequentially one after the other, it is the later that will prevail. (DP3) says that if revising S by φ causes χ to be accepted in the new belief state, then revising first by χ and then by φ can not possibly block the acceptance of χ. Finally, (DP4) captures the intuition that "no evidence can contribute to its own demise" [14]; if the revision of S by φ does not cause the acceptance of $\neg\chi$, then surely this should still be the case if S is first revised by χ before revised by φ.

Apart from their simplicity and intuitive appeal, postulates (DP1) - (DP4) also have a nice characterization in terms of systems-of-spheres dynamics. First however some more notation. Let S be a system of spheres and r, r' any two possible worlds. We shall write $r \sqsubseteq_S r'$ iff every sphere in S that contains r' also contains r (i.e. r is at least as plausible as r' in S); we shall write $r \sqsubset_S r'$ iff there is a sphere in S that contains r but not r' (i.e. r is strictly more plausible than r' wrt S). It is not hard to verify that $\sqsubseteq_S$ is a total preorder in $\mathbb{M}_L$ with the center of S as its minimal elements, while $\sqsubset_S$ is the strict part of $\sqsubseteq_S$. Darwiche and Pearl proved that there is a one-to-one correspondence between (DP1) - (DP4) and the following constraints on system-of-spheres dynamics:

(DPS1) If $r, r' \in [\varphi]$ then $r \sqsubseteq_{S*\phi} r'$ iff $r \sqsubseteq_S r'$.

(DPS2) If $r, r' \in [\neg\varphi]$ then $r \sqsubseteq_{S*\phi} r'$ iff $r \sqsubseteq_S r'$.

(DPS3) If $r \in [\varphi]$ and $r' \in [\neg\varphi]$ then $r \sqsubset_S r'$ entails $r \sqsubset_{S*\phi} r'$.

(DPS4) If $r \in [\varphi]$ and $r' \in [\neg\varphi]$ then $r \sqsubseteq_S r'$ entails $r \sqsubseteq_{S*\phi} r'$.

[35] Recall that for any sentence ψ, $c(\psi)$ denotes the smallest sphere in S intersecting $[\psi]$.

[36] The postulates are expressed in terms of the Katsuno and Mendelzon formalism [49]; herein however we have rephrased them in the AGM terminology.

Theorem 16 (Darwiche and Pearl [14]). *Let S be a belief state and $*$ a revision function satisfying the (DP-modified) AGM postulates. Then $*$ satisfies (DP1) - (DP4) iff it satisfies (DPS1) - (DPS4) respectively.*

In a way, Darwiche and Pearl were forced to make the shift from belief sets to belief states, for otherwise (DP2) would have conflicted with the AGM postulates (see [25, 68]).[37] Nayak, Pagnucco, and Peppas, [68], proposed another way to reconcile (DP2) with the AGM postulates that does not require moving away from belief sets. It does however require two other changes to the original formulation of belief revision. Firstly, $*$ is defined as a *unary* rather than a binary function, mapping sentences to theories. That is, each theory K is assigned its own revision function which for any sentence φ produces the revision of K by φ. We shall denote the unary revision function assigned to K by $*_K$ and the result of revising K by φ as $*_K(\varphi)$. This change in notation will serve as a reminder of the unary nature of revision functions adopted in [68]. Notice that this reformulation of revision functions does not require any modification to the AGM postulates, since all of them refer only to a single theory K.

The second modification to revision functions proposed in [68] is that they are *dynamic*; i.e. they could change as new evidence arrives. The implications of this modification are best illustrated in the following scenario. Consider an agent whose belief set at time t_0 is K_0, and who receives a sequence of new evidence $\varphi_1, \varphi_2, \ldots, \varphi_n$ and performs the corresponding n revisions that take him at time t_n to the belief set K_n. Suppose now that it so happens that $K_n = K_0$; i.e. after incorporating all the new evidence, the agent ended up with the theory she started with. Because of the dynamic nature of revision functions in [68], it is possible that the revision function assigned to K_0 at time t_0 is different from the one assigned to it at time t_n. Hence although the evidence $\varphi_1, \varphi_2, \ldots, \varphi_n$ did not change the agent's beliefs, they did alter her attitude towards new epistemic input.

These two modifications to revision functions take care of the inconsistency between (DP2) and the AGM postulates when applied to belief sets. There is however another problem with (DP1) - (DP4) identified in [68]. Nayak *et al.* argue that (DP1) - (DP4) are also too permissive; i.e. there are revision functions that comply with both the AGM and DP postulates and nevertheless lead to counter-intuitive results. Moreover, an earlier proposal by Boutilier, [7, 9], which strengthens (DP1) - (DP4) still fails to block the unintended revision functions (and introduces some problems of its own – see [14]). Hence Nayak *et al.* proposed the following addition to (DP1) - (DP4) instead, called the *Conjunction Postulate* :

(CNJ) If $\chi \wedge \varphi \nvdash \bot$, then $*^{\chi}_{*_K(\chi)}(\varphi) = *_K(\chi \wedge \varphi)$.

Some comments on the notation in (CNJ) are in order. As usual, K denotes the initial belief set, and $*_K$ the unary revision function associated with it. When K is revised by a sentence χ, a new theory $*_K(\chi)$ is produced. This however is not the only outcome of the revision of K by χ; a new revision function associated with $*_K(\chi)$ is also produced. This new revision function is denoted in (CNJ) by $*^{\chi}_{*_K(\chi)}$. The need for the superscript χ is due to the dynamic nature of $*$ (as discussed earlier, along a sequence of revisions, the same belief

[37] Although it should be noted that Darwiche and Pearl argue that this shift is not necessitated by technical reasons alone; conceptual considerations also point the same way.

set may appear more than once, each time with a different revision function associated to it, depending on the input sequence).

Postulate (CNJ) essentially says that if two pieces of evidence χ and φ are consistent with each other, then it makes no difference whether they arrive sequentially or simultaneously; in both cases the revision of the initial belief set K produces the same theory.

Nayak *et al.* show that (CNJ) is consistent with both AGM and DP postulates, and it blocks the counterexamples known at the time. In fact (CNJ) is strong enough to *uniquely* determine (together with (K*1) - (K*8) and (DP1) - (DP4)) the new revision function $*^{\chi}_{*_K(\chi)}$. A construction of this new revision function from $*_K$ and χ is given is [68].

