SPRINGER BRIEFS IN INTELLIGENT SYSTEMS
ARTIFICIAL INTELLIGENCE, MULTIAGENT SYSTEMS,
AND COGNITIVE ROBOTICS

Eduardo Fermé
Sven Ove Hansson

Belief Change
Introduction and Overview

Springer

SpringerBriefs in Intelligent Systems

Artificial Intelligence, Multiagent Systems, and Cognitive Robotics

This series covers the entire research and application spectrum of intelligent systems, including artificial intelligence, multiagent systems, and cognitive robotics. Typical texts for publication in the series include, but are not limited to, state-of-the-art reviews, tutorials, summaries, introductions, surveys, and in-depth case and application studies of established or emerging fields and topics in the realm of computational intelligent systems. Essays exploring philosophical and societal issues raised by intelligent systems are also very welcome.

More information about this series at http://www.springer.com/series/11845

Eduardo Fermé · Sven Ove Hansson

Belief Change

Introduction and Overview

Springer

Eduardo Fermé
Faculty of Exact Sciences
and Engineering
University of Madeira
Funchal
Portugal

Sven Ove Hansson
Division of Philosophy
Royal Institute of Technology (KTH)
Stockholm
Sweden

ISSN 2196-548X ISSN 2196-5498 (electronic)
SpringerBriefs in Intelligent Systems
ISBN 978-3-319-60533-3 ISBN 978-3-319-60535-7 (eBook)
https://doi.org/10.1007/978-3-319-60535-7

Library of Congress Control Number: 2018942194

Printed on acid-free paper

This Springer imprint is published by the registered company Springer International Publishing AG part of Springer Nature
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Preface

Belief change (belief revision) is a research area in formal philosophy that makes use of logic to produce models of how human and artificial agents change their beliefs in response to new information. The properties of these models are investigated in order to improve our understanding of how beliefs can be changed, with a particular emphasis on what it means to change one's beliefs in a rational way. This area of research arose in the 1970s and 1980s through a confluence of several approaches. Its first great synthesis was an article in the *Journal of Symbolic Logic* in 1985 (the so-called AGM paper). In the three decades that have followed since then, the area has developed rapidly and in many directions. This book is an introduction to the subject and at the same time an overview of some of its major ramifications. The model proposed in the 1985 paper will be used as a starting point, but we will also pay much attention to the criticism that has been raised against it and to several alternative models that have been proposed to supplement or replace it. Theorems and other formal results are presented without proofs; instead we give references to the research papers where the proofs can be found.

Belief change has turned out to be an unusually rich research area that provides us with new, dynamic ways to understand human rationality. We hope that the reader will see how full this field of research is of exciting openings for new discoveries.

In preparing the contents of this book we benefited from the help of more than fifty colleagues who answered our queries and provided us with information. In addition to its academic excellence, the belief revision community is a remarkably generous one. Special thanks go to Mauricio Reis, Rafael Testa, Hans van Ditmarsch and Ramón Pino Pérez for their help with Chapters 9, 10, 12 and 13 respectively.

Funchal and Stockholm,
November 2017

Eduardo Fermé[1]
Sven Ove Hansson

[1] In his work with this book Eduardo Fermé was partially supported by FCT UID/CEC/04516/2013 (NOVA LINCS),FCT SFRH/BSAB/127790/2016, and FAPESP 2016/13354-3

Contents

Chapter 1
Motivation

This is a brief introduction to the theory of belief change. It provides an example of a belief change problem, and lists some of the major issues that are investigated in this research area.

1.1 An Example of the Belief Change Problem

We consider the following set of sentences in natural language [94]: "Juan was born in Puerto Carreño" (α), "José was born in Puerto Ayacucho" (β), and "Two people are compatriots if they were born in the same country" (γ). We assume that this set represents all the currently available information about Juan and José. Suppose that we receive the following piece of new information: "Juan and José are compatriots" (δ). If we add the new information to our corpus of beliefs, then we obtain a new set of beliefs that contains the sentences α, β, γ and δ. We can define an operation of addition as one that takes a sentence and a set of previous beliefs and returns the minimal set that includes both the previous beliefs and the new sentence. This operation exemplifies the simplest way of changing a set of sentences. There are other types of change that are not that simple.

For example, suppose that upon consulting an atlas we discover to our surprise that Puerto Carreño is in Colombia (ϵ) and Puerto Ayacucho is in Venezuela (ϕ). If we add ϵ and ϕ to the set $\{\alpha,\beta,\gamma,\delta\}$, the result will be a set with contradictory information: Juan and José are compatriots but Puerto Carreño and Puerto Ayacucho do not belong to the same country. The addition does not satisfactorily reflect the notion of a *consistent revision*. If we wish to retain consistency, then some subset of the original set must be discarded or perhaps a part of the new information has to be rejected. In our example, there are several possible alternatives. The information about Juan's or José's birthplace could be wrong, and so could the atlas. Finally the claim that Juan and José are compatriots could be wrong. Any of these three options, either individually or combined, will allow us to solve the problem of the incompatibility among the original and the new information or beliefs. Consequently, we

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_1

can specify an operation that takes a set and a sentence and returns a new consistent set. The new set includes parts (or all) of the beliefs in the original set and it also includes the new sentence (if we are willing to accept it). The outcome of a revision can be expressed as a consistent subset of the outcome of the addition. This operation is based on two notions: *consistency* and a *selection* among the possible ways to perform the change.

There are other ways to change a set of beliefs. Suppose that we discover that γ is incorrect, and therefore wish to discard it from our set. The result should be a new set where γ is absent. We may for instance want it to be undetermined whether Juan and José are compatriots. Note that this is different from accepting as a fact that Juan and José are not compatriots. We can ask if the process of discarding information should behave as the inverse of the process of adding information: If after discarding some information we proceed to add it again, will we obtain the original set or not? Like revision, the operation of discarding requires the selection of one out of several possible results.

1.2 Some Questions About the Belief Change Problem[1]

Any formalization of belief change requires the selection of a language in which the beliefs are represented. In our previous example the information about Juan and José is represented by a set of sentences in natural language. The use of a linguistic representation of beliefs implies the acceptance of important idealizations. Whatever language is chosen, the question emerges how to use the language to represent the epistemic state: should it be represented by a single sentence or by a set (perhaps an infinite set) of sentences? In the latter case, should the set be closed under some notion of logical consequence or should it only be a simple enumeration of sentences? The second option implies the need to obtain in some way the consequences of these sentences and to differentiate between implicit and explicit information.

Can the belief state be changed spontaneously or does change require an external stimulus? In other words, is the belief state internally stable? If the belief state is changed only in response to external stimuli, should the belief state and the information that provokes the change be represented by the same or different types of formal structures? Should both be sentences or both be sets of sentences? How should the sentences be interpreted? If an epistemic interpretation of the sentences is chosen, what are the possible statuses of the sentences? Acceptance, rejection, indeterminateness, or perhaps degrees of acceptability? What types of information can be represented in the belief state?

Generally speaking, it seems to be fundamental to define operations that answer to the notion of minimal change, or maximal preservation of the belief state. That is to say, it is required in some way to "calculate the value" of the information to

[1] Borrowed from [11].

be discarded. Does a preference order exist that represents the credibility or informational value of expressions in the language? Is this order included in the belief state or is it intrinsic to the change operation? Should minimal change be defined quantitatively or qualitatively?

In what ways can a belief state be modified? Are they independent or interrelated? What is the relationship between the original and the updated belief state? How should an operation to revise the original belief state be constructed? What are the parameters of this operation? The original belief state and the new information are obvious such parameters, but are there any other parameters? Should the change operation take into account the history of the produced changes, or is each new change performed independently of those performed earlier?

These kinds of questions have encouraged several authors to propose different belief change models and to assume some of the above options and discard others. By far the most influential of these models was proposed by Carlos Alchourrón (1931–1996), Peter Gärdenfors, and David Makinson in their paper "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions". Many research papers have been called "seminal", but few deserve that designation as much as this article in the *Journal of Symbolic Logic* in 1985. It was the starting point of a large and rapidly growing literature that employs formal models in the investigation of changes in belief states and databases.

This book is an introduction on and an overview of the research that has been inspired by the AGM article.

Chapter 2
History

This is a brief history of belief change theory, showing how it emerged in the 1980s from work in philosophy and computer science and how it has impacted further developments in these two disciplines.

2.1 The Beginnings

In a wide sense, belief change has been a subject of philosophical reflection since antiquity. In the twentieth century, philosophers discussed the mechanisms by which scientific theories develop, and they proposed criteria of rationality for revisions of probability assignments. Beginning in the 1970s a more focused discussion of the requirements of rational belief change has taken place in the philosophical community.

Four authors had a big influence in the beginnings of the area: Georg Henrik von Wright (1916–2003), Jaakko Hintikka (1929–2015), William Harper and Isaac Levi. von Wright laid the foundations of the logic of beliefs, which was considerably developed and extended by Hintikka [325, 199]. Harper [197] pointed out that Carnap's inductive logic does not allow us to revise previously accepted evidence and created a model to accommodate rational conceptual change. Levi [220] provided much of the basic formal framework of belief change. Carlos Alchourrón (1931–1996) and David Makinson cooperated in studies of changes in legal codes, analysing the logical structure of the derogation procedure in which a norm is removed from a legal code. They tried to find the general principles that any derogation should satisfy, and defined a family of all the possible derogations [5]. The key idea was, given a code A, to create a partial order on the norms of A and induce an order on the set of parts of A. The maximal sets of A that did not imply the norm to be removed were called *remainders*. Later they extended the horizon of the problem, arguing that the problem was not limited only to sets of norms. The set A might be an arbitrary set of formulae, and the problem was how to eliminate one of the formulae or one of the consequences of the set [6]. Two different ways

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_2

to contract a theory by means of remainder sets were analyzed: *maxichoice* and *full meet*.

