Handbook of Philosophical Logic 16

Dov M. Gabbay
Franz Guenthner *Editors*

Handbook of Philosophical Logic

Second Edition

Springer

HANDBOOK OF PHILOSOPHICAL LOGIC
2ND EDITION

VOLUME 16

HANDBOOK
OF PHILOSOPHICAL LOGIC

2nd Edition

Volume 16

edited by D.M. Gabbay and F. Guenthner

Volume 1 – ISBN 0-7923-7018-X
Volume 2 – ISBN 0-7923-7126-7
Volume 3 – ISBN 0-7923-7160-7
Volume 4 – ISBN 1-4020-0139-8
Volume 5 – ISBN 1-4020-0235-1
Volume 6 – ISBN 1-4020-0583-0
Volume 7 – ISBN 1-4020-0599-7
Volume 8 – ISBN 1-4020-0665-9
Volume 9 – ISBN 1-4020-0699-3
Volume 10 – ISBN 1-4020-1644-1
Volume 11 – ISBN 1-4020-1966-1
Volume 12 – ISBN 1-4020-3091-6
Volume 13 – ISBN 978-1-4020-3520-3
Volume 14 – ISBN 978-1-4020-6323-7
Volume 15 – ISBN 978-94-007-0484-8

HANDBOOK OF PHILOSOPHICAL LOGIC

2nd EDITION

VOLUME 16

Edited by

D.M. GABBAY
King's College, London, U.K.

and

F. GUENTHNER
Centrum für Informations- und Sprachverarbeitung,
Ludwig-Maximilians-Universität München, Germany

Springer

Editors
Prof. Dr. Dov M. Gabbay
King's College London
Dept. Computer Science
London
United Kingdom
dov.gabbay@kcl.ac.uk

Franz Guenthner
University of Munich
Centrüm für Informations- und Sprachverarbeitung
Oettingenstr. 67
80538 Munich
Germany
gue@cis.uni-muenchen.de

ISBN 978-94-007-0478-7 e-ISBN 978-94-007-0479-4
DOI 10.1007/978-94-007-0479-4
Springer Dordrecht Heidelberg London New York

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

CONTENTS

PREFACE TO THE SECOND EDITION

It is with great pleasure that we are presenting to the community the second edition of this extraordinary handbook. It has been over 15 years since the publication of the first edition and there have been great changes in the landscape of philosophical logic since then.

The first edition has proved invaluable to generations of students and researchers in formal philosophy and language, as well as to consumers of logic in many applied areas. The main logic article in the Encyclopaedia Britannica 1999 has described the first edition as 'the best starting point for exploring any of the topics in logic'. We are confident that the second edition will prove to be just as good!

The first edition was the second handbook published for the logic community. It followed the North Holland one volume *Handbook of Mathematical Logic*, published in 1977, edited by the late Jon Barwise. The four volume *Handbook of Philosophical Logic*, published 1983–1989 came at a fortunate temporal junction at the evolution of logic. This was the time when logic was gaining ground in computer science and artificial intelligence circles.

These areas were under increasing commercial pressure to provide devices which help and/or replace the human in his daily activity. This pressure required the use of logic in the modelling of human activity and organisation on the one hand and to provide the theoretical basis for the computer program constructs on the other. The result was that the *Handbook of Philosophical Logic*, which covered most of the areas needed from logic for these active communities, became their bible.

The increased demand for philosophical logic from computer science and artificial intelligence and computational linguistics accelerated the development of the subject directly and indirectly. It directly pushed research forward, stimulated by the needs of applications. New logic areas became established and old areas were enriched and expanded. At the same time, it socially provided employment for generations of logicians residing in computer science, linguistics and electrical engineering departments which of course helped keep the logic community thriving. In addition to that, it so happens (perhaps not by accident) that many of the Handbook contributors became active in these application areas and took their place as time passed on, among the most famous leading figures of applied philosophical logic of our times. Today we have a handbook with a most extraordinary collection of famous people as authors!

The table below will give our readers an idea of the landscape of logic and its relation to computer science and formal language and artificial intelligence. It shows that the first edition is very close to the mark of what was needed. Two topics were not included in the first edition, even though

D. Gabbay and F. Guenthner (eds.),
Handbook of Philosophical Logic, Volume 16, vii–ix.
DOI 10.1007/978-94-007-0479-4, © Springer Science+Business Media B.V. 2011

they were extensively discussed by all authors in a 3-day Handbook meeting. These are:

- a chapter on non-monotonic logic
- a chapter on combinatory logic and λ-calculus

We felt at the time (1979) that non-monotonic logic was not ready for a chapter yet and that combinatory logic and λ-calculus was too far removed.[1] Non-monotonic logic is now a very major area of philosophical logic, alongside default logics, labelled deductive systems, fibring logics, multi-dimensional, multimodal and substructural logics. Intensive re-examinations of fragments of classical logic have produced fresh insights, including at time decision procedures and equivalence with non-classical systems.

Perhaps the most impressive achievement of philosophical logic as arising in the past decade has been the effective negotiation of research partnerships with fallacy theory, informal logic and argumentation theory, attested to by the Amsterdam Conference in Logic and Argumentation in 1995, and the two Bonn Conferences in Practical Reasoning in 1996 and 1997.

These subjects are becoming more and more useful in agent theory and intelligent and reactive databases.

Finally, fifteen years after the start of the Handbook project, I would like to take this opportunity to put forward my current views about logic in computer science, computational linguistics and artificial intelligence. In the early 1980s the perception of the role of logic in computer science was that of a specification and reasoning tool and that of a basis for possibly neat computer languages. The computer scientist was manipulating data structures and the use of logic was one of his options.

My own view at the time was that there was an opportunity for logic to play a key role in computer science and to exchange benefits with this rich and important application area and thus enhance its own evolution. The relationship between logic and computer science was perceived as very much like the relationship of applied mathematics to physics and engineering. Applied mathematics evolves through its use as an essential tool, and so we hoped for logic. Today my view has changed. As computer science and artificial intelligence deal more and more with distributed and interactive systems, processes, concurrency, agents, causes, transitions, communication and control (to name a few), the researcher in this area is having more and more in common with the traditional philosopher who has been analysing

[1] I am really sorry, in hindsight, about the omission of the non-monotonic logic chapter. I wonder how the subject would have developed, if the AI research community had had a theoretical model, in the form of a chapter, to look at. Perhaps the area would have developed in a more streamlined way!

such questions for centuries (unrestricted by the capabilities of any hardware).

The principles governing the interaction of several processes, for example, are abstract an similar to principles governing the cooperation of two large organisation. A detailed rule based effective but rigid bureaucracy is very much similar to a complex computer program handling and manipulating data. My guess is that the principles underlying one are very much the same as those underlying the other.

I believe the day is not far away in the future when the computer scientist will wake up one morning with the realisation that he is actually a kind of formal philosopher!