Yet, some authors have argued, [109, 47], that while (DP1) - (DP4) are too permissive, the addition of (CNJ) is too radical (at least in some cases). Accordingly, Jin and Thielscher proposed a weakening of (CNJ), which they call the *Independence postulate* [47]. The Independence postulate is formulated within the DP framework; that is, it assumes that belief states rather than belief sets are the primary objects of change:

(Ind) If $\neg\chi \notin S * \varphi$ then $\chi \in (S * \chi) * \varphi$.

The Independence postulate, apart from performing well in indicative examples (see [47]), also has a nice characterization in terms of system of spheres dynamics:

(IndR) If $r \in [\varphi]$ and $r' \in [\neg\varphi]$ then $r \sqsubseteq_S r'$ entails $r \sqsubset_{S*\phi} r'$.

Theorem 17 (Jin and Thielscher [47]). *Let S be a belief state and $*$ a revision function satisfying the (DP-modified) AGM postulates. Then $*$ satisfies (Ind) iff it satisfies (IndR).*

The Independence postulate can be shown to be weaker than (CNJ) and in view of Theorems 16, 17, it is clearly stronger than (DP3) and (DP4). Jin and Thielscher show that (Ind) is consistent with the AGM and DP postulates combined.

Other important works on iterated revision are, [6] which proposes a different strengthening of the DP approach, [45] that considers the interaction between iterated revisions and updates (see Section 8.8), [91] that defines belief revision in terms of distances between possible worlds and derives interesting properties for iterated revision, as well as [17, 19, 52, 56, 102].

8.7 Non-Prioritized Revision

A fundamental assumption in our discussion on belief revision so far has been that the new information the agent receives comes from a reliable source and therefore it should be accepted without second thoughts, no matter how implausible it may seem given the agent's initial beliefs.

This assumption is of course a rather strong one and a number of researchers have proposed alterations to the AGM paradigm in order to lift it. The resulting new type of belief change is called *non-prioritized belief revision*. Depending on a number of parameters, a non-prioritized belief revision operator may fully accept, partially accept, or even totally reject the new information.

One of the earliest proposals for non-prioritized belief revision is Makinson's *screened revision* [63]. The basic idea here is that the fate of the new information depends on its disposition towards a set of *core* beliefs. More precisely, a set of beliefs A is given *a priori*

that is considered to be immune to contraction. This set A together with the agent's initial belief set K determine the set of core beliefs defined as $A \cap K$. If the new information φ is inconsistent with $A \cap K$ then it is rejected as implausible; otherwise φ is accepted and K is revised accordingly. In the latter case however, the revision of K by φ should be such that none of the core beliefs are removed. Makinson denotes by $*_A$ an AGM revision function that satisfies the following condition:

(CR) If φ is consistent with $A \cap K$ then $A \cap K \subseteq K *_A \varphi$.

With the aid of $*_A$ Makinson defines a screened revision operator, denoted by $\#_A$, as follows:

$$\text{(SC)} \qquad K\#_A\varphi = \begin{cases} K *_A \varphi & \text{if } \varphi \text{ is consistent with } A \cap K \\ K & \text{otherwise} \end{cases}$$

Makinson then proceeds to introduce a more flexible variant of screened revision which he calls *relationally screened revision* . The main new feature of this variant is that the core beliefs are not fixed but they depend on the new information φ. In particular, instead of A, a binary relation $\lessdot$ is given *a priori* representing comparative credibility; i.e. if $\chi \lessdot \psi$ then χ is less credible than ψ. Then for input φ the set of core beliefs is defined as $\{\chi : \varphi \lessdot \chi\} \cap K$. Accordingly, the condition that defines a relationally screened revision, denoted $\#_<$, is the following:

$$\text{(RSC)} \qquad K\#_< \varphi = \begin{cases} K *_{\{\chi:\ \varphi < \chi\}} \varphi & \text{if } \varphi \text{ is consistent with } \{\chi :\ \varphi \lessdot \chi\} \cap K \\ K & \text{otherwise} \end{cases}$$

It is not hard to verify that screened revision is a special case of relationally screened revision. Simply set, for a given A, the binary relation $\lessdot$ to be $L \times A$.[38]

Hansson et al., [43], proposed a different approach to non-prioritized revision called *credibility-limited revision* . According to this approach, a set $\mathbb{C}$ of *credible sentences* is given *a priori* and any new information φ is accepted only if it belongs to $\mathbb{C}$:

$$\text{(CL)} \qquad K \circledcirc \varphi = \begin{cases} K * \varphi & \text{if } \varphi \in \mathbb{C} \\ K & \text{otherwise} \end{cases}$$

In the above condition $\circledcirc$ is the new credibility-limited revision operator and $*$ is an AGM revision function.[39]

Depending on the constraints that one places on $\mathbb{C}$ and $*$, a number of interesting results can be obtained for the induced operator $\circledcirc$. In particular, assume that $\mathbb{C}$ can be generated from a subset A of the initial belief set K by means of the following condition:

[38] If $\lessdot$ is required to be a strict order (i.e. transitive and antisymmetric), then things are not as simple but it is still possible (in principle) to reduce screened revision to relationally screened revision.

[39] To be precise, in [43] the function $*$ does not have to satisfy the AGM postulates; when it does, the induced operator $\circledcirc$ is called a *credibility-limited AGM revision*. Herein we focus only on such operators and therefore, for the sake of readability, we have dropped the AGM advert from the title of $\circledcirc$.

(CCL) $\varphi \in \mathbb{C}$ iff $A \nvdash \neg\varphi$

The credibility-limited revision operator induced from such a $\mathbb{C}$ is called a *core belief revision* operator and can be characterized both axiomatically and constructively (see [43]). Below we briefly review a constructive model of core belief revision based on system of spheres.[40]

Let S be a system of spheres centered on $[K]$ and assume that S contains $[A]$ as one of its spheres. Consider the following construction of $\circledcirc$ (recall that for any consistent sentence φ, $c(\varphi)$ denotes the smallest sphere in S intersecting $[\varphi]$):

$$\text{(S}\circledcirc\text{)} \qquad K \circledcirc \varphi = \begin{cases} \bigcap(c(\varphi) \cap [\varphi]) & \text{if } c(\varphi) \subseteq [A] \\ K & \text{otherwise} \end{cases}$$

Intuitively the sphere $[A]$ circumscribes the set of "entertainable" worlds; any world outside $[A]$ is so implausible that it should never be accepted as a possible state of affairs. Consequently, says condition (S$\circledcirc$), any sentence φ that takes us to the "forbidden land" of non-entertainable worlds (i.e. any sentence φ for which all φ-world are outside $[A]$) should be rejected; otherwise it is business as usual and the next belief set is determined by the minimal φ-worlds. Hansson *et al.* in [43] show that the operators constructed through (S$\circledcirc$) coincide with the family of core belief revision operators.