Peter Gärdenfors's early work was concerned with the connections between belief change and conditional sentences (if-sentences). He was looking for a model of explanations. Gärdenfors thought that explanations can be expressed as different types of conditional sentences. He was influenced by Levi and Harper (see above), and this led him to make a thorough study of epistemic conditionals [128]. Gärdenfors constructed a semantic account of epistemic conditionals that is based on belief states and belief changes [129]. He defined a set of postulates that change functions must satisfy [130].

Gärdenfors's postulates were closely related with the ideas developed by Alchourrón and Makinson. With combined forces the three wrote a paper that provided a new, much more general and versatile formal framework for studies of belief change, now known as the AGM model. The cooperation of the three philosophers was explained by Gärdenfors in [137]:

> "I became acquainted with the work of Alchourrón and Makinson on *derogations* of legal systems first via their 1981 paper on "Hierarchies of regulations and their logic". When they in 1982 submitted an article to *Theoria* with the title "On the logic of theory change: Contraction functions and their associated revision functions" I was then the managing editor of the journal and I immediately saw the similarities between our programs. We soon joined forces and via a series of letters, sent between Buenos Aires, Lund and Beirut, we developed what was to become the AGM paper." [137]

Since the paper was published in 1985, its major concepts and constructions have been the subject of significant elaboration and development.

2.2 The Impact in Philosophy

Belief change theory has opened up a new area for philosophical reflection. Many of the issues mentioned in Section 1.2 are philosophical questions that became accessible to study through the introduction of formal models of belief change. Precise philosophical discussions–for instance, about what types of belief changes there are, whether some of them can be defined in terms of the others [115, 146, 182], how (changes in) all-or-nothing beliefs relate to (changes) in probabilistic beliefs [244], and whether belief changes can be perfectly reversible [161]–have become possible through the introduction of formal models of belief change. In addition, new light has been thrown on philosophical issues that were already discussed long before belief change theory was developed. We will mention two particularly important examples of this.

First, the discussion on the meaning of conditional sentences turned out to be closely related to the logic of belief change. An important class of conditional sentences seems to satisfy the so-called Ramsey test. The test stipulates that a rational agent believes in the sentence "If α then β" if and only if she would believe in β if she were brought to believe in α. Due to this connection, the logic of belief change provided new tools for studies of the properties of conditional sentences. (We will return to this in Section 11.3.)

Secondly, philosophers of science have long discussed the mechanisms of theory change and theory replacement in science. Although belief change models have usually been constructed to represent individual belief systems, they can also be applied to joint or collectively held thought systems such as scientific theories. Recently, belief change models have been used to investigate this and other issues in the philosophy of science [263].

2.3 The Impact in Artificial Intelligence

The AGM model appeared at a moment when the area of artificial intelligence was suffering from a crisis. This crisis was clearly elucidated by Allen Newell in his seminal address *The Knowledge Level*, delivered when he assumed the presidency of the American Association for Artificial Intelligence [258]. Newell pointed out three indicators of the crisis:

> "A first indicator comes from our continually giving to representation a somewhat magical role. What is indicative of underlying difficulties is our inclination to treat representation like a homunculus, as the locus of real intelligence."

> "A second indicator is the great theorem-proving controversy of the late sixties and early seventies. Everyone in AI has some knowledge of it, no doubt, for its residue is still very much with us."

> "The results of a questionnaire promoted in 79/80 by Brachman and Smith which was sent to the AI community: 'The main result was overwhelming diversity–a veritable jungle of opinions. There was no consensus on any question of substance.' [...] As one [of the respondents] said, 'Standard practice in the representation of knowledge is the scandal of AI.'"

Newell claimed that Knowledge Representation and Reasoning (KRR) must be a priority on the AI agenda and postulated the existence of a *knowledge level*:

> "[...] there exists a distinct computer systems level, lying immediately above the symbol level, which is characterized by knowledge as the medium and the principle of rationality as the law of behavior. "

The knowledge level rationalizes the agent's behaviour, while the symbol level mechanizes it. The knowledge level is more related with the question *what?* whereas the symbol level is more related with *how?* For example, in a computer system, a knowledge level is related with the functional requirements and the symbol level with the non-functional requirements and implementation features.

Newell's proposal had an enormous influence on AI community. In 1983, Fagin, Ullman and Vardi pointed out the necessity of defining the dynamics of the process of update in a database [86]:

> "The ability of the database user to modify the content of the database, the so-called *update* operation, is fundamental to all database management systems. "

> "[W]e consider the problem of updating arbitrary theories by inserting into them or deleting from them arbitrary sentences. The solution involves two key ideas: when replacing an old theory by a new one we wish to minimize the change in the theory, and when there are several theories that involve minimal changes, we look for a new theory that reflects that ambiguity."

In his keynote at the Theoretical Aspects of Rationality and Knowledge Workshop in 1986 [145], Joseph Halpern also refered to the dynamics of beliefs:

> "Most of the work discussed above has implicitly or explicitly assumed that the messages received are consistent. The situation gets much more complicated if messages may be inconsistent. This quickly leads into a whole complex of issues involving belief revision and reasoning in the presence of inconsistency. Although *I won't attempt to open this can of worms* here, these are issues that must eventually be considered in designing a knowledge base."

The principles exposed by the cited authors were very similar to those expressed by the AGM trio for belief states. The rapid propagation of the AGM ideas in the AI community after Gärdenfors and Makinson presented them on TARK 88 [138] was no surprise. The can of worms was open.[1]

[1] For a more thorough overview of the impact of AGM in artificial intelligence see [59].

Chapter 3
The AGM Model

The purpose of this chapter is to introduce the AGM account of belief change, originally developed by Alchourrón, Gärdenfors and Makinson [4]. In Sections 3.1–3.3 we introduce the formal apparatus of belief sets and in Section 3.4 the operations of change. In Sections 3.5–3.6 we introduce the axioms of the AGM model. In Section 3.7 the relations between contraction and revision are specified and in Section 3.8 we introduce the basic constructive method of the AGM model, partial meet contraction and revision functions.

3.1 The Language

Beliefs are expressed in a language $\mathcal{L}$ that is called the object language of our model. As is common in logic, we consider the language to be identical to the set of sentences that can be expressed in it; thus $p \in \mathcal{L}$ if and only if p is a sentence in $\mathcal{L}$.

The language may be either finite or infinite, unless we explicitly specify that it is finite. We also assume that the language contains the usual truth functional connectives: negation ($\neg$), conjunction ($\wedge$), disjunction ($\vee$), implication ($\rightarrow$) and equivalence ($\leftrightarrow$). $\bot$ denotes an arbitrary contradiction and $\top$ an arbitrary tautology. $\mathcal{L}$ is closed under truth-functional operations (for example, if $p \in \mathcal{L}$ and $q \in \mathcal{L}$, then $p \vee q \in \mathcal{L}$, etc.). We say that two sentences p and q are logically independent if and only if all combinations of truth values are logically possible for them. The letters $p, p_i, q, \ldots$ will be used to denote sentences of $\mathcal{L}$. $A, A_i, B, \ldots$ denote sets of sentences in $\mathcal{L}$. K and H are reserved to represent sets of sentences that are closed under logical consequence, i.e., each such set contains all sentences that follow logically from it. Such a set is called a *belief set* or *theory*. $\mathcal{K}$ is the set of belief sets.

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_3

3.2 Logical Consequence

To express the logical relationships among the sentences in the object language $\mathcal{L}$ we introduce a *consequence operation* Cn. It is a function that takes us from any set of sentences X in $\mathcal{L}$ to the set $Cn(X)$ that consists, intuitively speaking, of all the logical consequences of X. We will assume that the logic encoded by Cn includes classical truth-functional logic, i.e., the elementary logic that is usually taught through truth tables. This means for instance that if $p \in X$ then $p \vee s \in Cn(X)$ for all sentences $s \in \mathcal{L}$, and it also means that if both $p \in X$ and $q \in X$, then $p \wedge q \in Cn(X)$. There may also be other logical principles encoded in Cn, in addition to classical truth-functional logic, but it will be left open what these logical principles are. Formally, the consequence operation is introduced as follows:

Definition 3.1. [321] A *consequence operation* on a language $\mathcal{L}$ is a function Cn that takes each subset of $\mathcal{L}$ to another subset of $\mathcal{L}$, such that:

- $A \subseteq Cn(A)$ (inclusion)
- $Cn(A) = Cn(Cn(A))$ (iteration)
- If $A \subseteq B$, then $Cn(A) \subseteq Cn(B)$. (monotony)

We are going to assume that Cn satisfies the following three properties:

- If p can be derived from A by classical truth-functional logic, then $p \in Cn(A)$. (supraclassicality)
- $q \in Cn(A \cup \{p\})$ if and only if $(p \rightarrow q) \in Cn(A)$. (deduction)
- If $p \in Cn(A)$, then $p \in Cn(A')$ for some finite subset A' of A. (compactness)

We write $Cn(p)$ for $Cn(\{p\})$ when $p \in \mathcal{L}$. We use $\vdash p$ as an alternative notation for $p \in Cn(\varnothing)$, $A \vdash p$ for $p \in Cn(A)$ and $p \vdash q$ for $q \in Cn(p)$. Note that $K \vdash p$ if and only if $p \in K$. We will use both notations interchangeably. $K_\perp$ denotes the inconsistent belief set, and it follows that $K_\perp = \mathcal{L}$.