The projected number of volumes for this Handbook is about 18. The subject has evolved and its areas have become interrelated to such an extent that it no longer makes sense to dedicate volumes to topics. However, the volumes do follow some natural groupings of chapters.

I would like to thank our authors are readers for their contributions and their commitment in making this Handbook a success. Thanks also to our publication administrator Mrs J. Spurr for her usual dedication and excellence and to Kluwer Academic Publishers for their continuing support for the Handbook.

Dov M. Gabbay
King's College London

Logic	**IT**			
	Natural language processing	**Program control specification, verification, concurrency**	**Artificial intelligence**	**Logic programming**
Temporal logic	Expressive power of tense operators. Temporal indices. Separation of past from future	Expressive power for recurrent events. Specification of temporal control. Decision problems. Model checking.	Planning. Time dependent data. Event calculus. Persistence through time—the Frame Problem. Temporal query language. temporal transactions.	Extension of Horn clause with time capability. Event calculus. Temporal logic programming.
Modal logic. Multi-modal logics	generalised quantifiers	Action logic	Belief revision. Inferential databases	Negation by failure and modality
Algorithmic proof	Discourse representation. Direct computation on linguistic input	New logics. Generic theorem provers	General theory of reasoning. Non-monotonic systems	Procedural approach to logic
Non-monotonic reasoning	Resolving ambiguities. Machine translation. Document classification. Relevance theory	Loop checking. Non-monotonic decisions about loops. Faults in systems.	Intrinsic logical discipline for AI. Evolving and communicating databases	Negation by failure. Deductive databases
Probabilistic and fuzzy logic	logical analysis of language	Real time systems	Expert systems. Machine learning	Semantics for logic programs
Intuitionistic logic	Quantifiers in logic	Constructive reasoning and proof theory about specification design	Intuitionistic logic is a better logical basis than classical logic	Horn clause logic is really intuitionistic. Extension of logic programming languages
Set theory, higher-order logic, λ-calculus, types	Montague semantics. Situation semantics	Non-well-founded sets	Hereditary finite predicates	λ-calculus extension to logic programs

Imperative vs. declarative languages	Database theory	Complexity theory	Agent theory	Special comments: A look to the future
Temporal logic as a declarative programming language. The changing past in databases. The imperative future	Temporal databases and temporal transactions	Complexity questions of decision procedures of the logics involved	An essential component	Temporal systems are becoming more and more sophisticated and extensively applied
Dynamic logic	Database updates and action logic	Ditto	Possible actions	Multimodal logics are on the rise. Quantification and context becoming very active
Types. Term rewrite systems. Abstract interpretation	Abduction, relevance	Ditto	Agent's implementation rely on proof theory.	
	Inferential databases. Non-monotonic coding of databases	Ditto	Agent's reasoning is non-monotonic	A major area now. Important for formalising practical reasoning
	Fuzzy and probabilistic data	Ditto	Connection with decision theory	Major area now
Semantics for programming languages. Martin-Löf theories	Database transactions. Inductive learning	Ditto	Agents constructive reasoning	Still a major central alternative to classical logic
Semantics for programming languages. Abstract interpretation. Domain recursion theory.		Ditto		More central than ever!

Classical logic. Classical fragments	Basic background language	Program synthesis	A basic tool	
Labelled deductive systems	Extremely useful in modelling		A unifying framework. Context theory.	Annotated logic programs
Resource and substructural logics	Lambek calculus		Truth maintenance systems	
Fibring and combining logics	Dynamic syntax	Modules. Combining languages	Logics of space and time	Combining features
Fallacy theory				
Logical Dynamics	Widely applied here			
Argumentation theory games		Game semantics gaining ground		
Object level/ metalevel			Extensively used in AI	
Mechanisms: Abduction, default relevance			ditto	
Connection with neural nets				
Time-action-revision models			ditto	

	Relational databases	Logical complexity classes	The workhorse of logic	The study of fragments is very active and promising.
	Labelling allows for context and control.		Essential tool.	The new unifying framework for logics
Linear logic			Agents have limited resources	
	Linked databases. Reactive databases		Agents are built up of various fibred mechanisms	The notion of self-fibring allows for self-reference
				Fallacies are really valid modes of reasoning in the right context.
			Potentially applicable	A dynamic view of logic
				On the rise in all areas of applied logic. Promises a great future
			Important feature of agents	Always central in all areas
			Very important for agents	Becoming part of the notion of a logic
				Of great importance to the future. Just starting
			A new theory of logical agent	A new kind of model

ODINALDO RODRIGUES, DOV GABBAY AND
ALESSANDRA RUSSO

BELIEF REVISION

1 INTRODUCTION AND HISTORICAL PERSPECTIVE

The investigation of how humans reach conclusions from given premises has long been the subject of intense research in the literature. It was the basis of the development of classical logic, for instance. The investigation of how humans change their minds in the face of new contradictory information is however somewhat more recent. Early accounts include the work of Ramsey [Ramsey, 1931; Ramsey, 1990] in his insights into conditional statements, for instance, and subsequently the work on conditionals by Stalnaker [Stalnaker, 1968b] and by Lewis [Lewis, 1973], among others. More recent work on the formalisation of *common-sense reasoning*, sometimes also called *non-monotonic reasoning*, include [McCarthy, 1958; Brewka, 1990; Łukaszewicz, 1990; Reiter, 1980].

The now trademark term AGM is an acronym formed with the initials of the main proposers of this theory of belief change, namely, Carlos Alchourrón, Peter Gärdenfors and David Makinson. Alchourrón and Makinson had worked jointly on *theory change* in the early 80s [Alchourrón and Makinson, 1982; Alchourrón and Makinson, 1985], and, independently, Gärdenfors had been working on *belief change* in the late 70's and early 80s [Gärdenfors, 1978; Gärdenfors, 1982]. After collaborations in various combinations of the three authors, they published together the paper "On the Logic of Theory Change" [Alchourrón *et al.*, 1985], which provided the basis for what is now known as the AGM theory of belief revision.[1]

The main object of study of the theory of belief revision is the dynamics of the process of belief change: when an agent is faced with new information which contradicts her current beliefs, she will have to retract some of the old beliefs in order to accommodate the new belief consistently. In general, this can be done in a number of ways. Belief revision is concerned with how to make the process *rational*. The AGM theory stipulates some rationality principles to be observed — the so-called *AGM postulates for belief change*. These are discussed in more detail in Section 2.

[1] A more comprehensive recapitulation of the early history of AGM done by Makinson himself can be found in [Makinson, 2003]

D. Gabbay and F. Guenthner (eds.),
Handbook of Philosophical Logic, Volume 16, 1–114.
DOI 10.1007/978-94-007-0479-4_1,

Many other articles followed the initial proposal of the AGM theory analysing properties of belief change operations and how they relate to each other. Subsequently, Gärdenfors published a book entitled "Knowledge in Flux" [Gärdenfors, 1988], which is an excellent reference to the early work in the area. Since then, the work on belief revision has flourished and diversified into several different subareas.