Both screened revision and credibility-limited revision work in two stages: firstly they check whether the new information φ should be accepted (each with its own decision mechanism) and then, if φ is credible, they revise the initial belief set K by φ. As a result, φ is either accepted in its entirety or not at all; there is no middle ground (such as accepting part of φ). Hansson in [40] proposed non-prioritized belief revision operators that escape this black-and-white attitude towards φ.

The basic idea is the following: add the new information φ to the initial beliefs without checking its credibility and then remove all inconsistencies that may result. Of course in the process of restoring consistency, one may also lose φ. Even so, it may still be possible to keep some *parts* of φ; i.e. non-tautological sentences ψ that follow logically from φ and which were not among the initial beliefs. Hansson calls this operation *semi-revision* and it is clearly more flexible in its treatment of φ than any of the operators discussed so far. It should be noted that semi-revision is defined over *belief bases* rather than belief sets. The extra structure of a belief base is used to guide the restoration of consistency after the addition of φ. Formally the semi-revision of a belief base B by a sentence φ, which we denote by $B \circledcirc \varphi$, is defined as follows:

(SR) $B \circledcirc \varphi = (B \cup \{\varphi\}) \dot{-} \bot$

In the condition above, $\dot{-}$ is a belief base contraction operator, and depending on the constraints one places on $\dot{-}$, different types of semi-revision functions are produced. Of particular interest are the class of semi-revision operators induced from kernel contractions, and the class generated from partial meet belief base contractions; both these classes have been characterized axiomatically in [43].

[40] This constructive model is slightly different from the one discussed in [43] but it is nevertheless equivalent to it.

A totally different approach to non-prioritized belief revision was proposed by Schlechta in [92]. Schlechta's proposal is based on a notion of distance between possible worlds. In this context, the distance between two worlds r' and r'' does not have some numerical value, but it is defined in reference to a third world r. In particular, a ternary relation between worlds is introduced such that whenever it holds between the worlds r, r', and r'', it means that r' is closer to r than is r''. Based on this ternary relation, Schlechta defines the non-prioritized belief revision of K by φ to be the belief set determined by the set of K-worlds and φ-worlds that have minimal distance between them among all pairs of K-worlds and φ-worlds.

Yet another important approach to non-prioritized belief revision can be found in [12, 34], while the process of *extraction* reported in [104] can also be used to this end. See also Hansson's survey on this subject [42].

We shall conclude this section with a quick look at *Belief Merging* which started with a similar agenda to non-prioritized belief revision [82, 58] but quickly developed into a fully-fledged research area of its own addressing much more general and diverse issues.[41]

In Belief Merging one starts with a set of belief bases $B = \{B_1, B_2, \ldots, B_n\}$ (possibly with weights assigned to each B_i or with some other structure expressing relative importance) and has to produce an aggregate belief base $\Delta(B)$ that is in some sense the result of rationally merging all B_i's. What makes the problem non-trivial is that in principle $\bigcup B_i$ is inconsistent, whereas the aggregate belief base $\Delta(B)$ is required to be consistent. Moreover, a set of *integrity constraints IC* is typically given together with B, to which $\Delta(B)$ needs to adhere to.

Most work in Belief Merging can be classified either as *model-based* [82, 58, 51] or *syntax-based* [3, 52, 5]. In the first case $\Delta(B)$ is defined in terms of the *most preferred* models of IC. Preference in turn is defined according to some criterion that depends on B – usually a notion of distance between possible worlds and B with the worlds closest to B being the most preferred.

Syntax-based approaches on the other hand typically select consistent subsets of $\bigcup B$ taking into account the syntax of the belief bases B_i and any additional preference information that might be given.

Recently, S. Konieczny, J. Lang and P. Marquis, [53], developed a unifying framework that can encompass many of the existing merging operators both from the model-based and the syntax-based families.

8.8 Belief Update

In this final section we shall examine a type of belief change that was initially mistaken to be identical with belief revision, but it turns out to be different from it.

Consider the following scenario. Philippa is looking through an open door at a room with a table, a magazine, and a book. One of the two items is on the table and the other on the floor, but because of poor lighting Philippa cannot distinguish which is which. Let us represent by b the proposition that "the book is on the table", and by m the proposition that "the magazine is on the table". Philippa's belief set is then represented by $K = Cn((b \wedge \neg m) \vee (\neg b \wedge m))$. Suppose now that Philippa instructs a robot standing beside her to enter the room and make sure that the book is placed on the floor. The robot will approach

[41] Nevertheless, many would still classify Belief Merging as a sub-area of Belief Revision.

the table and if the book is on the table the robot will place it on the floor; otherwise it will do nothing. In either case the robot will go back to Philippa and report "mission accomplished!".

What would be Philippa's belief set K' after the robot reports back to her that the book is on the floor? Presumably it will be the initial belief set K modified by $\neg b$. Suppose now that we use an AGM revision function to perform the modification. Notice that $\neg b$ is consistent with K, and therefore by $(K * 3)$ - $(K * 4)$, $K * \neg b = K + \neg b = Cn(\neg b \wedge m)$. So according to the AGM paradigm, if the book was initially on the table, putting it on the floor somehow makes the magazine jump onto the table!

This counter-intuitive behavior of AGM revision functions was first observed by Katsuno and Mendelzon in [48] who also proposed a solution to the problem. Their solution does not dismiss (or even alter) the AGM paradigm; it simply carefully defines its range of applicability.

According to Katsuno and Mendelzon, the reason that the AGM postulates fail to produce the right results in the book/magazine example is because they were never meant to deal with these situations in the first place. Belief revision should *only* be used to modify an incomplete or incorrect belief set K in the light of new information φ that was previously inaccessible to the agent. It *should not* be used in cases where an agent needs to bring her belief set K up-to-date with changes in the world that brought about φ; in the latter case a new type of belief change takes place called *belief update*. In a nutshell, the difference between belief revision and belief update is that the former is used when new information φ is received about a static world, and the latter is used when the agent is informed that a change in the world has occurred that brought about φ; in the first case the initial belief set K needs to be modified because it is incorrect or incomplete, whereas in the latter case K is modified because it is out-of-date (it was initially correct but in the meantime changes have occurred in the world).