The consequence operation satisfies the following properties [163]:

1. If Cn satisfies *iteration, monotony, supraclassicality*, and *deduction* then $Cn(p \vee q) = Cn(p) \cap Cn(q)$.
2. If Cn satisfies *iteration, monotony, supraclassicality*, and *deduction* then: If $q \in Cn(A \cup \{p_1\})$ and $q \in Cn(A \cup \{p_2\})$, then $q \in Cn(A \cup \{p_1 \vee p_2\})$. (introduction of disjunction into premises).
3. If Cn satisfies *deduction* then: $Cn(A) \vdash \neg p$ if and only if $Cn(A \cup \{p\}) \vdash \perp$.
4. If Cn satisfies *iteration* and *monotony* then: If $A \subseteq B \subseteq Cn(A)$ then $Cn(A) = Cn(B)$.
5. If Cn satisfies *monotony* then $Cn(A) \cap Cn(B) = Cn(Cn(A) \cap Cn(B))$.
6. If Cn satisfies *inclusion, iteration*, and *monotony* then $Cn(A \cup B) = Cn(A \cup Cn(B))$.

7. If Cn satisfies *inclusion, iteration*, and *monotony* then: if A and B are logically closed, then $A \cup B$ is logically closed if and only $A \subseteq B$ or $B \subseteq A$.[1]
8. If Cn satisfies *inclusion, iteration, monotony*, and *supraclassicality* then: If $p \in Cn(A)$ and $\neg p \in Cn(A)$, then $Cn(A) = \mathcal{L}$.
9. If Cn satisfies *inclusion, iteration, monotony*, and *supraclassicality* then $Cn(\{p \wedge q\}) = Cn(\{p, q\})$.

3.3 Belief Sets and Possible Worlds

As already mentioned, to represent an epistemic agent's belief state we will use a set of sentences that is logically closed, i.e., it contains all its own logical consequences. Such a set will be called a *belief set* and denoted by K. Its characteristic logical property is:

$$K = Cn(K)$$

Obviously, the use of a logically closed set to represent an individual's belief state is an idealization, since the individual may fail to draw all the logical conclusions that follow from her beliefs. However, it should be observed that Cn does not include all of logic (or mathematics); we have left it open whether it includes anything more than classical truth-functional logic. It should also be observed that belief change theory is mainly concerned with rational belief change, and rationality can be conceived as requiring that one draws the logical conclusions that are available.

Another important class of subsets of $\mathcal{L}$ are its inclusion-maximal consistent subsets, more commonly called *possible worlds*. A subset μ of $\mathcal{L}$ is a possible world if and only if it satisfies two conditions. First, μ is consistent, i.e., $\bot \notin Cn(\mu)$. Secondly, it is so large that no sentence can be added to it without making the new set inconsistent, i.e., if $p \notin \mu$ then $\bot \in Cn(\mu \cup \{p\})$. The set of possible worlds will be denoted by $\mathfrak{W}$.

Two important properties of possible worlds should be noted. First, they are logically closed, i.e., if μ is a possible world then $\mu = Cn(\mu)$. Secondly, they are determinate in the sense that for each sentence $p \in \mathcal{L}$, either $p \in \mu$ or $\neg p \in \mu$. It is due to the latter property that they are called possible worlds; they give a complete description (as far as the language $\mathcal{L}$ allows) of what is true or false in a particular state of the world.

There is an important two-way connection between belief sets and sets of possible worlds. In one direction, the intersection of any set of possible worlds (i.e., the set of sentences that are elements in all these worlds) forms a belief set; thus if $\mathfrak{X}$ is a set of possible worlds, then $\bigcap \mathfrak{X}$ is a belief set. In the other direction, for any belief set K we can identify the set of possible worlds that contain it, namely

[1] $A \cup B = Cn(A \cup B)$ is not true in general even if A and B are logically closed: Let p and q be logically independent sentences, $A = Cn(\{p\})$ and $B = Cn(\{p \rightarrow q\})$. Then $q \notin A \cup B$, but $q \in Cn(A \cup B)$.

$\{\mu \in \mathfrak{W} \mid K \subseteq \mu\}$. The following observation connects belief sets and possible worlds tightly to each other:

Observation 3.1. Let K be a belief set. Then $\bigcap\{\mu \in \mathfrak{W} \mid K \subseteq \mu\} = K$. (For a proof, see [163] p. 52.)

This observation provides us with a one-to-one correspondence between, on the one hand, belief sets, and on the other hand sets of possible worlds. Therefore, sets of possible worlds (often called propositions) can be used as an equivalent alternative to belief sets for representing belief states. It is conventionally assumed that $\bigcap\varnothing = \mathcal{L}$, which means that the inconsistent belief set also corresponds to a subset of $\mathfrak{W}$, namely the empty set. In what follows we will use the following notation (for any belief set K and any sentence p):

$\| K \| = \{\mu \in \mathfrak{W} \mid K \subseteq \mu\}$ is the set of K-worlds.

$\| p \| = \{\mu \in \mathfrak{W} \mid p \in \mu\}$ is the set of p-worlds.

Both belief sets and sets of possible worlds have been used extensively in the AGM framework. As we will see, they can both be used to construct models of belief change. But before studying these models we will consider the syntactical approach to AGM, i.e., the approach that focuses on connections among sentences describing what is believed or not believed before and after a belief change.

3.4 Basic Ideas of AGM Theory

In the AGM model, belief sets are used to represent epistemic states.[2] Provided that the belief set is consistent, the epistemic agent can have exactly three epistemic attitudes to a sentence p, each defined from the belief set:

belief in p if $p \in K$.
disbelief in p if $\neg p \in K$.
unsettledness about p if $p \notin K$ and $\neg p \notin K$.

The epistemic state is assumed to be internally stable, and all changes result from inputs. Inputs in the AGM framework always take the form of a sentence together with an instruction on what to do with it. This is either an instruction to include or not to include the sentence in the resulting new belief set. Based on the threefold classification into belief, disbelief, and unsettledness, there are six forms of belief change, as illustrated in Figure 3.1. The operations denoted by $+p$ and $+\neg p$ are expansions. They consist in the addition to a belief set of a sentence that does not contradict it. Nothing has to be removed in order to retain consistency. The operations denoted

[2] This section is based on [181].

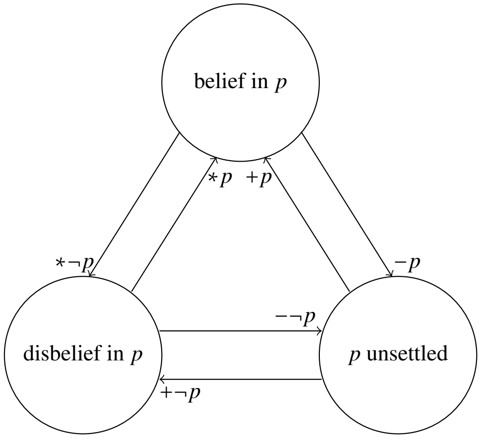

Fig. 3.1: The six types of belief change (assuming $\nvdash p$ and $\nvdash \neg p$).

$*p$ and $*\neg p$ are revisions. They consist in the addition to a belief set of a sentence that contradicts it. In order to retain consistency, some previous beliefs will have to be removed. Finally, the operations $-p$ and $-\neg p$ are contractions. They consist in the removal of a sentence from a belief set without introducing its negation. These are the major forms of belief change in the AGM model, and the symbols $+$, $*$ and $-$ are used to denote them. In performing these operations, the following general rationality criteria are applied as far as possible:

- *logical closure*: The outcome is a logically closed set, just like the original belief set.[3]
- *success*: (i) A sentence to be added is included in the outcome. (ii) A sentence to be contracted is not included in the outcome.
- *consistency preservation*: The outcome is consistent, just like the original belief set.
- *conservatism*: (i) When a new sentence is added, no sentences are removed. (ii) In contraction, no sentences are added [181].
- *informational economy*: As much as possible of the previous information is retained.
- *non-arbitrariness*: If there is more than one candidate for the revised belief set that satisfies the other rationality criteria, then one of them should not be arbitrarily chosen.

In contraction, these requirements are all compatible if the sentence to be removed is non-tautologous. If the sentence to be removed is a tautology, then *logical closure* and *success* are incompatible (but each of them is compatible with the other conditions). In the AGM model this is solved by giving higher priority to logical closure

[3] This is a special case of the principle of categorial matching, according to which the belief state should be represented in the same way after as before the operation of change [139].

than to success, i.e., the outcome of contraction by a tautology is a logically closed set and therefore it does not satisfy the success criterion. ($K - \top = K$.)

For changes consisting in adding a new sentence, we have to distinguish between three cases. First, if the sentence to be added is consistent with the original belief set (i.e., if $K \cup \{p\}$ is consistent), then all requirements are compatible. This is the expansion case referred to above. It is often treated as a limiting case of revision.