One of the first specialisations was related to the status given to certain beliefs held by an agent. Following the usual terminology, we will call the collection of beliefs of an agent at a given moment in time her *belief set*. According to the *coherentist view*, an agent has no mechanism other than logical coherence for keeping track of the 'reasons' why a given belief is supported. Consequently, belief change operations need only to describe the relationship between belief sets at two adjacent moments in time. In the *foundationalist view*, however, beliefs can be held by an agent only if they have a proper *justification* — if a justification becomes untenable and is retracted, then all beliefs that rely on it must also be retracted. Therefore, belief change operations need to specify a mechanism for maintaining the justifications for the beliefs. In the simplest form of foundationalism, some beliefs are regarded as requiring no justification and called *basic* (sometimes also called *foundational* [Harman, 1986, page 31]). A variation of this approach with special interest to computer science makes a simple distinction between the set of beliefs supported by an agent (her *belief set*) and the set of beliefs from which these are derived (her *belief base*). Changes are made to the belief base, which, in general, is a finite set. The reader is referred to [Gärdenfors, 1990; Doyle, 1992] for a more comprehensive discussion of the differences between the two paradigms. Obviously, the problem of belief revision can be approached from different perspectives as well. We will present some of these throughout this chapter.

A related area of investigation on theory change is concerned with the formalisation of the effects of the execution of actions in the real world. Operations of this types are usually called *updates* and this problem is related with other areas of artificial intelligence, including planning, logical databases and robotics. In a very influential article [Fagin *et al.*, 1983], Fagin et al. investigated how to update logical databases presenting many of the principles of theory revision and update now widely accepted. One of these is the idea of preservation of old information through the definition of some minimality criteria. They also realised early on the importance of considering the *logical consequences* associated with a database.

Some other early work in that area is also worth mentioning. For instance, Ginsberg and Smith's articles on the formalisation of the reasoning about actions [Ginsberg and Smith, 1988a; Ginsberg and Smith, 1988b] and the well known article written by Winslett "Reasoning about action using a possible models approach" [Winslett, 1988a]. The latter highlighted the importance of semantical considerations in the achievement of rational changes of information caused by the execution of actions.

Even though belief revision and updates are clearly distinct, they have similarities between them. In particular, the so-called *principle of informational economy*. After all, it does not seem rational for an agent to discard all of the knowledge accumulated about the world in the face of new contradictory information. The similarities between updates and belief revision (as well as those between other forms of non-monotonic reasoning) have been extensively investigated (see, for instance [Makinson, 1993; Katsuno and Satoh, 1991; Makinson and Gärdenfors, 1989]). Analogously, the differences between the two were emphasised by Katsuno and Mendelzon [Katsuno and Mendelzon, 1991a; Katsuno and Mendelzon, 1992]. In an analogy to the AGM trio, they proposed some postulates for update operations. Further investigation on specialised types of update operations appeared in [Brewka and Hertzberg, 1993].

We note that several formalisms deal independently with belief revision [Dalal, 1988a; Gärdenfors, 1988; Alchourrón and Makinson, 1985], updates in databases [Brewka and Hertzberg, 1993; Winslett, 1988a; Ginsberg and Smith, 1988a; Ginsberg and Smith, 1988b], default reasoning [Reiter, 1980; Poole, 1988a; Brewka, 1989b; Brewka, 1991], conditional reasoning [Lewis, 1973; Nute, 1984; Stalnaker, 1968b; Grahne, 1991a; Grahne, 1991b], argumentation [Bench-Capon and Dunne, 2007; Besnard and Hunter, 2008; Modgil, 2009], etc. However, there is very little work on the combination of these. Such integration is arguably essential to the modelling of a rational agent that needs to deal with multiple forms of common-sense reasoning.

One issue that often arises in the problem of belief revision is the choice of what beliefs to give up during a revision operation. One approach relies on the representation of how strongly an agent feels about her beliefs so that when a choice needs to be made, those beliefs on which she less strongly believes will go first. In order to support this approach, some of the work on belief revision includes the investigation of the representation of priorities associated with the beliefs in the belief set (or base). The usual mechanism is a preference relation associated with the beliefs. The reasoning about the preferences can itself be quite complex, because in the general (and interesting) case, an agent has only partial information about these preferences.

Another issue is the investigation of the relation between belief sets obtained after successive belief change operations. The original AGM theory had little to say about the iteration of the process of belief change. The study of the properties of *iterated revision* started in the mid 90's and is, of course, of great importance to both computer science and philosophy in general. More profound considerations started with [Darwiche and Pearl, 1994; Freund and Lehmann, 1994; Lehmann, 1995; Friedman and Halpern, 1996; Darwiche and Pearl, 1997; Rodrigues, 1998; Friedman and Halpern, 1999] and more recent work includes [Konieczny and Pérez, 2000; Herzig *et al.*, 2003; Rodrigues, 2005]. The problems of iteration will be discussed in Section 4.

Finally, some effort has also been directed towards the study of the complexity involved in the implementation of belief revision systems. Some results can be found in [Gärdenfors and Rott, 1995, page 98] and in [Eiter and Gottlob, 1992a; Eiter and Gottlob, 1993; Nebel, 1991a; Nebel, 1992; Nebel, 1998; Gärdenfors and Rott, 1995; Eiter and Gottlob, 1992b; Nebel, 1992; Nebel, 1998]. As we shall see in Section 2, one of the postulates for belief revision operations stipulates that the resulting belief set is consistent provided that the revising information is not itself contradictory. Thus, the problem of belief revision is at least as hard as the problem of deciding the satisfiability of a set of formulae. Reasoning about preferences can also add to the complexity of the problem and, as a consequence, many belief revision formalisms constrain themselves to a fragment of first-order logic or, in most cases, to propositional logic [Dalal, 1988a; Dargam, 1996; Katsuno and Mendelzon, 1991b; Gabbay and Rodrigues, 1997; Rodrigues, 2003]. Some complexity results will be briefly presented in Section 6.

In the sections that follow we provide a review of the AGM framework, discuss some philosophical problems of the process of belief change not addressed by the theory, and present some representative work proposed to tackle the issues.

This chapter is structured as follows: alternative formalisations of the problem of belief revision are presented and discussed in Section 2. Some well known belief revision operators are presented in Section 3. In Section 4, we discuss the problems of iterated revision and alternative ways of dealing with it. Section 5 presents some special types of revision formalisms. This is followed by a survey of complexity issues associated with the revision operation in Section 6; some applications of belief revision in Section 7 and a discussion of challenges and open issues in Section 8.