Following the AGM tradition, Katsuno and Mendelzon characterized the process of belief update (or simply update) in terms of a set of postulates, now known as the KM postulates. Like the AGM postulates, the KM postulates are also motivated by the principle of minimal change. However in this context the notion of minimal change applies to world states, *not* to belief sets; when an agent updates her beliefs in response to a minimal change in the world, her new belief set does not necessarily differ minimally from the original. This is a subtle point that has been the source of some confusion before Winslett (see [105]) and finally Katsuno and Mendelzon set things straight.

For ease of comparison we have rephrased the KM postulates in the tradition of the AGM paradigm:

$(K \diamond 1)$ $K \diamond \varphi$ is a theory of L.

$(K \diamond 2)$ $\varphi \in K \diamond \varphi$.

$(K \diamond 3)$ If $\varphi \in K$ then $K \diamond \varphi = K$.

$(K \diamond 4)$ If K and φ are individually consistent then $K \diamond \varphi$ is consistent.

$(K \diamond 5)$ If $\vdash \varphi \leftrightarrow \psi$ then $K \diamond \varphi = K \diamond \psi$.

$(K \diamond 6)$ $K \diamond (\varphi \wedge \psi) \subseteq (K \diamond \varphi) + \psi$.

($K \diamond 7$) If $\psi \in K \diamond \varphi$ and $\varphi \in K \diamond \psi$ then $K \diamond \varphi = K \diamond \psi$.

($K \diamond 8$) If K is complete then $K \diamond (\varphi \vee \psi) \subseteq Cn((K \diamond \varphi) \cup (K \diamond \psi))$.

($K \diamond 9$) $K \diamond \varphi = \bigcap_{r \in [K]} r \diamond \varphi$.

Postulates ($K \diamond 1$), ($K \diamond 2$), ($K \diamond 5$), and ($K \diamond 6$) are identical with ($K * 1$), ($K * 2$), ($K * 6$) and ($K * 7$) respectively, and need no further explanation. Postulate ($K \diamond 3$) is a restricted version of the postulates ($K * 3$) and ($K * 4$) combined; it says that if the new proposition φ is already in the initial belief set K then updating K by φ changes nothing. Notice however that ($K \diamond 3$) puts no constraints on updates when φ is *consistent* with, but not a member of K. This liberty of ($K \diamond 3$) is the first main difference between revision and update (recall that for such cases ($K * 3$) and ($K * 4$) uniquely determine the result of revision to be $K + \varphi$). The book/magazine example mentioned above falls into this category.

Postulate ($K \diamond 4$) is the update analog of ($K * 5$) highlighting the importance of reaching consistency after update. Once again however ($K \diamond 4$) is more liberal than ($K * 5$) since it does not apply when the initial belief set K is inconsistent; ($K \diamond 4$) only *preserves* consistency, it doesn't *generate* it. Postulate ($K \diamond 7$) essentially says that there is only one way to *minimally* change the world to bring about φ. To see this, consider two sentences φ and ψ for which the precondition of ($K \diamond 7$) holds. Because $\psi \in K \diamond \varphi$, updating by ψ does not produce more change than updating by φ; conversely, since $\varphi \in K \diamond \psi$, updating by φ is not more "expensive" (in terms of induced change) than updating by φ. Consequently, says ($K \diamond 7$), since the two sentences induce the same degree of change to K, they actually produce exactly the *same* change.

For the last two postulates, recall that an update at K is triggered by the occurrence of an action in the world. Hence ($K \diamond 8$) relates the agent's belief set $K \diamond (\varphi \vee \psi)$ after the occurrence of a non-deterministic action with possible effects φ or ψ, with the belief sets $K \diamond \varphi$ and $K \diamond \psi$ resulting from deterministic actions with direct effect φ and ψ respectively. ($K \diamond 8$) states that the former can not be larger than the union of the latter two belief sets, with the provision that the original belief set K is complete.

The last postulate ($K \diamond 9$) reduces the update of any belief set K to the update of all K-worlds. To see the motivation behind this postulate, suppose that $r_1, r_2, \ldots, r_n$ are all the consistent complete theories in L that are compatible with the agent's initial belief set K; i.e. $[K] = \{r_1, r_2, \ldots, r_n\}$. Then, as far as the agent knows, any of $r_1, r_2, \ldots, r_n$ could be the initial state of the world. Consequently, after the occurrence of an action with direct effect φ, the world can be at any of the state $r_1 \diamond \varphi, r_2 \diamond \varphi, \ldots, r_n \diamond \varphi$. Thus the agent's new belief set is $K \diamond \varphi = \bigcap_{r \in [K]} r \diamond \varphi$ (see Figure 8.4).

Apart from their postulates, Katsuno and Mendelzon also introduced semantics for update which, like Grove's semantics for revision, are based on preorders on possible worlds. More precisely, consider a theory K of L, and let $\leq$ be a function that assigns to every world r compatible with K (i.e. $r \in [K]$), a preorder on $\mathbb{M}_L$ denoted $\leq_r$. The function $\leq$ is called a *faithful assignment* iff for every $r \in [K]$ it satisfies the following two conditions: (i) r is the minimum element of $\mathbb{M}_L$ with respect to $\leq_r$ (i.e. for all $r' \in \mathbb{M}_L$, if $r \neq r'$ then $r <_r r'$)[42], and, (ii) for any consistent sentence φ, the set $[\varphi]$ has a minimal element with respect to

[42] As usual, $<_r$ denotes the strict part of $\leq_r$.

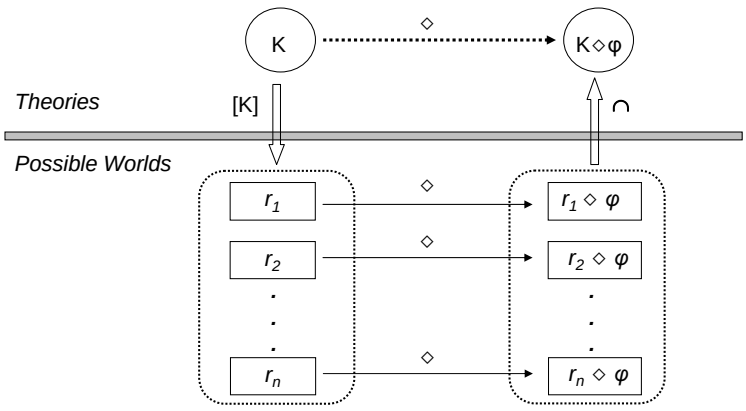

Figure 8.4: Updating a Theory via Possible Worlds

$\leq_r$.[43] Intuitively, $\leq_r$ represents the comparative similarity of possible worlds with respect to r; the further away a world is from r the less similar it is to r.