Secondly, if p is inconsistent, then *consistency preservation* and *success* cannot both be satisfied. This is traditionally solved by giving priority to *success* (which is compatible with the other conditions). Thus $K * \bot = \mathcal{L}$.

Thirdly and finally, if p is consistent in itself but inconsistent with the original belief set, then any two of the three conditions *consistency preservation*, *success*, and *conservatism* are compatible, but not all three of them. (*Logical closure* is compatible with each of these combinations.) There are two standard solutions to this. One is to give up *consistency preservation*, usually by just letting the outcome be $Cn(K \cup \{p\})$. This operation is called expansion. The other solution is to give up *conservatism*, and remove enough elements from the original belief set K to ensure that p can be added without giving rise to inconsistency. This type of operation is called revision.

3.5 A Syntactic Approach to Contraction

A contraction of a belief set by a sentence p requires that p be removed but no new belief be added. We write $-$ to denote a contraction function from $\mathcal{L}$ to $\mathcal{K}$. Hence, $K - p$ denotes the contraction of K by a sentence p.

According to the principle of categorial matching, when we contract a belief set K by a sentence p, the outcome should be logically closed:

$K - p = \mathrm{Cn}(K - p)$ (closure)

No new belief is added to the belief set in the contraction:

$K - p \subseteq K$ (inclusion)

As far as possible, the objective of the operation must be carried out; i.e., if it is possible to eliminate the sentence from the theory then it must be eliminated. The only sentences that cannot be contracted are the tautologies.

If $\nvdash p$ then $p \notin K - p$. (success)

Together, closure, success, and inclusion signify that the outcome of a contraction is a belief set included in the original belief set, which does not, if this can be avoided, contain the sentence to be contracted. These conditions provide a convenient demarcation of contractions from other types of belief change [163]:

Definition 3.2. [163] An operation $-$ for a belief set K is a contraction if and only if it satisfies *closure, inclusion* and *success*.

Next, let us consider two limiting cases (extreme cases) of contraction. In the first of these, p is logically true. We have already made an exemption from the success postulate for this case. When we are presented with the impossible task to take away something that cannot be taken away, the most reasonable solution is to leave the original belief set unchanged.

If $\vdash p$ then $K - p = K$. (failure) [121]

The second limiting case is when the sentence to be contracted is not implied by the original belief set. In this case the minimal way to eliminate p from K is to do nothing:

If $K \nvdash p$ then $K - p = K$. (vacuity)

The contraction operation should be independent of the syntactic representation of the sentences. In other words, logically equivalent sentences should yield the same result:

If $\vdash p \leftrightarrow q$ then $K - p = K - q$. (extensionality)

Extensionality guarantees that the logic of contraction is extensional in the sense of allowing logically equivalent sentences to be freely substituted for each other.

Definition 3.3. [239] An operation $-$ for a belief set K is a *withdrawal* if and only if it satisfies *closure, success, inclusion, vacuity* and *extensionality*.

The criterion of informational economy requires that $K - p$ be a *large* subset of K. For example,

$$K - p = \begin{cases} K & \text{if } p \notin K \\ Cn(\varnothing) & \text{otherwise} \end{cases}$$

satisfies the five postulates for withdrawal. However, such an operation seems extreme, since for all non-vacuous cases of contraction it returns only the minimal theory. According to the AGM rationality criteria, contraction should not only be successful, it should also be conservative in the sense of not leading to unnecessary losses of previous beliefs. An additional postulate is necessary to achieve that. The AGM trio proposed as a rule of conservativity the postulate of *recovery*, according to which it is enough to add (by expansion) the eliminated sentence to recover totally the original theory.

$K \subseteq (K - p) + p$ (recovery)[4]

Recovery is the most debated postulate of belief change [239, 99]. We will turn to this debate in Section 5.1. The postulates discussed thus far may be called *elementary* (or basic) since they concern changes (contractions) by one single sentence. Another category of postulates, to which we will now turn, is those that compare changes by different sentences. They will be called *comparative* (or supplementary) postulates. The AGM trio provided comparative postulates for contraction in the form of postulates for contraction by a conjunction. In order to contract a conjunction $p \wedge q$ from a theory K, we must either cease believing p or cease believing q. Now, if a sentence r in K is neither removed in the contraction of K by p nor in the contraction of K by q, then r should not be removed in the contraction of K by $p \wedge q$:

$(K - p) \cap (K - q) \subseteq K - (p \wedge q)$ (conjunctive overlap)

On the other hand, if p is removed if we contract by $p \wedge q$, then we expect that if a sentence r must to be removed in order to remove p then it will also be removed when $p \wedge q$ is contracted:

If $p \notin K - (p \wedge q)$ then $K - (p \wedge q) \subseteq K - p$. (conjunctive inclusion)

The last two postulates are called the supplementary AGM (or Gärdenfors) postulates. In presence of the basic postulates, the supplementary postulates are equivalent to the following:

$$K - (p \wedge q) = \begin{cases} K - p, \text{ or} \\ K - q, \text{ or} \\ K - p \cap K - q \end{cases} \qquad \text{(conjunctive factoring)}$$

Observation 3.2. [4] Let K be a belief set and $-$ an operation on K that satisfies *closure, inclusion, vacuity, success, extensionality*, and *recovery*. Then $-$ satisfies both *conjunctive overlap* and *conjunctive inclusion* if and only if $-$ satisfies *conjunctive factoring*.

The intuition behind this observation and the conjunctive factoring postulate is one of the pillars of the AGM theory. If we wish to contract the belief set by a conjunction and there exists some preference between the conjuncts, then this contraction is equivalent to contraction by the non-preferred conjunct. In the case of indifference between the conjuncts, the outcome of contracting by the conjunction equals the intersection of the outcomes of contractions by the conjuncts.

[4] The converse of *recovery*, $(K - p) + p \subseteq K$, follows from *inclusion* for the case when $p \in K$.

3.6 A Syntactic Approach to Revision

Revision is related to expansion, in the sense that it incorporates new beliefs. However, as opposed to expansion, consistency is preserved in revision (unless the new information is itself inconsistent). Consequently, the revision process must eliminate enough sentences to avoid contradiction with the new belief. Just as for contraction, there is no plausible way to define revision uniquely, but it can be constrained by a set of postulates.

We write $*$ to refer to a revision function from $K \times \mathcal{L}$ to $\mathcal{K}$. Hence $K * p$ denotes the belief set that is the outcome of the revision of K by the sentence p. Again, the result of the change must be a belief set:

$K * p = \mathrm{Cn}(K * p)$ (closure)

According to the principle of success, the new sentence must be incorporated into the outcome of the revision.

$p \in K * p$ (success)

The revised belief set contains nothing that does not follow from the original belief set in combination with the new belief. This is guaranteed by the following postulate:

$K * p \subseteq K + p$ (inclusion)

Note that if $\neg p \in K$, then $K + p$ is the inconsistent belief set. In the case when the new belief does not contradict any of the sentences in K, there is no reason to remove any of them:

If $K \nvdash \neg p$ then $K + p \subseteq K * p$. (vacuity)

According to the "consistency preservation" criterion, unless the new belief is itself inconsistent, the result of the revision must be consistent.

If $p \nvdash \bot$ then $K * p \nvdash \bot$. (consistency)

Thus, if p is consistent, then $K * p$ is consistent even if K is inconsistent. Like contraction, the revision operation should be independent of the syntactic representation of the sentences. In other words, logically equivalent sentences must yield the same result:

If $\vdash p \leftrightarrow q$ then $K * p = K * q$. (extensionality)

The above are the six basic AGM (Gärdenfors) postulates for revision. Let us now analyze revision of a belief set K by a conjunction $p \wedge q$. According to the AGM authors, if q does not contradict $K * p$, then revising K by $p \wedge q$ yields the same result as expanding $K * p$ by q [139]. This follows from the following postulates:

$K * (p \wedge q) \subseteq (K * p) + q$ (superexpansion)

If $K * p \nvdash \neg q$ then $(K * p) + q \subseteq K * (p \wedge q)$. (subexpansion)

Note that when $\neg q \in K * p$, then $(K * p) + q = K_{\perp}$. Therefore the condition $K * p \nvdash \neg q$ is not needed in *superexpansion*. *Superexpansion* and *subexpansion* are called the supplementary AGM (or Gärdenfors) postulates for revision. They are presented in terms of revision by a conjunction, but they are equivalent to a pair of postulates that refer to revision by a disjunction [133]. The first of these postulates says that if a sentence r is incorporated both in the revision of K by p and in the revision of K by q, then r is also incorporated in the revision of K by $p \vee q$:

$(K * p) \cap (K * q) \subseteq K * (p \vee q)$ (disjunctive overlap)

The second of the disjunctive postulates says that when $K * (p \vee q) \nvdash \neg p$, the sentences in K that remain in $K * (p \vee q)$ are also retained in $K * p$:

If $K * (p \vee q) \nvdash \neg p$, then $K * (p \vee q) \subseteq K * p$. (disjunctive inclusion)

As already indicated, there is a direct correspondence between the conjunctive and the disjunctive postulates for revision:

Observation 3.3. [133] Let K be a belief set and let $*$ be an operation for K that satisfies *closure, success, inclusion, vacuity, consistency* and *extensionality*. Then:

1. $*$ satisfies *disjunctive overlap* if and only if it satisfies *superexpansion*.
2. $*$ satisfies *disjunctive inclusion* if and only if it satisfies *subexpansion*.