2 FORMALISATION OF THE PROBLEM OF BELIEF REVISION

In order to discuss belief revision in more detail, it will be useful to introduce some terminology first. Let K be a set of formulae representing the beliefs of an agent in the language of some logic L, with consequence relation Cn.[2] K is called a *belief set* when it is closed under Cn, i.e., $K = \mathrm{Cn}(K)$. Given a belief set K and a belief φ, we say that φ is *accepted* in K when $\varphi \in K$.

As mentioned previously, the framework of belief revision was developed around some desiderata of the operations on belief sets, called the *AGM postulates for belief change*, whose main purpose is to model *rational changes of belief*. The AGM theory defines three main types of belief change:

- *Expansion*: the incorporation of a new belief φ into a belief set K. The new belief set is represented by $K{+}\varphi$ and defined simply as $\mathrm{Cn}(K \cup \{\varphi\})$. Notice that $K{+}\varphi$ will be inconsistent if K is inconsistent; or if φ is contradictory; or if they are both independently satisfiable although $K \cup \{\varphi\}$ is not jointly satisfiable. Since belief sets are closed under the consequence relation, the inconsistent belief set is unique and equivalent to the set of all formulae in the language. The inconsistent belief set will be denoted by $K_\perp$.

- *Contraction*: the retraction of a belief from a belief set. Since belief sets are closed under the consequence relation, in order to retract a belief φ from K, it is also necessary to remove other beliefs in K that imply φ. A contraction of K by φ is represented by $K{-}\varphi$.

- *Revision*: the incorporation of a belief φ into a belief set K in such a way that the resulting belief set is consistent unless φ is itself contradictory. The interesting case is when φ is not contradictory but inconsistent with K. The main issue in this case is to determine which of the beliefs in K to retract in order to consistently accept φ. The revision of a belief set K by a belief φ is represented by $K{\circ}\varphi$.

As can be seen, the interesting belief change operations are contractions and revisions. In fact, it is possible to define one of the operations in terms of the other. The *Levi identity* defines revisions in terms of contractions and the *Harper identity* defines contractions in terms of revisions.[3]

Levi identity: $K{\circ}\varphi{=}(K{-}\neg\varphi){+}\varphi$

[2] We will use the terms "belief" and "formula" interchangeably.

[3] Note that these operations are all defined at the metalevel. For an investigation on how to bring contraction to the object level see [Gabbay *et al.*, 2002; Gabbay *et al.*, 2004].

Harper identity: $K{-}\varphi{=}K \cap (K{\circ}\neg\varphi)$

In this chapter, we will concentrate only on the revision process although we state the following important result [Gärdenfors, 1988] that will prove useful in the next sections:

THEOREM 1. *If a contraction function verifies the AGM postulates for contraction and an expansion function verifies the AGM postulates for expansion and a revision function* $\circ$ *is defined in terms of both according to the Levi identity above, then* $\circ$ *verifies the AGM postulates for revision (presented in the next section).*

The intuition behind revisions defined via the Levi identity is that one should first give up all beliefs that are in conflict with the new information before adding it (if consistency is to be maintained). Naturally, when forced to give up some of the beliefs, one should try and minimise the loss of information involved in the process. This requirement is commonly referred to as the *principle of minimal change* or the *principle of informational economy* [Gärdenfors, 1988, page 49]:

> "... when we change our beliefs, we want to retain as much as possible of our old beliefs — information is not in general gratuitous, and unnecessary losses of information are therefore to be avoided."

Without considering preferences between beliefs, in general there will be several possibilities of contracting a belief set K in such a way that it can consistently accept a new belief φ. In order to comply with the principle of minimal change, we will be interested only in those contractions that minimise the loss of beliefs. This can be formalised in the following way:

DEFINITION 2 (Maximal subsets that fail to imply a sentence). Let K be a belief set and $\neg\varphi$ a belief. A set K' is a maximal subset of K that fails to imply $\neg\varphi$ iff the following conditions are met:

- $K' \subseteq K$
- $\neg\varphi \notin \mathrm{Cn}(K')$
- $\forall K''$, $K' \subset K'' \subseteq K$ implies $\neg\varphi \in \mathrm{Cn}(K'')$

In other words, any subset of K larger than K' (and containing it) would result in the derivation of $\neg\varphi$. It should always be possible to find such subsets unless $\neg\varphi$ is a tautology. Following the usual convention found in the literature, the set of all subsets of K that do not imply $\neg\varphi$ will be denoted $K_{\perp}\neg\varphi$. A *maxichoice contraction* of K by $\neg\varphi$ is an operation that returns one of the elements of $K_{\perp}\neg\varphi$ when there is at least one or K itself if $\neg\varphi$ is a tautology. A *full meet contraction* of K by $\neg\varphi$ is an operation that returns the intersection of all elements of $K_{\perp}\neg\varphi$ or K itself if $K_{\perp}\neg\varphi$ is empty. Finally, a *partial meet contraction* of K by $\neg\varphi$ is an operation that returns the intersection of *some* appropriately selected elements of $K_{\perp}\neg\varphi$ if it is non-empty[4] or as before K itself, otherwise. Based on these contraction functions, a number of revision functions can be defined via the Levi identity. However, Alchourrón and Makinson showed that maxichoice contraction functions produce belief sets that are too large [Alchourrón and Makinson, 1982] and, as a result, revision operations defined in terms of these contractions will be *maximal*:[5]

THEOREM 3. *If a revision function* $\circ$ *is defined from a maxichoice contraction function* $-$ *via the Levi identity, then for any belief* φ *such that* $\neg\varphi \in K$, $K\circ\varphi$ *will be maximal.*

On the other hand, full meet contractions are too restrictive, and analogously Alchourrón and Makinson showed that revisions defined in terms of this kind of contractions will in general produce belief sets that are too "small":

THEOREM 4. *If a revision function* $\circ$ *is defined from a full meet contraction function* $-$ *via the Levi identity, then for any belief* φ *such that* $\neg\varphi \in K$, $K\circ\varphi = \mathrm{Cn}(\{\varphi\})$.

Thus, on the one hand, if maxichoice revisions are used, the arrival of some conflicting information makes an agent omniscient. On the other hand, if full meet revisions are used, the arrival of new conflicting information causes the agent to lose all previous information. So it seems that the only realistic revisions that can be defined in terms of contractions and the Levi identity are the ones that use partial meet contractions. The difficulty with this type of revisions is that they rely on a selection mechanism that is *external* to the agent's own representation of beliefs. Although arguably not as elegant from the philosophical point of view, the need for some extra information supporting the beliefs of an agent will become evident in the sections that follow. In particular, a discussion to motivate the employment of an agent's *epistemic state* (as opposed to her *belief* state) is presented in Section 4.

[4]The selection function may take other criteria such as epistemic entrenchment (see Section 2.5) into consideration during the process.

[5]A belief set K is maximal if for any belief φ, either $\varphi \in K$ or $\neg\varphi \in K$.