Based on a faithful assignment $\leq$ for a theory K, Katsuno and Mendelzon define constructively the update of K by a sentence φ as follows:

$$\text{(KM)} \qquad K \diamond \varphi = \begin{cases} \cap(\cup_{r \in [K]} \min([\varphi], \leq_r)) & \text{if } [K] \neq \emptyset \text{ and } [\varphi] \neq \emptyset \\ L & \text{otherwise} \end{cases}$$

In the above definition, $\min([\varphi], \leq_r)$ represents the set of minimal elements in $[\varphi]$ with respect to $\leq_r$; i.e. $\min([\varphi], \leq_r) = \{z \in [\varphi] : \text{there is no } z' \in [\varphi] \text{ such that } z' <_r z\}$.

Let us take a closer look at the above construction of belief updates. Katsuno and Mendelzon tell us that to update K by φ, we first need to consider every possible world r compatible with K *individually* and identify the minimal φ-worlds with respect to $\leq_r$ (i.e. $\min([\varphi], \leq_r)$. The principle of minimal change tells us that these minimal φ-worlds are the states that can result from the occurrence of an action bringing about φ at r. But the agent is not certain that r is the initial state of the world; as far as the agent knows, initially the world can be at any state in $[K]$. Consequently, after the occurrence of φ the world can be at any of the minimal φ-worlds w.r.t to some $r \in [K]$ (i.e. at any state in $\cup_{r \in [K]} \min([\varphi], \leq_r)$).

Katsuno and Mendelzon proved the following representation result showing that their semantics is sound and complete with respect to their postulates for update:

Theorem 18 (Katsuno and Mendelzon [48]). *Let K be a theory of L. If $\leq$ is a faithful assignment for K then the function $\diamond$ induced from $\leq$ by means of (KM) satisfies the KM postulates ($K \diamond 1$) - ($K \diamond 9$). Conversely, for any function $\diamond : \mathbb{K}_L \times L \mapsto \mathbb{K}_L$ that satisfies the KM postulates ($K \diamond 1$) - ($K \diamond 9$) there exists a faithful assignment $\leq$ for K such that (KM) is satisfied.*

[43] Once again, the definition of a faithful assignment presented herein is slightly different, in its phrasing but not in essence, from the original one in [48].

As mentioned already, the Katsuno and Mendelzon semantics for update is quite similar to Grove's system-of-sphere semantics for revision. There are however two major differences between the two: firstly, to a fixed theory K, Katsuno and Mendelzon assign a whole *family* of preorders on possible worlds (one for each world compatible with K) as opposed to a *single* preorder – alias system of spheres – assigned by Grove; secondly, Grove's preorders are always *total* whereas the preorders used by Katsuno and Mendelzon are (in general) partial. For more details on the relationship between belief revision and update, see [75, 76, 78].

8.9 Conclusion

Clearly it is not possible to provide a detailed account of all the work in Belief Revision in a few pages; an entire book would be needed for that. Instead our aim in this chapter was to expose the reader to some of the main ideas and results of the field.

Nevertheless, a few of the missing topics need to be mentioned, even if only in passing.

A large amount of work exists on variations of the AGM postulates and appropriate adjustments to the corresponding constructive models [18, 23, 24, 65, 71, 72, 79, 81, 87, 88, 96]. Moreover, specific belief change operators have been proposed in [13, 90, 97], and their computational complexity has been studied in a seminal article by Eiter and Gottlob [22]. Interesting applications of Belief Revision can be found in [111, 103, 54, 55, 95].

Finally, there is an important body of work on the relationship between Belief Revision and other research areas in Knowledge Representation. Numerous results have been established that reveal profound connections between Belief Revision and areas like *Nonmonotonic Reasoning* [10, 33, 62, 64, 107], *Reasoning about Action* [8, 16, 36, 44, 46, 73, 74, 75, 78, 93, 110], *Conditionals* [26, 30, 35, 37, 60, 89, 83] and *Possibility Theory* [20, 21, 4].

8.10 Acknowledgments

I am grateful to Mary-Anne Williams, Jim Delgrande, Michael Thielscher, and Tommie Meyer for their valuable comments on this chapter, and to the editors of this book, Frank van Harmelen, Vladimir Lifschitz, and Bruce Porter, for their excellent work in coordinating our joint efforts.

Bibliography

[1] C. Alchourron and P. Gardenfors and D. Makinson, "On the Logic of Theory Change: Partial Meet Functions for Contraction and Revision", *Journal of Symbolic Logic*, vol 50, pp 510-530, 1985.

[2] C. Alchourron and D. Makinson, "On the Logic of Theory Change: Safe contractions", *Studia Logica* vol. 44, pp 405-422, 1985.

[3] C. Baral, S. Kraus and J. Minker, "Combining Multiple Knowledge Bases", *IEEE Transactions on Knowledge and Data Engineering*, vol 3(2) , pp 208-220, 1991.

[4] S. Benferhat, D. Dubois, H. Prade, and M.-A. Williams, "A Practical Approach to Revising Prioritized Knowledge Bases" *Studia Logica*, vol. 70, pp 105-130, 2002.

[5] S. Benferhat, S. Kaci, D Le Berre, and M.-A. Williams, " Weakening Conflicting Information for Iterated Revision and Knowledge Integration" *Artificial Intelligence*, vol. 153(1-2), pp 339-371, 2004.

[6] R. Booth and T. Meyer, "Admissible and Restrained Revision", *Journal of Artificial Intelligence Research*, vol. 26, pp. 127-151, 2006.

[7] C. Boutilier, "Revision Sequences and Nested Conditionals", *Proceedings of the 13th International Joint Conference in Artificial Intelligence*, pp 519-525, 1993.

[8] C. Boutilier, "Generalised update: Belief change in dynamic settings, *Proceedings of the 14th International Joint Conference in Artificial Intelligence*, pp. 1550-1556, Montreal, 1995.