Finally, in the presence of the basic postulates, the combination of *superexpansion* (*disjunctive overlap*) and *subexpansion* (*disjunctive inclusion*) is equivalent to the following postulate:

$$K * (p \vee q) = \begin{cases} K * p, \text{ or} \\ K * q, \text{ or} \\ K * p \cap K * q \end{cases} \quad \text{(disjunctive factoring)}$$

Observation 3.4. [133] Let K be a belief set and let $*$ be an operation for K that satisfies *closure, success, inclusion, vacuity, consistency* and *extensionality*. Then $*$ satisfies both *superexpansion* and *subexpansion* if and only if $*$ satisfies *disjunctive factoring*.

The intuition behind disjunctive factoring is that if we wish to revise by a disjunction and there is some preference between the disjuncts, then this revision is equivalent to revising by the preferred disjunct. In the case of indifference, revising by the disjunction returns the beliefs that are common to the outcomes of revising by each member of the disjunction.

When the language is finite, K can be represented by a single sentence φ (which represents the conjunction of all the sentences in K). For this case, Katsuno and Mendelzon [205] presented an alternative (and equivalent) axiomatic characterization of revision:

(R1) $\varphi * p \vdash p$
(R2) If $\varphi \wedge p \nvdash \bot$ then $\vdash \varphi * p \leftrightarrow \varphi \wedge p$.
(R3) If $p \nvdash \bot$ then $\varphi * p \nvdash \bot$.
(R4) If $\vdash \varphi_1 \leftrightarrow \varphi_2$ and $\vdash p_1 \leftrightarrow p_2$ then $\vdash \varphi_1 * p_1 \leftrightarrow \varphi_2 * p_2$.
(R5) $(\varphi * p) \wedge q \vdash \varphi * (p \wedge q)$
(R6) If $(\varphi * p) \wedge q \nvdash \bot$ then $\varphi * (p \wedge q) \vdash (\varphi * p) \wedge q$.

3.7 Relations Between Contraction and Revision

We have seen that contraction and revision are characterized by two different sets of postulates. These postulates are independent in the sense that the postulates of revision do not refer to contraction and vice versa. However, it is possible to define revision functions in terms of contraction functions, and the other way around. We can define revision in terms of contraction by means of the Levi identity:

$K * p = K - \neg p + p$ (Levi identity)

The idea behind the Levi identity is that before adding p we have to reduce K so that p can be consistently added. It is possible to add p consistently to a set if and only if that set does not imply $\neg p$. Thus, our reduction of K should remove $\neg p$. This we can do by contracting $\neg p$ from K.

Observation 3.5. Let K be a theory and $-$ an operation for K that satisfies the contraction postulates *closure, inclusion, success, vacuity* and *extensionality*. Let $*$ be defined from $-$ via the Levi identity. Then:

1. [4] $*$ satisfies the revision postulates *closure, success, inclusion, vacuity, consistency* and *extensionality*.
2. [4, 97] If $-$ also satisfies *conjunctive inclusion*, then $*$ satisfies *subexpansion*.
3. [4] If $-$ also satisfies *conjunctive overlap* and *recovery*, then $*$ satisfies *superexpansion*.
4. [97] If $-$ satisfies *conjunctive overlap* but not *recovery*, then $*$ does not in general satisfy *superexpansion*.

Note that *recovery* is not needed to obtain the basic revision postulates. This means that each withdrawal generates, via the Levi identity, a revision that satisfies the six basic AGM postulates. If $-_1$ and $-_2$ are two withdrawals that generate the same revision they are called *revision equivalent* [239].

To define contraction in terms of revision we use the following identity:

$K - p = (K * \neg p) \cap K$ (Harper identity)

Similar relationships between the postulates hold in this direction as well.

Observation 3.6. [4] Let K be a theory and $*$ an operation for K that satisfies the revision postulates *closure, success, inclusion, vacuity, consistency* and *extensionality*. Let $-$ be defined from $*$ via the Harper identity. Then:

1. $-$ satisfies the contraction postulates *closure, inclusion, success, vacuity, extensionality* and *recovery*.
2. If $*$ also satisfies *subexpansion*, then $-$ satisfies *conjunctive inclusion*.
3. If $*$ also satisfies *superexpansion*, then $-$ satisfies *conjunctive overlap*.

The Levi and Harper identities provide us with a nice one-to-one correspondence between operations of revision and contraction:

Definition 3.4. [239] Let K be a belief set. Then $\mathbb{R}$ and $\mathbb{C}$ are functions from and to sentential operations on K such that

(1) for every sentential operation $-$ for K, $\mathbb{R}(-)$ is the operation such that for all $p \in \mathcal{L}$: $K(\mathbb{R}(-))p = Cn((K - \neg p) \cup \{p\})$.

(2) for every sentential operation $*$ for K, $\mathbb{C}(*)$ is the operation such that for all $p \in \mathcal{L}$: $K(\mathbb{C}(*))p = (K * \neg p) \cap K$.

Observation 3.7. [239] Let K be a logically closed set and $-$ an operation for K that satisfies the contraction postulates *closure*, *inclusion*, *vacuity*, *extensionality*, and *recovery*. Then $\mathbb{C}(\mathbb{R}(-)) = -$.

Observation 3.8. [239] Let K be a logically closed set and $*$ an operation for K that satisfies the revision postulates *closure*, *success*, *inclusion*, and *extensionality*. Then $\mathbb{R}(\mathbb{C}(*)) = *$.

3.8 Partial Meet Operations

Already in the original AGM paper in 1985, the operations were presented both syntactically, as we saw in Sections 3.5-3.6, and through an explicit construction. In this section we will present the direct construction that was provided in the 1985 paper [4]. As in that paper, the main focus will be on contraction.

If we wish to apply the criterion of *informational economy* uncompromisingly, then the contracted belief set $K - p$ should be as large a subset of K as it can be without implying p. In order to express this more precisely, the following notation is useful:

Definition 3.5. [5] Let K be a belief set and p a sentence. The set $K \perp p$ ("*K remainder p*") is the set of sets such that $H \in K \perp p$ if and only if:

$$\begin{cases} H \subseteq K \\ H \nvdash p \\ \text{There is no set } H' \text{ such that } H \subset H' \subseteq K \text{ and } H' \nvdash p \end{cases}$$

Hence, $K \perp p$ is the set of maximal subsets of K that do not imply p. $K \perp p$ is called a *remainder set* and its elements are the *remainders of K by p*. There is a special remainder set $\mathcal{L} \perp\perp$ that consists of all the maximal consistent subsets of the language, i.e., possible worlds. We will use $\mathcal{L} \perp\perp$ as an alternative notation for $\mathfrak{W}$.

Remainder sets satisfy the following properties:

1. $K \perp p = \{K\}$ if and only if $K \nvdash p$.
2. $K \perp p = \varnothing$ if and only if $\vdash p$.
3. If $H \subseteq K$ and $H \nvdash p$, then there exists some $H' \in K \perp p$ such that $H \subseteq H'$ [5] [5].
4. If $p \in K$ and $\nvdash p$, then for all H in $K \perp p$, $H + \neg p$ is an element of $\mathcal{L} \perp\perp$ (i.e., a possible world) [6].

A first tentative approach to constructing an operation of contraction is to choose just one element from $K \perp p$ for each input sentence [6]:

$K - p \in K \perp p$ when $\nvdash p$, and otherwise $K \perp p = K$. (maxichoice contraction)

Though it seems to be intuitive, maxichoice contraction generates belief sets that are "too large", since they satisfy the following property [238]:

If $p \in K$, then for any $q \in \mathcal{L}$, either $p \vee q \in K - p$ or $p \vee \neg q \in K - p$. (saturability)

The following example shows the implausibility of this property:

Example 3.1. I believe that "it is four o'clock" (p). Then I discover that my watch has stopped. After that I must contract my belief p (but not revise by $\neg p$). According to *saturability* I must retain either "it is four o'clock or there is a life after death" ($p \vee q$) or "it is four o'clock or there is no life after death" ($p \vee \neg q$), but I have no reason to make this choice.

As was noted by Makinson [238, p. 357], neither $p \vee q$ nor $p \vee \neg q$ should be retained in general in the process of eliminating p from K, unless there is "some reason" in K for its continued presence.

Maxichoice contraction satisfies the following postulate [133, 163]:

If $q \in K$ and $q \notin K - p$ then $\nvdash p$ and $q \to p \in K - p$. (fullness)

In the presence of *closure* and *success*, *fullness* implies *recovery*. Using *fullness* we can obtain an axiomatic characterization of maxichoice contraction:

[5] This property depends on compactness and the axiom of choice.

Observation 3.9. [133] Let K be a belief set. An operation $-$ on K is a maxichoice contraction if and only if $-$ satisfies *closure, success, inclusion, vacuity, extensionality*, and *fullness*.

At the other extreme, we can consider an operation that returns only the sentences that are common to all the elements of $K \perp p$ [6]:

$$K - p = \bigcap(K \perp p) \text{ when } \nvdash p, \text{ and otherwise } K - p = K. \text{ (full meet contraction)}$$

Contrary to maxichoice contraction, full meet contraction generates belief sets that are "too small", since they satisfy the following property [97]:

$$\text{If } p \in K \text{ then for any } q \in \mathcal{L},\ p \vee q \in K - p \text{ if and only if } \vdash p \vee q. \text{ (devastation)}$$

Full meet contraction is characterized as follows:

Observation 3.10. [6] If $-$ is the full meet contraction on K and $p \in K$, then $K - p = Cn(\{\neg p\}) \cap K$.