The postulates for the revision operation as given in [Gärdenfors, 1988, pp. 54–56] are now presented. In the following presentation φ and ψ denote beliefs and the symbol $K_\perp$ denotes the inconsistent belief set as before. From now on, $\circ$ will be subscripted to denote a specific belief revision operation. In particular, $\circ_a$ will be used to denote a revision operation complying with the AGM postulates.

2.1 AGM postulates for belief revision

(K°1) $K \circ_a \varphi$ is a belief set
(K°2) $\varphi \in K \circ_a \varphi$
(K°3) $K \circ_a \varphi \subseteq \mathrm{Cn}(K \cup \{\varphi\})$
(K°4) If $\neg\varphi \notin K$, then $\mathrm{Cn}(K \cup \{\varphi\}) \subseteq K \circ_a \varphi$
(K°5) $K \circ_a \varphi = K_\perp$ only if φ is contradictory
(K°6) If $\varphi \equiv \psi$, then $K \circ_a \varphi \equiv K \circ_a \psi$
(K°7) $K \circ_a (\varphi \wedge \psi) \subseteq \mathrm{Cn}((K \circ_a \varphi) \cup \{\psi\})$
(K°8) If $\neg\psi \notin K \circ_a \varphi$, then $\mathrm{Cn}(K \circ_a \varphi \cup \{\psi\}) \subseteq K \circ_a (\varphi \wedge \psi)$

Postulate (K°1) requires the result of the revision operation to be a belief set, i.e., that the revised set be closed under the consequence relation Cn. In more general terms, (K°1) requires that the operation preserves the defining properties of the original belief set.

Postulate (K°2) is known as the *success postulate*, but sometimes also referred to as the *principle of the primacy of the update* [Dalal, 1988b]. It basically requires the revision process to be successful in the sense that the new belief is effectively accepted after the revision operation is applied. The controversy is that the new belief may be itself contradictory, in which case (K°2) requires the new belief set to be inconsistent. Since the logic used to model the belief sets is classical and AGM adopts the coherentist view, all beliefs become accepted after such a revision is performed. The reliance of AGM framework on the consistency notion is discussed in more detail in [Gabbay *et al.*, 2000; Rodrigues *et al.*, 2008].

(K°3) says sets an expansion as the upper bound of a revision operation. (K°4) on the other hand, specifies that provided that the new belief is not inconsistent with the current belief set, the revision operation will include all of the consequences of the old belief set together with the new belief. Thus, it sets a lower bound for the operation in the case when the new belief is consistent with the current belief set.

(K°5) is sometimes referred to as the *recovery postulate*. It guarantees that the result of a revision is consistent provided that the revising sentence itself is non-contradictory. To understand what (K∘3)–(K∘5) say, two cases need to be considered:

Case 1: $K \cup \{\varphi\}$ is consistent.

In this case, (K∘3) and (K∘4) require that $K \circ_a \varphi = \mathrm{Cn}(K \cup \{\varphi\})$, since by (K∘3), $K \circ_a \varphi \subseteq \mathrm{Cn}(K \cup \{\varphi\})$ and by (K∘4), $\mathrm{Cn}(K \cup \{\varphi\}) \subseteq K \circ_a \varphi$. (K∘5) is vacuously true.

Case 2: $K \cup \{\varphi\}$ is inconsistent.

In this case, (K∘3) does not say much about $K \circ_a \varphi$. If $K \cup \{\varphi\}$ is classically inconsistent, then any theory whatsoever is included in $\mathrm{Cn}(K \cup \{\varphi\})$, because this theory is simply $K_\perp$. Similarly, (K∘4) says little about $K \circ_a \varphi$. Since $K \cup \{\varphi\}$ is inconsistent, it follows that $\neg\varphi \in K$ (since K is a closed theory), and hence (K∘4) is satisfied vacuously. As for (K∘5), two subcases can be considered:

1. φ is non-contradictory. In this case, (K∘5) requires $K \circ_a \varphi$ to be consistent, but gives us no clue as to what $K \circ_a \varphi$ should look like — minimal requirements are given by (K∘1) and (K∘2).

2. φ is contradictory. In this case, (K∘5) says nothing about $K \circ_a \varphi$. However, (K∘1) and (K∘2) jointly force $K \circ_a \varphi = K_\perp$.

The above case analysis shows that the AGM postulates (K∘3)–(K∘5) have something to say only when $K \cup \{\varphi\}$ is consistent, or when it is inconsistent even though φ is non-contradictory. The particular way of writing the postulates given above makes use of technical properties of classical logic (the way inconsistent theories prove everything). Also notice that classical inconsistency is the (only) trigger of the revision process — if the new belief and the current belief set are jointly consistent, the revision simply amounts to an expansion.

When considering the AGM postulates for logics other than classical logic, the notion of *acceptability* needs to be employed instead of consistency whenever the latter is missing. In that case, one needs to decide when a revision is required according to what is reasonable in the non-classical logic. In classical logic, the postulates do not give any clue beyond (K∘5) as to what to require when $K \cup \{\varphi\}$ is inconsistent. These issues have been analysed in detail in [Gabbay *et al.*, 2000; Rodrigues *et al.*, 2008; Gabbay *et al.*, 2010].

To summarise, postulates (K∘3)–(K∘4) effectively mean the following:

($K^\circ_{3,4}$) If φ is consistent with K, then $K \circ_a \varphi = \mathrm{Cn}(K \cup \{\varphi\})$.

If K is finitely representable, it can be taken as a formula and the postulate above corresponds to postulate (R2) in Katsuno and Mendelzon's rephrasing of the AGM postulates for belief sets represented by finite bases [Katsuno and Mendelzon, 1992, p. 187] (see also Section 2.4 below).

(K°6) specifies that the revision process should be independent of the syntactic form of the sentences involved. It is called the *principle of irrelevance of syntax* by many authors, including [Dalal, 1988b].

(K°7) and (K°8) are the most interesting and controversial postulates. They try to capture the informational economy principle outlined before. In order to understand these postulates, consider the following semantical interpretation and assume one has some mechanism to evaluate similarity between models (i.e., valuations of the logic L). In order to keep as much as possible of the *informational content* of a belief set K, we need to look at the valuations that most resemble the models of K itself (in symbols, $\mathrm{mod}(K)$). If a new belief φ is also to be accepted, we will then be looking at the models of φ that most resemble *some* model of K. (K°7) says that if a model I of φ is among the valuations that are most similar to models of K and it happens that I is also a model of a belief ψ, then I should also be among the models of $\varphi \wedge \psi$ which are most similar to models of K.