[9] C. Boutilier, "Iterated Revision and Minimal Change of Conditional Beliefs", *Journal of Philosophical Logic*, vol. 25(3), 1996.

[10] G. Brewka, "Belief Revision in a Framework of Default Reasoning", *The Logic of Theory Change*, LNCS 465, Springer-Verlag, 1991.

[11] C. Chang and H. Keisler. *Model Theory*. Studies in Logic and the Foundation of Mathematics, vol 73, Elsevier Science Publications, 1991.

[12] S. Chopra, A. Ghose, and T. Meyer, Non-Prioritized Ranked Belief Change, *Journal of Philosophical Logic*, vol 32(4), pp 417-443, 2003.

[13] M. Dalal, "Investigations into a Theory of Knowledge Base Revision: Preliminary Report", *Proceedings of the 7th National (USA) Conference on Artificial Intelligence (AAAI'88)*, pp 475-479, 1988.

[14] A. Darwiche and J. Pearl, "On the Logic of Iterated Belief Revision", *Artificial Intelligence*, vol. 89, pp 1-29, 1997.

[15] A. del Val, "On the Relation between the Coherence and Foundations Theories of Belief Revision", in *Proceedings of the 12th National (USA) Conference on Artificial Intelligence (AAAI'94)*, pp. 909-914, MIT Press, July 1994.

[16] A. del Val and Y. Shoham, "Deriving Properties of Belief Update from Theories of Action", in *Journal Logic, Language and Information*, pp 81-119, vol. 3(2), 1994.

[17] J. Delgrande and T. Schaub, "A Consistency-Based Approach for Belief Change", *Artificial Intelligence*, vol. 151, pp 1-41, 2003.

[18] J. Delgrande, A. Nayak and M. Pagnucco, "Gricean Belief Change", *Studia Logica*, 79(1), 2005.

[19] J. Delgrande, D. Dubois, J. Lang, "Iterated Revision as Prioritized Merging", in *Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning*, pp. 210-220, Morgan Kauffman, 2006.

[20] D. Dubois and H. Prade, "Epistemic Entrenchment and Possibilistic Logic" *Artificial Intelligence*, vol 50, p. 223-239, 1991.

[21] D. Dubois and H. Prade, "Belief Change and Possibility Theory", in P. Gardenfors (ed), *Belief Revision*, p. 142-182, Cambridge University Press, 1992.

[22] T. Eiter and G. Gottlob, "On the Complexity of Propositional Knowledge Base Revision Updates and Counterfactuals", *Artificial Intelligence*, vol 57, pp 227-270, 1992.

[23] E. Ferme, H. Rott, "Revision by comparison", *Artificial Intelligence*, vol 157, 2004.

[24] G. Flouris, D. Plexousakis, G. Antoniou, "On Generalizing the AGM Postulates ", *Proceedings of the 3rd European Starting AI Researcher Symposium*, IOS Press, Riva del Garda, 2006

[25] M. Freund, D. Lehmann, "Belief Revision and Rational Inference", *Technical Report 94-16*, Leibniz Center for Research in Computer Science, Institute of Computer

Science, Hebrew University, 1994.
[26] N. Friedman and J. Halpern, "Conditional Logics of Belief Change", *Proceedings of the 12th National (USA) Conference on Artificial Intelligence (AAAI'94)*, Seattle, 1994.
[27] A. Fuhrmann, "Theory contraction through base contraction", *Journal of Philosophical Logic*, vol. 20, pp. 175-203, 1991.
[28] A. Fuhrmann and S. Hansson, "A Survey on Multiple Contraction", *Journal of Logic, Language, and Information*, vol. 3, pp. 39-76, Kluwer Academic Publishers, 1994.
[29] P. Gardenfors, "Epistemic Importance and Minimal Changes of Belief", *Australasian Journal of Philosophy*, vol 62, pp 136-157, 1984.
[30] P. Gardenfors, "Belief Revision and the Ramsey Test for Conditionals", *The Philosophical Review*, vol 95, pp 81-93, 1986.
[31] P. Gardenfors, *Knowledge in Flux*. MIT Press, 1988.
[32] P. Gardenfors and D. Makinson, "Revisions of knowledge systems using epistemic entrenchment", *Proceedings of Theoretical Aspects of Reasoning about Knowledge*, 83–95. Morgan-Kaufmann, 1988
[33] P. Gardenfors and D. Makinson, "Nonmonotonic Inference based on Expectations", *Artificial Intelligence*, vol. 65, pp 197-245, 1994.
[34] A. Ghose and G. Goebel, "Belief States as Default Theories: Studies in Non-Prioritized Belief Change", *Proceedings of the 13th European Conference on Artificial Intelligence (ECAI'98)*, pp 8-12, 1998.
[35] L. Giordano, V. Gliozzi, and N. Olivetti "Weak AGM Postulates and Strong Ramsey Test: A Logical Formalization", *Artificial Intelligence*, vol 168, pp 1-37, 2005.
[36] M. Goldszmidt and J. Pearl, "Rank-Based Systems: A Simple Approach to Belief Revision, Belief Update and Reasoning about Evidence and Actions", *Proceedings 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR'92)*, pp. 661-672, Cambridge, MA, 1992.
[37] G. Grahne, "Updates and counterfactuals", *Proceedings of the 2nd International Conference on the Principles of Knowledge Representation and Reasoning (KR'91)*, pp. 269-276, Cambridge, MA, 1991.
[38] A. Grove, "Two modellings for theory change", *Journal of Philosophical Logic* vol. 17, pp. 157-170, 1988.
[39] S.-O. Hansson, "Kernel contraction", *Journal of Symbolic Logic*, vol. 59, 1994.
[40] S.-O. Hansson, "Semi-Revision", *JANCL*, vol. 7, pp. 151-175, 1997.
[41] S.-O. Hansson, *A Textbook of Belief Dynamics. Theory Change and Database Updating*. Kluwer Academic Publishers, 1999.
[42] S.-O. Hansson, "A Survey of Non-Prioritized Belief Revision", *Erkenntnis*, vol 50, pp 413-427, Kluwer Academic Publishers, 1998.
[43] S.-O. Hansson, E. L. Ferme, J. Cantwell, M. A. Falappa, "Credibility Limited Revision", *Journal of Symbolic Logic*, vol. 66(4), pp. 1581-1596, 2001.
[44] A. Herzig, "The PMA Revisited" *Proceedings of the International Conference on the Principles of Knowledge Representation and Reasoning*, Cambridge, MA, 1996.
[45] A. Hunter and J. Delgrande, "Iterated Belief Change: A Transition System Approach", *Proceedings of the 19th International Joint Conference in Artificial Intelligence*, pp 460-465, Edinburgh, 2005.
[46] A. Hunter and J. Delgrande, "An Action Description Language for Iterated Belief Change", *Proceedings of the 20th International Joint Conference in Artificial Intelligence*, pp. 2498-2503, Hyderabad, 2007.