In terms of postulates, full meet contraction can be characterized as follows:

Observation 3.11. [4] Let K be a belief set. An operation $-$ is the full meet contraction on K if and only if $-$ satisfies *closure, success, inclusion, vacuity, extensionality, recovery*, and

$$K - (p \wedge q) = (K - p) \cap (K - q). \text{ (meet identity)}$$

Although full meet contraction is not an appropriate contraction, it provides the *lower bound* for the *recovery* postulate. We can formalize this concept in the following way:[6]

Observation 3.12. [238] Let K be a belief set, $\sim$ the operation of full meet contraction for K and $-$ an operation for K. Then $-$ satisfies *recovery* if and only if $K \sim p \subseteq K - p$ for all p.

A third approach is to generate the contraction outcome from the intersection of only some of the elements of $K \perp p$. To do this we need to define a selection function for $K \perp p$.

Definition 3.6. [4] Let K be a belief set. A *selection function* for K is a function γ such that for all sentences p:

1. If $K \perp p$ is non-empty, then $\gamma(K \perp p)$ is a non-empty subset of $K \perp p$.
2. If $K \perp p$ is empty, then $\gamma(K \perp p) = K$.

[6] For a detailed study of full meet contraction see [166] and [173].

We can further specify properties of the selection function to ensure that the "best" elements of $K \perp p$ are selected. For this purpose, we need to introduce a preference relation on $K \perp p$:

Definition 3.7. [4] A selection function γ for a belief set K is *relational* if and only if there is a relation $\sqsubseteq$ on the remainders of K[7] such that for all sentences p, if $K \perp p$ is non-empty, then:

$$\gamma(K \perp p) = \{B \in K \perp p \mid B' \sqsubseteq B \text{ for all } B' \in K \perp p\}$$

γ is *transitively relational* if and only if $\sqsubseteq$ is a transitive relation.

Partial meet contraction is defined in terms of the selection function γ:

Definition 3.8. [4] Let K be a belief set and γ a selection function for K. An operation $-$ on K is a partial meet contraction if and only if there is a selection function γ for K such that for all sentences p:

$$K - p = \bigcap \gamma(K \perp p)$$

Furthermore, $-$ is (transitively) relational if and only if it can be generated from a (transitively) relational selection function.

Maxichoice contraction is the special case of partial meet contraction when for all sentences p, $\gamma(K \perp p)$ has exactly one element. Full meet contraction is the special case when $\gamma(K \perp p) = K \perp p$ whenever $K \perp p$ is non-empty.

One of the major achievements of AGM theory is the characterization of partial meet contraction, and its transitively relational variant, in terms of postulates:

Theorem 3.2. [4] Let K be a belief set. An operation $-$ on K is a partial meet contraction if and only if $-$ satisfies *closure, success, inclusion, vacuity, recovery*, and *extensionality*. If $-$ is a relational partial meet contraction then it satisfies *conjunctive overlap*. Furthermore, $-$ is a transitively relational partial meet contraction if and only if it satisfies both *conjunctive overlap* and *conjunctive inclusion*.

As we saw in Section 3.7, we can define a corresponding operation of revision by means of the Levi identity:

Definition 3.9. [4] Let K be a belief set. Let $*$ and $-$ be operations on K such that for all sentences p:

$$K * p = (K - \neg p) + p$$

Then:

1. $*$ is a *maxichoice revision* if and only if $-$ is a *maxichoice contraction*.
2. $*$ is a *full meet revision* if and only if $-$ is a *full meet contraction*.

[7] X is a remainder of K if and only if there exists some p with $X \in K \perp p$.

3. $*$ is a *partial meet revision* if and only if $-$ is a *partial meet contraction*.
4. $*$ is a *(transitively) relational partial meet revision* if and only if $-$ is a *(transitively) relational partial meet contraction*.

We have observed that *maxichoice contraction* and *full meet contraction* are implausible operations, but useful as upper and lower bounds of *partial meet contraction*. This implausibility is even more evident in revision, as can be inferred from the following observations:

Observation 3.13. [6] Let $*$ be a *maxichoice revision* for a belief set K. If $\neg p \in K$, then $K * p \in \mathcal{L} \perp \perp$.

Observation 3.14. [6] Let $*$ be a *full meet revision* for a belief set K. If $\neg p \in K$, then $K * p = Cn(p)$.

As we did for contraction, we can characterise *partial meet revision* in terms of postulates:

Theorem 3.3. [4] Let K be a belief set. An operation $*$ on K is a partial meet revision if and only if $*$ satisfies *closure, success, inclusion, vacuity, consistency*, and *extensionality*. If $*$ is a relational partial meet revision then it satisfies *superexpansion*. Furthermore $*$ is a transitively relational partial meet revision if and only if it satisfies both *superexpansion* and *subexpansion*.

Chapter 4
Equivalent Characterizations

One of the most remarkable features of the AGM model is its capability of being expressed in several different, seemingly quite dissimilar ways. In this chapter we present five different ways to characterise AGM operations: *possible world models*, *epistemic entrenchment*, *specified meet contraction*, *kernel contraction* and *safe contraction*.

4.1 Possible World Models

In Section 3.3 we introduced possible worlds in the logical sense, viz. inclusion-maximal consistent subsets of the language. We noted that there is a one-to-one correspondence between belief sets and sets of possible worlds. We also introduced the notation $\|K\|$ for the set of all possible worlds that contain the belief set K and similarly $\|p\|$ for the set of all possible worlds that contain the sentence p. The following observation summarizes some useful properties of possible worlds.

Observation 4.1. [143, 163] Let K and H be logically closed sets and p and q sentences. Then the following properties hold:

$H \subseteq K$ if and only if $\|K\| \subseteq \|H\|$
If $\mu \in \mathcal{L} \perp\!\!\perp$, then $\mu \in \|p\|$ if and only if $\mu \notin \|\neg p\|$.
$\|Cn(K \cup H)\| = \|K\| \cap \|H\|$
$\|K\| \cup \|H\| \subseteq \|K \cap H\|$.
$\|p\| \subseteq \|q\|$ if and only if $\vdash p \rightarrow q$.
$\|p \wedge q\| = \|p\| \cap \|q\|$
$\|p \vee q\| = \|p\| \cup \|q\|$.

One of the major advantages of possible world models is that they can be represented graphically in a very intuitive way. Usually, a rectangular space such as the square in Figure 4.1 represents the set of all possible worlds. Each possible world is represented by a point in the rectangle, and each set of possible worlds (proposition) is thought of as a subarea of the rectangle. Hence, in Figure 4.1, the central

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_4

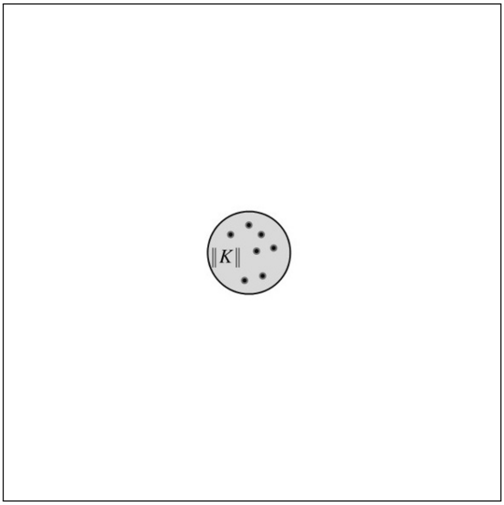

Fig. 4.1: A possible worlds model.

circle corresponds to the belief set K. It is important to note that in these diagrams, a larger area represents a smaller belief set (since a smaller belief set is compatible with a larger number of possible states of the world). The smallest belief set, $Cn(\varnothing)$, is represented by the whole square since it is a subset of all possible worlds. The largest belief set, $Cn(\bot)$, is represented by the empty set since it is not a subset of any possible world.

Expansion is easily represented in a possible worlds model. $\|K+p\|$ should be the set of worlds that are both K-worlds and p-worlds; thus:

$$\|K+p\| = \|K\| \cap \|p\|$$

This is illustrated in Figure 4.2.

The contraction of K by p takes the form of an addition of some $\neg p$-worlds to the set $\|K\|$. We will use the concept of a *propositional selection function* to present the possible worlds semantics for partial meet contractions.

Definition 4.1. [163] Let $\mathcal{M}$ be a proposition. A *propositional selection function* for $\mathcal{M}$ is a function f such that for all sentences p: (1) $f(\|p\|) \subseteq \|p\|$, (2) If $\|p\| \neq \varnothing$ then $f(\|p\|) \neq \varnothing$ and (3) If $\mathcal{M} \cap \|p\| \neq \varnothing$, then $f(\|p\|) = \mathcal{M} \cap \|p\|$.

We can think of $\mathcal{M}$ as the set of K-worlds for some belief set K. Then this definition says that f selects some p-worlds if there are any, and that if there are some p-worlds that are also K-worlds, then it selects exactly those.

We are now in a position to present the following observation, showing how partial meet contraction can be constructed in terms of possible worlds. (See Figure 4.3.)