Similarly, to understand the intuitive meaning of (K°8) consider the following situation: suppose that $(K \circ_a \varphi) \wedge \psi$ is satisfiable. It follows that some models of φ which are most similar to models of K are also models of ψ. These models are obviously in $\mathrm{mod}(\varphi \wedge \psi)$, since by (K°1), $\mathrm{mod}(K \circ_a \varphi) \subseteq \mathrm{mod}(\varphi)$. Now, every model in $\mathrm{mod}(\varphi \wedge \psi)$ which is most similar to a model of K must also be a model of $(K \circ_a \varphi) \wedge \psi$. This situation is depicted in Figure 1, where valuations are represented around K according to their degree of similarity. The closer it is to $\mathrm{mod}(K)$, the more similar to K is a valuation (the exact nature of the similarity notion is irrelevant to the understanding of the postulate now). The figure provides an illustration of (K°8) using Grove's modelling of theory change [Grove, 1988] presented in Section 2.3.

Another way of seeing (K°7) and (K°8) is by considering the restrictions they impose on the acceptance of beliefs φ and ψ as a sequence (revising by φ, then expanding by ψ), as compared to revising by φ and ψ at the same time (i.e., revising by $\varphi \wedge \psi$). One of the main criticisms to the AGM framework is the fact that they do not constrain enough properties of sequences of revisions. (K°7) and (K°8) impose the bare minimum restrictions (see Section 4 below). We distinguish the following three cases:

Case 1: φ is consistent with K.

In this case, $K \circ_a \varphi = \text{Cn}(K \cup \{\varphi\})$ (by previous postulates). Three possible subcases with respect to the sentence ψ are considered.

1. ψ is consistent with $K \circ_a \varphi$. In this case, the antecedent of (K°8) holds and (K°7) and (K°8) together effectively say that $\text{Cn}((K \circ_a \varphi) \cup \{\psi\}) = K \circ_a (\varphi \wedge \psi)$. A more thorough analysis reveals more about AGM in this case, namely, that $(K \circ_a \varphi) \circ_a \psi = \text{Cn}(K \circ_a \varphi \cup \{\psi\})$.

2. ψ is inconsistent with $K \circ_a \varphi$, but ψ itself is non-contradictory. In this case, $\text{Cn}((K \circ_a \varphi) \cup \{\psi\})$ is $K_\perp$. (K°7) holds because the right hand side of the inclusion is the set of all well-formed formulae and any set of formulae is included in this set. (K°8) holds vacuously, since the antecedent of the implication is false.

3. ψ is itself contradictory. The postulates effectively say nothing new in this case, since $K \circ_a (\varphi \wedge \psi) = \text{Cn}((K \circ_a \varphi) \cup \{\psi\}) = K_\perp$. (K°7) holds trivially and (K°8) holds vacuously.

Case 2: φ is not consistent with K, but φ is itself non-contradictory.

In this case, $K \circ_a \varphi$ can be any consistent theory (by previous postulates), such that $\varphi \in K \circ_a \varphi$. As before, there are three possibilities:

1. ψ is consistent with $K \circ_a \varphi$.

2. ψ is inconsistent with $K \circ_a \varphi$, but ψ itself is non-contradictory.

3. ψ is itself contradictory.

These three cases follow, respectively, the same reasoning of cases (1.1), (1.2) and (1.3) above.

Case 3: φ is itself contradictory.

In this case, $K \circ_a \varphi = K_\perp$. Whether or not ψ is contradictory is irrelevant in the postulates in this case. $\text{Cn}(K \circ_a \varphi \cup \{\psi\}) = K_\perp$ and as for case (1.2) above (K°7) holds because any set of formulae is included in $K_\perp$. (K°8) holds vacuously, since the antecedent of the implication is false.

Summary of (K°7)–(K°8)

Postulates (K°7)–(K°8) do not tell us anything new (beyond what we can deduce from earlier postulates), except in the case where ψ is consistent with $K \circ_a \varphi$ (case 1.1), when (K°7) and (K°8) together are equivalent to the postulate below:

($K^\circ_{7,8}$) If ψ is consistent with $K \circ_a \varphi$, then $\text{Cn}((K \circ_a \varphi) \cup \{\psi\}) = K \circ_a (\varphi \wedge \psi)$

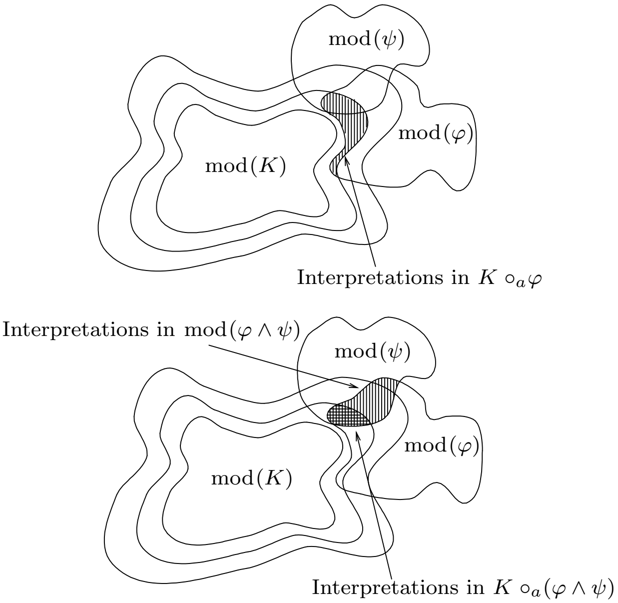

Figure 1. Illustrating postulate (K°8).

Several representation theorems exist for the AGM postulates, for more details the reader is referred to [Grove, 1988; Katsuno and Mendelzon, 1992; Boutilier, 1994].

The realisation of a belief set as an infinite set of formulae poses some problems for computer science applications. In order to overcome this, many authors concentrate instead on a set of basic beliefs from which the belief set is derived. In this case, the basic set of beliefs is called the *belief base* and the belief change process called *base revision* instead. A reformulation of the AGM postulates for finite belief bases will be discussed in Section 2.4.

2.2 Counterfactual statements and the Ramsey Test

There are references to the work on information change since the early 30's [Ramsey, 1931] as well as in subsequent decades [Chisholm, 1946; Stalnaker, 1968a; Stalnaker and Thomason, 1970; Lewis, 1973]. In particular, there is work on the so-called *counterfactual statements*. The best way to introduce the intuition behind counterfactual statements is by presenting an example borrowed from Lewis' book on the subject [Lewis, 1973]:

"If kangaroos had no tails, they would topple over."

Since the antecedent of the sentence is false, its evaluation as an implication in classical logic is trivially *true*. However the *intended* meaning of such a sentence is, as described by Lewis, something like "in any possible state of affairs in which kangaroos have no tails, and which resembles our actual state of affairs *as much as kangaroos having no tails permits it to*, kangaroos topple over" [Lewis, 1973].

The evaluation of such statements has been the object of investigation by many authors [Stalnaker, 1968a; Adams, 1975; Pollock, 1981; Nute, 1984]. Speaking generically, a counterfactual is a sentence of the form

(CNT) "If it were the case that φ, then it would also be the case that ψ."