[47] Y. Jin and M. Thielscher, "Iterated Revision, Revised", *Proceedings of the 19th International Joint Conference in Artificial Intelligence*, pp 478-483, Edinburgh, 2005.

[48] H. Katsuno and A. Mendelzon, "On the difference between updating a knowledge base and revising it", in *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kauffman, 1991.

[49] H. Katsuno and A. Mendelzon, "Propositional Knowledge Base Revision and Minimal Change", in *Artificial Intelligence*, vol. 52, pp 263-294, 1991.

[52] Konieczny and R. P. Perez, "A framework for iterated revision", in *Journal of Applied Non-Classical Logics*, vol. 10(3-4), 2000.

[51] S. Konieczny and R. Perez, "Merging with Integrity Constraints", *Proceedings of the 5th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU99)*, Lecture Notes in Artificial Intelligence, vol. 1638, pp. 233-244, Springer, Berlin, 1999.

[52] S. Konieczny, "On the Difference Between Merging Knowledge Bases and Combining Them", *Proceedings of the 7th International Conference on Principles of Knowledge Representation and Reasoning (KR00)*, pp. 135-144, Breckenridge, 2000.

[53] S. Konieczny, J. Lang and P. Marquis, "DA^2 Merging Operators", *Artificial Intelligence*, vol 157, pp 49-79, 2004.

[54] R.Y.K. Lau, "Belief Revision for Adaptive Recommender Agents in E-commerce", *Proceedings of the 4th International Conference on Intelligent Data Engineering and Automated Learning (IDEAL'2003)*, Hong Kong, 2003.

[55] R.Y.K. Lau, "Context-Sensitive Text Mining and Belief Revision for Intelligent Information Retrieval on the Web", *Journal of Web Intelligence and Agent Systems*, vol. 1(3-4), pp 151-172, 2003.

[56] D. Lehmann, "Belief Revision, Revised", *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pp. 1534-1540, Montreal, 1995.

[57] D. Lewis, *Counterfactuals*, Harvard University Press, Cambridge, Mass., 1973.

[58] P. Liberatore, M.Schaerf, "Arbitration (or How to Merge Knowledge Bases)", *IEEE Transactions on Knowledge and Data Engineering*, vol 10(1), pp 76-90, 1998.

[59] S. Lindstrom. "A semantic approach to nonmonotonic reasoning: Inference operations and choice", Technical report, Dept of Philosophy, Uppsala University, 1991.

[60] S. Lindstrom, and W. Rabinowicz, "The Ramsey Test revisited", in G. Crocco, L. Farias del Cerro, A. Herzig (eds.), *Conditionals from Philosophy to Computer Science*, Oxford University Press, Oxford, 1995.

[61] D. Makinson, "On the Status of the Postulate of Recovery in the Logic of Theory Change", *Journal of Philosophical Logic*, vol 16, pp 383-394, 1987.

[62] D. Makinson, "Five Faces of Minimality", *Studia Logica*, vol 52, pp 339-379, 1993.

[63] D. Makinson, "Screened Revision", *Theoria*, vol. 63, pp 14-23, 1997.

[64] D. Makinson, *Bridges from Classical to Nonmonotonic Logic*, Texts in Computing, King's College Publications, 2005.

[65] T. Meyer, J. Heidema, W. Labuschagne and L. Leenen, "Systematic Withdrawal", *Journal of Philosophical Logic*, vol 31(5), pp 415-443, 2002.

[66] A. Nayak, "Foundational belief change", *Journal of Philosophical Logic*, vol. 23, pp 495-533, 1994.

[67] A. Nayak, "Iterated Belief Change Based on Epistemic Entrenchment", *Erkenntnis*, vol. 41, pp 353-390, 1994.

[68] A. Nayak, M. Pagnucco, and P. Peppas, "Dynamic Belief Revision Operators", *Arti-*

ficial Intelligence, vol. 146, pp 193-228, 2003.

[69] B. Nebel, "Base Revision Operations and Schemes: Semantics, Representation, and Complexity", in *Proceedings of the 11th European Conference on Artificial Intelligence*, pp. 341-345, Wiley, August 1994.

[70] B. Nebel, "How Hard is it to Revise a Belief Base?", in D. Dubois and H. Prade (eds), *Handbook of Defeasible Reasoning and Uncertainty Management Systems, Vol. 3: Belief Change*, pp. 77-145, Kluwer Academic, 1998.

[71] M. Pagnucco, *The Role of Abductive Reasoning within the Process of Belief Revision*, PhD Thesis, Department of Computer Science, University of Sydney February 1996.

[72] R. Parikh, "Beliefs, Belief Revision, and Splitting Languages", in J. Lawrence Moss, M. de Rijke, (eds)., *Logic, Language, and Computation*, vol 2, pp 266–268, CSLI Lecture Notes No. 96,. CSLI Publications, 1999.

[73] P. Peppas, N. Foo and W. Wobcke, "Events as Theory Operators"', in *Proceedings of the 1st World Conference on the Fundamentals of Artificial Intelligence*, pp. 413-426, Angkor, Paris, 1991.

[74] P. Peppas, and W. Wobcke, "On the Use of Epistemic Entrenchment in Reasoning about Action", in *Proceedings of the 10th European Conference on Artificial Intelligence (ECAI'92)*, pp. 403-407, Wiley & Sons, Vienna, 1992.

[75] P. Peppas, *Belief Change and Reasoning about Action*, PhD Thesis, Basser Dept of Computer Science, University of Sydney, 1994.

[76] P. Peppas and M. Williams, "Constructive Modelings for Theory Change", Notre Dame Journal of Formal Logic, vol. 36, Num. 1, pp. 120-133, 1995.

[77] P. Peppas, "Well Behaved and Multiple Belief Revision", in *Proceedings of the 12th European Conference on Artificial Intelligence (ECAI'96)*, pp. 90-94, Budapest, 1996.