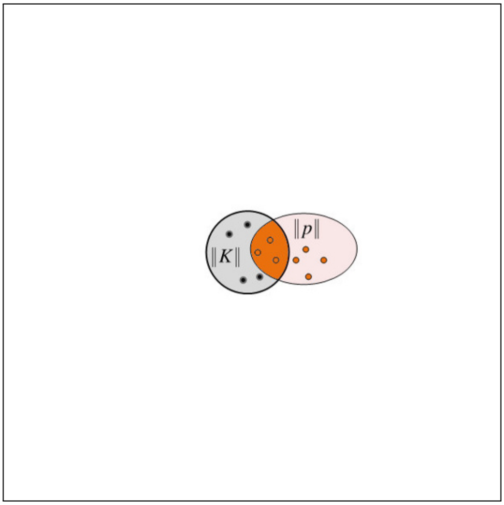

Fig. 4.2: Expansion. The dark orange area is the outcome of the operation $\|K+p\|$.

Observation 4.2. [143, 163] Let K be a belief set. An operation $-$ on K is a *partial meet contraction* if and only if there is a propositional selection function f for $\|K\|$ such that for all sentences p: $K - p = \bigcap(\|K\| \cup f(\|\neg p\|))$.

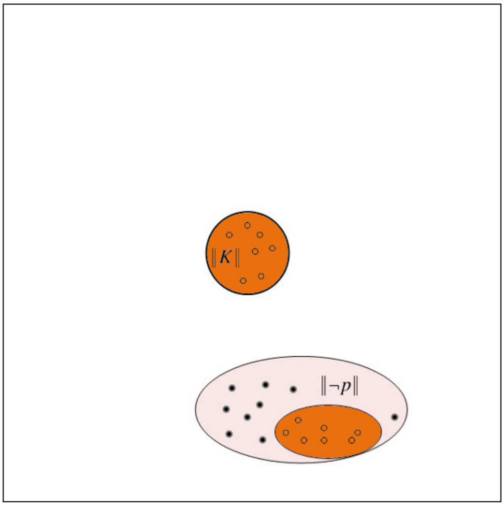

Fig. 4.3: Partial meet contraction in a case with $p \in K$ and $\nvdash p$. The outcome of the operation $\|K-p\|$ is represented by the union of the dark orange areas.

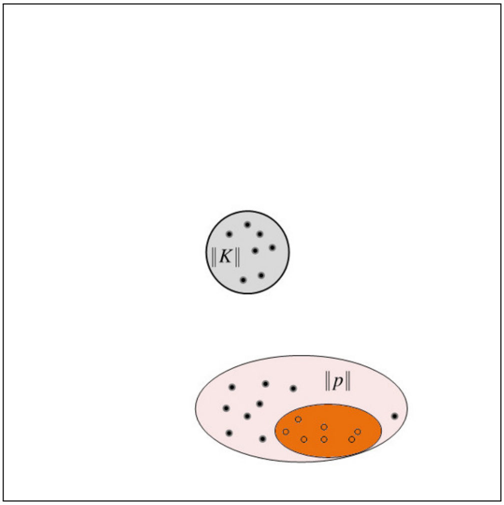

Fig. 4.4: Partial meet revision. The dark orange area represents the outcome of the operation $\|K * p\|$.

Due to the Levi identity and the definitions of f and $+$ we can define $K * p = (K - \neg p) + p = \bigcap((\|K\| \cup f(\|p\|)) \cap \|p\|) = \bigcap(f(\|p\|))$. (See Figure 4.4.)

To capture the supplementary postulates either in contraction or in revision, and consequently transitively relational partial meet contraction (revision), we need tools that are much more sophisticated. There are two closely related ways to do this: by Grove's sphere systems or by means of a total preorder on worlds.

Adam Grove [143] defined a sphere system centred around $\|K\|$ as a collection $\mathbb{S}$ of subsets of $\mathcal{L}\perp\perp$ that are ordered by inclusion. Figuratively, the distance from a possible world to the centre of the system reflects its plausibility relative to the theory K. Formally:

Definition 4.2. [143] Let $\mathcal{X}$ be a subset of $\mathcal{L}\perp\perp$. A *system of spheres* centred on $\mathcal{X}$ is a collection $\mathbb{S}$ of subsets of $\mathcal{L}\perp\perp$, i.e., $\mathbb{S} \subseteq \mathcal{P}(\mathcal{L}\perp\perp)$, that satisfies the following conditions:

($\mathbb{S}$1) $\mathbb{S}$ is totally ordered with respect to set inclusion; that is, if $\mathcal{U}, \mathcal{V} \in \mathbb{S}$, then $\mathcal{U} \subseteq \mathcal{V}$ or $\mathcal{V} \subseteq \mathcal{U}$.
($\mathbb{S}$2) $\mathcal{X} \in \mathbb{S}$, and if $\mathcal{U} \in \mathbb{S}$ then $\mathcal{X} \subseteq \mathcal{U}$.
($\mathbb{S}$3) $\mathcal{L}\perp\perp \in \mathbb{S}$ (and so it is the largest element of $\mathbb{S}$).
($\mathbb{S}$4) For every $p \in \mathcal{L}$, if there is some element of $\mathbb{S}$ intersecting with $\|p\|$, then there is also a smallest element of $\mathbb{S}$ intersecting with $\|p\|$. The smallest sphere in $\mathbb{S}$ intersecting with $\|p\|$ is denoted by $\mathbb{S}_p$.

The elements of $\mathbb{S}$ are called *spheres*. See Figure 4.5 for a graphical illustration of a sphere system.

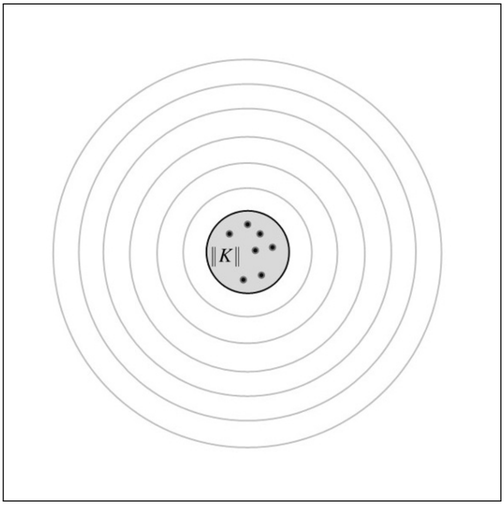

Fig. 4.5: A system of spheres centred around K.

Definition 4.3. [143] A propositional selection function f for a proposition $\mathcal{X}$ is *sphere-based* if and only if there is a system of spheres $\mathbb{S}$ such that for all p: If $\|p\| \neq \emptyset$, then $f(\|p\|) = \mathbb{S}_p \cap \|p\|$.

Observation 4.3. [143] Let K be a belief set.

1. An operation $-$ on K is a *transitively relational partial meet contraction* if and only if there is a sphere-based propositional selection function f for $\|K\|$ such that for all sentences p: $K - p = \bigcap(\|K\| \cup f(\|\neg p\|))$.
2. An operation $*$ on K is a *transitively relational partial meet revision* if and only if there is a sphere-based propositional selection function f for $\|K\|$ such that for all sentences p: $K * p = \bigcap(f(\|p\|))$.

These operations are illustrated in Figures 4.6 and 4.7.

Alternatively, the semantics of the AGM model can be characterized by a total preorder between possible worlds, i.e., a reflexive, transitive and total relation on $\mathcal{L}\perp\perp$.

Definition 4.4. [205] A total preorder $\leq_K$ on possible worlds, with the strict part $<_K$ and the symmetric part $\simeq_K$, is a *global faithful assignment* if and only if the following conditions hold:

1. If $K \subseteq \omega$ and $K \subseteq \omega'$, then $\omega \simeq_K \omega'$
2. If $K \subseteq \omega$ and $K \nsubseteq \omega'$, then $\omega <_K \omega'$

The notion of a global faithful assignment allows us to define partial meet operations:

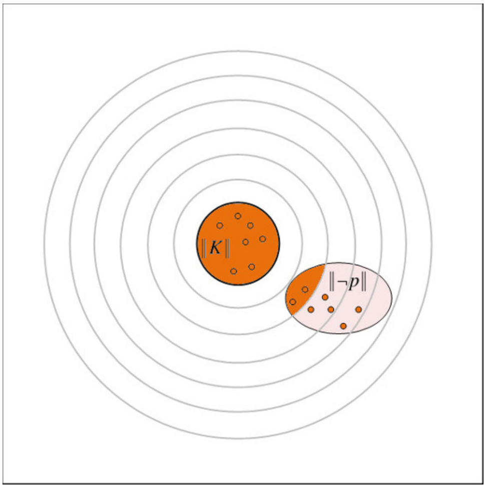

Fig. 4.6: Transitively relational partial meet contraction. The outcome of the operation $\|K-p\|$ is represented by the union of the dark orange areas.

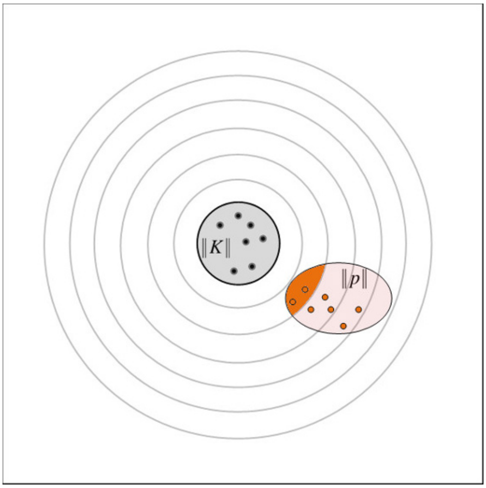

Fig. 4.7: Transitively relational partial meet revision. The dark orange area represents the outcome of the operation $\|K * p\|$.