Following Lewis, we will represent (CNT) by the expression $\varphi \Box\!\!\rightarrow \psi$. It is natural to ask how one should evaluate the truth-values of such sentences. The intended meaning described above suggests that one should accept the belief in φ, changing as little as possible one's current state of beliefs in order to maintain consistency, and then check whether ψ follows from the resulting belief set. This corresponds to the well known *Ramsey Test*, inspired by one of Ramsey's philosophical papers [Ramsey, 1990; Ramsey, 1931], and generalised to its present form by Stalnaker [Stalnaker, 1968b]. One could be easily mislead to think that belief revision could be the operation employed in the Ramsey Test, by taking $\varphi \Box\!\!\rightarrow \psi$ as accepted in a belief set K whenever ψ is accepted in $K \circ \varphi$. In symbols,

(RT) $K \vdash \varphi \Box\!\!\rightarrow \psi$ iff $K \circ \varphi \vdash \psi$

However, it is well known that belief revision cannot be used to evaluate counterfactual statements [Gärdenfors, 1986] [Gärdenfors, 1988, Section 7.4]. *Gärdenfors's impossibility theorem* showed us that whereas (RT) forces the belief change operation to be monotonic, belief revision is intrinsically non-monotonic. To see the first, assume (RT) is accepted, that $\circ$ is the belief change operation used to evaluate counterfactual statements and suppose that for belief sets K_1 and K_2, it is the case that $K_1 \subseteq K_2$. We show that $K_1 \circ \varphi \subseteq K_2 \circ \varphi$. Take any ψ such that $\psi \in K_1 \circ \varphi$. By (RT), $\varphi \Box\!\!\rightarrow \psi \in K_1 \therefore \varphi \Box\!\!\rightarrow \psi \in K_2$, and by (RT) again $\psi \in K_2 \circ \varphi$. To see that belief revision is incompatible with monotonicity, recall the following postulates and consider Example 5.

($K^\circ 2$) $\varphi \in K \circ_a \varphi$

($K^\circ_{3,4}$) If φ is consistent with K, then $K \circ_a \varphi = \mathrm{Cn}(K \cup \{\varphi\})$

($K^\circ 5$) $K \circ_a \varphi = K_\perp$ only if φ is contradictory

EXAMPLE 5. Consider three formulae φ, ψ and (non-contradictory) $\neg\varphi \vee \neg\psi$ and three belief sets K_1, K_2 and K_3, such that $K_1 = \mathrm{Cn}(\{\varphi\})$, $K_2 = \mathrm{Cn}(\{\psi\})$, and $K_3 = \mathrm{Cn}(\{\varphi, \psi\})$. It can be easily seen that $K_1, K_2 \subseteq K_3$. By ($\mathrm{K}^{\circ}_{3,4}$), $K_1 \circ_a(\neg\varphi \vee \neg\psi) = \mathrm{Cn}(\{\varphi, \neg\varphi \vee \neg\psi\}) = \mathrm{Cn}(\{\varphi, \neg\psi\})$; $K_2 \circ_a(\neg\varphi \vee \neg\psi) = \mathrm{Cn}(\{\psi, \neg\varphi\})$; and since $\neg\varphi \vee \neg\psi$ is non-contradictory, $K_3 \circ_a(\neg\varphi \vee \neg\psi)$ is satisfiable. However,

i) $\varphi \in K_1 \circ_a(\neg\varphi \vee \neg\psi)$

ii) $\neg\varphi \in K_2 \circ_a(\neg\varphi \vee \neg\psi)$

and hence, either $K_1 \circ_a(\neg\varphi \vee \neg\psi) \not\subseteq K_3 \circ_a(\neg\varphi \vee \neg\psi)$ or $K_2 \circ_a(\neg\varphi \vee \neg\psi) \not\subseteq K_3 \circ_a(\neg\varphi \vee \neg\psi)$, since $\{\varphi, \neg\varphi\} \not\subseteq K_3 \circ_a(\neg\varphi \vee \neg\psi)$.

In semantical terms, the reason can be understood by recalling Lewis's formulation of satisfiability of counterfactuals via systems of spheres. Let us first introduce the notion of a *centred system of spheres* [Lewis, 1973]:

DEFINITION 6 (Centred system of spheres). Let $\mathcal{I}$ be a set of worlds. A centred system of spheres \$ is an assignment from $\mathcal{I}$ to a set of subsets of $\mathcal{I}$, $\$_I$, where for each $I \in \mathcal{I}$:

① $\{I\} \in \$_I$. (*centring*)

② For all $S, T \in \$_I$, either $S \subseteq T$ or $T \subseteq S$. (*nesting*)

③ $\$_I$ is closed under unions.

④ $\$_I$ is closed under nonempty intersections.

Systems of spheres are used to represent the degree of similarity between worlds. The smaller a sphere containing a world J in $\$_I$ is, the closer to world I world J is. The centring condition ① can be interpreted as "there is no world more similar to world I than I itself". ① can be replaced by

①′ For all $S \in \$_I$, $I \in S$.

This condition is often called *weak centring.*[6] If, in addition to conditions ①–④ above, we also have that for all I, $\bigcup \$_I = \mathcal{I}$, then we say that \$ is *universal.*

In terms of a system of spheres \$, a world I satisfies $\varphi \Box\!\!\rightarrow \psi$, according to the following rules:

[6] Update operations as semantically characterised by Katsuno and Mendelzon in [Katsuno and Mendelzon, 1991b], require strong centring (i.e., the innermost sphere in $\$_I$ contains just I itself).

DEFINITION 7 (Satisfiability of counterfactuals via systems of spheres). Let $\mathcal{I}$ be a set of worlds, $I \in \mathcal{I}$ and \$ a centred system of spheres for $\mathcal{I}$:

$I \Vdash_{\$} \varphi \Box\!\!\rightarrow \psi$ iff

1. either $\forall S \in \$_I \ \mathrm{mod}(\varphi) \cap S = \emptyset$;
2. or $\exists S \in \$_I$ such that $\mathrm{mod}\,(\varphi) \cap S \neq \emptyset$, and $\forall I \in S,\ I \Vdash \varphi \rightarrow \psi$.