[78] P. Peppas, A. Nayak, M. Pagnucco, N. Foo, R. Kwok, and M. Prokopenko, Revision vs Update: Taking a Closer Look", in *Proceedings of the 12th European Conference on Artificial Intelligence (ECAI'96)*, pp. 95-99, Wiley & Sons, Budapest, 1996.

[79] P. Peppas, N. Foo, and A. Nayak, "Measuring Similarity in Belief Revision", *Journal of Logic and Computation*, vol. 10(4), 2000.

[80] P. Peppas, "The Limit Assumption and Multiple Revision", *Journal of Logic and Computation*, pp. 355-371, vol. 14(3), Oxford University Press, 2004.

[81] P. Peppas, S. Chopra, and N. Foo, "Distance Semantics for Relevance-Sensitive Belief Revision", *Proceedings of the 9th International Conference on the Principles of Knowledge Representation and Reasoning (KR2004)*, Whistler, Canada, June 2004.

[82] P.Z. Revesz, "On the Semantics of Theory Change: Arbitration between Old and New Information" *Proceedings of the 12th ACM SIGACT SIGMOD SIGART Symposium on Principles of Database Systems (PODS-93)*, pp. 71-82, 1993.

[83] H. Rott, "Conditionals and Theory Change: Revisions, Expansions, and Additions", *Synthese*, vol 81, pp 91-113, 1989.

[84] H. Rott, "Two methods of constructing contractions and revisions of knowledge systems", *Journal of Philosophical Logic*, 20:149-173, 1991.

[85] H. Rott, "A nonmonotonic conditional logic for belief revision", *The Logic of Theory Change*, Lecture Notes in Artificial Intelligence, vol 465, Springer-Verlag, 1991.

[86] H. Rott, "Belief Contraction in the Context of the General Theory of Rational Choice", *Journal of Symbolic Logic*, vol 58, 1993.

[87] H. Rott and M. Pagnucco "Severe Withdrawal (and Recovery)", *Journal of Philosophical Logic*, vol 28(5), 1999.

[88] H. Rott, *Change, Choice and Inference*, Oxford University Press, 2001.
[89] M. Ryan, P. Schobbens, "Intertranslating Counterfactuals and Updates", *Journal of Logic, Language and Information*, pp. 123-146, vol. 6(2), 1997.
[90] K. Satoh, "Nonmonotonic reasoning by Minimal Belief Revision, *Proceedings of the International Conference on 5th Generation Computer Systems*, Springer, 1988.
[91] K. Schlechta, D. Lehmann, M. Magidor, "Distance Semantics for Belief Revision", *Proceedings of: Theoretical Aspects of Rationality and Knowledge (TARK VI)*, pp 137-145, Morgan Kaufmann, San Francisco, 1996.
[92] K. Schlechta, "Non-Prioritized Belief Revision Based on Distances Between Models", *Theoria*, vol 63, pp 34-53, 1997.
[93] S. Shapiro, M. Pagnucco, Y. Lesperance, and H. Levesque, "Iterated Belief Change in the Situation Calculus", *Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning (KR'00)*, Morgan Kaufmann, 2000.
[94] W. Spohn, "Ordinal conditional functions: A dynamic theory of epistemic states", in *Causation in Decision, Belief Change and Statistics*, vol. 2, pp. 105-134, Kluwer Academic Publishers, 1988.
[95] D. Stavrinoudis, M. Xenos, P. Peppas, and D. Christodoulakis, "Early Estimation of Users Perception of Software Quality", *Software Quality Journal*, vol. 13(2), pp. 155-175, 2005.
[96] R. Wassermann, "Resource Bounded Belief Revision", *Erkenntnis*, vol 50, 1999.
[97] A. Weber, "Updating Propositional Formulas", in *Expert Database Systems - Proceedings From the 1st International Conference*, Benjamin/Cummings, 1986.
[98] M.-A. Williams, "Two Operators for Theory Bases", *Proceedings of the Australian Joint Conference on Artificial Intelligence*, pp. 525 - 332, World Scientific, 1992.
[99] M.-A. Williams, *Transmutations of Knowledge Systems*, PhD Thesis, University of Sydney, Australia, 1993.
[100] M.-A. Williams, "On the Logic of Theory Base Change", In C. MacNish, D. Pearce, and L. Pereira, editors, *Logics in Artificial Intelligence - European Workshop JELIA'94*, pp. 86-105, Springer-Verlag, 1994.
[101] M.-A. Williams, "Transmutations of Knowledge Systems", *Proceedings of the 4th International Conference on Principles of Knowledge Representation and Reasoning*, pp. 619-629, Morgan Kaufmann, 1994.
[102] M.-A. Williams, "Iterated Theory Base Change: A Computational Model", *Proceedings of the 14th International Joint Conference in Artificial Intelligence*, pp. 1541-1547, 1995.
[103] M.-A. Williams, "Applications of Belief Revision", *Lecture Notes in Computer Science*, vol. 1472, 1998.
[104] M.-A. Williams and A. Sims, "SATEN: An Object-Oriented Web-Based Revision and Extraction Engine", *CoRR cs.AI/0003059*, 2000.
[105] M. Winslett, "Reasoning about action using a possible models approach", in *Proceedings of the 7th National (USA) Conference on Artificial Intelligence (AAAI'88)*, pp. 89-93, 1988.
[106] D. Zhang, "Belief revision by sets of sentences", *Journal of Computer Science and Technology*, vol. 11(2), pp. 108-125, 1996.
[107] D. Zhang, S. Chen, W. Zhu, and H. Li, "Nonmonotonic Reasoning and Multiple Belief Revision", *Proceedings of the 15th International Joint Conference in Artificial Intelligence*, 1997.

[108] D. Zhang and N. Foo. Infinitary Belief Revision. *Journal of Philosophical Logic*, vol. 30(6), pp. 525-570, 2001.

[109] D. Zhang, "Properties of Iterated Multiple Belief Revision", *Proceedings of the 7th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR)*, pp. 314-325. Springer, 2004.

[110] Y. Zhang, N.Y. Foo, "Update with Disjunctive Information: From Syntactical and Semantical Perspective", *Journal of Computational Intelligence*, vol 16 (1), 2000.

[111] D. Zowghi, A. Ghose and P. Peppas, "A Framework for Reasoning about Requirements Evolution", *Proceedings of the 4th Pacific Rim International Conference on Artificial Intelligence*, Springer-Verlag, Cairns, 1996.