In case (1) above, we say that φ is not *entertainable* at I. That is, there is no sphere around I which intersects any worlds where φ is true. If \$ is universal, this happens only if $\mathrm{mod}(\varphi) = \emptyset$. The set of models of a counterfactual $\varphi \Box\!\!\rightarrow \psi$ can be defined as $\mathrm{mod}(\varphi \Box\!\!\rightarrow \psi) = \{I \in \mathcal{I} \,|\, \forall S \in \$_I\ (\mathrm{mod}(\varphi) \cap S \neq \emptyset$ implies $\forall J \in S\ J \Vdash \varphi \rightarrow \psi)\}$. As for case (2), since $\$_I$ is nested, it is sufficient to check whether $\varphi \rightarrow \psi$ is satisfied by every world in the innermost sphere S for which $S \cap \mathrm{mod}(\varphi)$ is non-empty. Intuitively, this intersection corresponds to the models of φ which are more similar (or closer) to I. Now, if we want to evaluate whether a counterfactual $\varphi \Box\!\!\rightarrow \psi$ is entailed by a belief set K, we have to check whether for each $I \in \mathrm{mod}(K)$, (2) holds, that is, whether the models of φ that are more similar to *each* of the models of K are also models of ψ.

It is not surprising that belief revision cannot be used to evaluate counterfactuals, since it fails to consider each model of K individually — which the operation of *update* does. Indeed, the relationship between counterfactual statements and updates has been pointed out many times [Grahne, 1991b; Rodrigues *et al.*, 1996; Ryan and Schobbens, 1996]. Updates can be used to evaluate conditional statements, and the properties of the resulting conditional logic will depend on the properties of the specific update operation considered.

An alternative way of evaluating counterfactuals is by employing an ordering relation on the set of worlds that determines similarity with respect to a given world I. This is done via the definition of a *comparative similarity system* [Lewis, 1973]:

DEFINITION 8 (Comparative similarity system). A *comparative similarity system* is a function that assigns to each world I a tuple $\langle \leq_I, S_I \rangle$, where S_I is a set of worlds, representing the worlds that are accessible from I; and $\leq_I$ is a binary relation on worlds, representing the comparative similarity of worlds with respect to I, such that

① $\leq_I$ is transitive

② $\leq_I$ is strongly connected

③ $I \in S_I$

④ For any world J, $J \neq I$ implies $I <_I J$

⑤ $K \notin S_I$ implies K is $\leq_I$-maximal

⑥ For any J, K, $J \in S_I$ and $K \notin S_I$ implies $J <_I K$.

The intended meaning for $\leq_I$ is the following: if $J \leq_I K$, then world J is at least as similar to world I as world K is.

Lewis proved that there is a correspondence between the satisfiability of counterfactuals via a system of spheres and their satisfiability via a comparative similarity system. If we consider only *universal comparative similarity systems* (i.e., $S_I = \mathcal{I}$), the truth conditions for counterfactuals can be simplified as follows:

DEFINITION 9 (Satisfiability of counterfactuals via comparative similarity systems). Let $\mathcal{I}$ be a set of worlds and take $I \in \mathcal{I}$. Given a comparative similarity system as in Definition 8:

$I \Vdash \varphi \Box\!\!\rightarrow \psi$ iff

1. either $\text{mod}(\varphi) = \emptyset$;
2. or $\exists M \in \text{ mod } (\varphi)$ such that for any $N \in \mathcal{I}$, $N \leq_I M$ implies $N \Vdash \varphi \rightarrow \psi$

These early results were very influential on the work of theory change carried out in the 80's and beyond.

We now turn to another very important semantical characterisation of belief change operations.

2.3 Grove's systems of spheres

In a very influential paper, Grove proposed a semantical characterisation of the AGM theory based on the so-called *systems of spheres* [Grove, 1988]. The idea is similar to that of Lewis' own systems of spheres presented in Section 2.2 above, except for a few modifications. Firstly, the spheres in Lewis' systems contain *worlds*, whereas in Grove's formulation they contain *theories*. In addition, Grove's systems of spheres can contain a *collection* of theories in their centre (a form of weak centring).

Interestingly enough, Grove was one of the first to notice that Lewis' formulation was incompatible with belief revision [Grove, 1988]. The relationship between the types of system of spheres proposed by Lewis and

Grove on the one hand and formalisms for theory change on the other was explored in more detail in [Katsuno and Mendelzon, 1991a; Katsuno and Mendelzon, 1992; Rodrigues *et al.*, 1996] and as it turns out only strongly centred systems of spheres can be used to model *updates* of a knowledge base in the reasoning about the effects of actions [Winslett, 1988b; Katsuno and Mendelzon, 1992].

The starting point in Grove's formulation is the set $M_L^\top$ of all maximal consistent sets of L. These in fact correspond to all (consistent) complete theories of L. Amongst these, some are of particular interest for a given (not necessarily complete) belief set K — the ones that *extend* it. The set of all such extensions is denoted by $|K|$ and formally defined as $\{m \in M_L^\top \mid K \subseteq m\}$. Notice that if K is $K_\perp$, then $|K|$ is simply $\emptyset$. Analogously, given a set S of maximal consistent sets of L, the set $t(S)$ is defined as $\bigcap\{S_i \in S\}$ or $K_\perp$ if $S = \emptyset$. It follows that $t(S)$ is also closed under logical consequence.[7]

In semantical terms, one can think of the set $\mathcal{I}$ of all valuations instead of $M_L^\top$. Analogously, $|K|$ would correspond to $\mathrm{mod}(K)$ and for a given set of valuations $S \subseteq \mathcal{I}$, $t(S) = \{\varphi \mid I \vDash \varphi \text{ for all } I \in S\}$. According to this view of the formulation, if K is $K_\perp$, then $|K| = \mathrm{mod}(K) = \emptyset$ and if $K = \mathrm{Cn}(\emptyset)$, then $|K| = \mathcal{I}$, as expected. However, viewing the revision process in terms of sets of formulae as done by Grove makes the relationship with the AGM postulates immediate, whereas viewing it semantically, i.e., in terms of valuations, gives us an interesting insight into the process.[8]

DEFINITION 10 (Grove's systems of spheres). Let $\mathcal{S}$ be a collection of subsets of $M_L^\top$ and take $S \subseteq M_L^\top$. $\mathcal{S}$ is a called a *system of spheres centred on* S if it satisfies the following conditions:

(S1) $\mathcal{S}$ is totally ordered by $\subseteq$

(S2) S is the $\subseteq$-minimum of $\mathcal{S}$

(S3) $M_L^\top$ is the $\subseteq$-maximum of $\mathcal{S}$

(S4) For any wff φ, if $M_L^\top \cap |\varphi| \neq \emptyset$, then there is a smallest sphere in $\mathcal{S}$ intersecting $|\varphi|$.

[7] It is possible to construct a lattice by ordering theories of a logic L under set inclusion. In this case, $\mathrm{Cn}(\emptyset)$ will be the minimum and $K_\perp$ will be the maximum. The only inconsistent theory in the lattice is $K_\perp$ itself with the elements of $M_L^\top$ sitting immediately below it. The only way to extend an element of $M_L^\top$ and retain closure under logical consequence is to jump to $K_\perp$.

[8] The semantical view is essentially the basis for the work by Katsuno and Mendelzon seen in Section 2.4 [Katsuno and Mendelzon, 1991a; Katsuno and Mendelzon, 1992].