Observation 4.4. [205] Let K be a belief set.

1. An operation $-$ on K is a *transitively relational partial meet contraction* if and only if there is a global faithful assignment $\leq_K$ for K such that $K-p = \bigcap(\|K\| \cup min(\|\neg p\|, \leq_K))$.
2. An operation $*$ on K is a *transitively relational partial meet revision* if and only if there is a global faithful assignment $\leq_K$ for K such that:

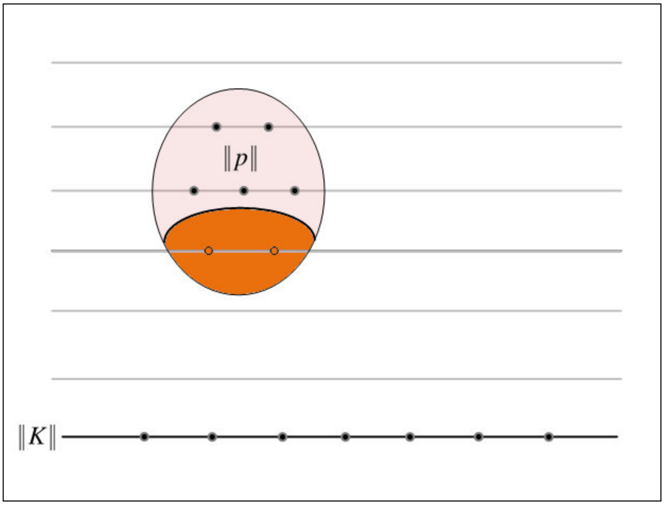

Fig. 4.8: Transitively relational partial meet revision in the model of global faithful assignment. The dark orange area represents the outcome of the operation $\|K * p\|$.

$$K * p = \bigcap min(\|p\|, \leq_K)$$

Revision by global faithful assignment is illustrated on Figure 4.8.[1]

4.2 Epistemic Entrenchment

> "Even if all sentences in a knowledge set are accepted or considered as facts, this does not mean that all sentences are of equal value for planning or problem solving purposes. Certain pieces of knowledge and belief about the world are more important than others when planning future actions, conducting scientific investigations or reasoning in general. We will say that some sentences in a knowledge system have a higher degree of *epistemic entrenchment* than others. The degree of entrenchment will, intuitively, have a bearing on what is abandoned from a knowledge set and what is retained, when a contraction or revision is carried out." [138]

Epistemic entrenchment was introduced by Peter Gärdenfors in [131], [138] and [133]. It is a binary relation $\leq$ on the sentences in the belief set K such that in contraction, giving up beliefs with lower entrenchment is preferred to giving up those with higher entrenchment. $p \leq q$ denotes that p is at most as sentrenched as q and $p < q$ that p is less entrenched than q. From a formal point of view, $<$ is the strict part of $\leq$, i.e., $p < q$ holds if and only if $p \leq q$ and $q \nleq p$.

[1] This graphical notation for global faithful assignment is due to Sébastien Konieczny and Ramón Pino Pérez.

Gärdenfors proposed a set of five postulates for epistemic entrenchment. The first postulate simply states that the order is transitive.

If $p \leq q$ and $q \leq r$, then $p \leq r$. (transitivity)

If q follows logically from p, then p has to be given up whenever q is given up. We can therefore assume that p is at most as entrenched as q in this case.

If $p \vdash q$, then $p \leq q$. (dominance)

Removing $p \wedge q$ necessarily implies removing either p or q. It is therefore natural to assume that either p or q is at most as entrenched as $p \wedge q$.

$p \leq (p \wedge q)$ or $q \leq (p \wedge q)$ (conjunctiveness)

The minimality postulate states that non-beliefs are all minimally entrenched.

If $K \neq K_{\perp}$, then $p \notin K$ if and only if $p \leq q$ for all q. (minimality)

The maximality postulate, on the other hand, states that the maximally entrenched beliefs are (exactly) the logical truths.

If $q \leq p$ for all q, then $\vdash p$. (maximality)

A relation satisfying transitivity, dominance, conjunctiveness, minimality and maximality is a *standard entrenchment ordering*.

Observation 4.5. [109, 138, 163, 296] Let $\leq$ (with the strict part $<$) be a standard entrenchment ordering. Then it satisfies:

4.1 *$p \leq p$ (reflexivity)*
4.2 *$p \leq q$ or $q \leq p$ (connectedness)*
4.3 *If $q \wedge r \leq p$, then $q \leq p$ or $r \leq p$.*
4.4 *$p < q$ if and only if $p \wedge q < q$.*
4.5 *If $r \leq p$ and $r \leq q$, then $r \leq p \wedge q$.*
4.6 *If $p \leq q$, then $p \leq p \wedge q$.*
4.7 *If $p < q$ and $q < r$, then $p < r$.*
4.8 *If $p \leq q$ and $q < r$, then $p < r$.*
4.9 *If $q < r$ and $q < p$, then $q < p \wedge r$.*
4.10 *If $p < q$, then $p < p \vee q$.*
4.11 *If $p < q$, then $p \wedge r < q$.*
4.12 *If $p \leq q$, then $p \wedge r \leq q$.*
4.13 *If $q \wedge r < p$, then $q < p$ or $r < p$.*
4.14 *If $p \leq q$, then $p \wedge r \leq q \wedge r$.*
4.15 *If $\nvdash p$ and $\vdash q$, then $p < q$.*
4.16 *$p \notin K$ if and only if $p < q$ for all $q \in K$.*
4.17 *If $p < q$, then $p \leq p \wedge q$.*
4.18 *$p \vee q < p \vee \neg q$ if and only if $p < p \vee \neg q$.*
4.19 *If $\vdash p \leftrightarrow p'$ and $\vdash q \leftrightarrow q'$, then: $p \leq q$ if and only if $p' \leq q'$. (intersubstitutivity)*
4.20 *If $K \neq K_{\perp}$, then: $p \in K$ if and only if $\perp < p$.*
4.21 *If $\top \leq p$, then $\vdash p$. (maximality 2)*
4.22 *If $p \notin K$, then $p \leq q$. (K-minimality)*

4.23 *If* $p \in K$ *and* $p \leq q$*, then* $q \in K$. *(K-representation)*

4.24 $p \wedge q \leq r$ *if and only if* $p \leq q \wedge r$ *or* $q \leq p \wedge r$. *(choice)*

Gärdenfors proposed the following connections between orders of epistemic entrenchment and operations of contraction:

$(G_{\leq})$ $\quad q \in K - p$ if and only if $q \in K$ and either $\vdash p$ or $p < (p \vee q)$.

$(C_{\leq})$ $\quad p \leq q$ if and only if $p \notin K - (p \wedge q)$ or $\vdash (p \wedge q)$.

Theorem 4.25. [133, 138] Let $\leq$ be a standard entrenchment ordering on a consistent belief set K. Furthermore let $-$ be the entrenchment-based contraction on K defined from $\leq$ by condition $G_{\leq}$. Then $-$ is a transitively relational partial meet contraction, and $C_{\leq}$ also holds.

Theorem 4.26. [133, 138] Let $-$ be a transitively relational partial meet contraction on a consistent belief set K. Furthermore let $\leq$ be the relation defined from $-$ by condition $C_{\leq}$. Then $\leq$ satisfies the standard entrenchment postulates and $G_{\leq}$ also holds.

The crucial clause of $G_{\leq}$ is $p < (p \vee q)$. This clause can be justified with reference to the *recovery* postulate [138]. If $-$ satisfies *recovery* and *closure*, then it holds for all $q \in K$ that $p \rightarrow q \in K - p$. Since $p \rightarrow q$ and $p \vee q$ together imply q, it follows that if $p \vee q \in K - p$, then $q \in K - p$. Furthermore, if $q \in K - p$, then $p \vee q \in K - p$. Thus, $q \in K - p$ if and only if $p \vee q \in K - p$.

We can see from $C_{\leq}$ (excluding the limiting case) that:

$p \vee q \notin K - ((p \vee q) \wedge p)$ if and only if $p \vee q \leq p$.

Using *extensionality*, we obtain:

$p \vee q \notin K - p$ if and only if $p \vee q \leq p$.

Negating both sides of the equivalence, using the connectedness of $\leq$, we obtain:

$p \vee q \in K - p$ if and only if $p < p \vee q$.

We showed above that given *recovery* and *closure*, $p \vee q \in K - p$ is equivalent to $q \in K - p$, thus:

$q \in K - p$ if and only if $p < p \vee q$,

as desired. However, it does not seem possible to intuitively justify $G_{\leq}$ unless one accepts the *recovery* postulate [138, pp. 89–90]. In fact, Gärdenfors admitted: “The comparison is somewhat counterintuitive” [136].

However, although *recovery* is needed to prove that the relation $\leq$ obtained from $-$ through $C_{\leq}$ satisfies $G_{\leq}$, it is not needed to show that $\leq$ is a standard entrenchment relation. Consequently, the role of *recovery* in epistemic entrenchment is located in the definition of the rule of contraction. Rott [288] proposed an alternative construction of contraction from $\leq$, later called *severe withdrawal* [302]:

$$K - p = \{q \in K \mid p < q\}$$

Severe withdrawal has been axiomatized [302, 104]. It has also been shown to satisfy the following implausible postulate [163]:

If $\nvdash p$ and $\nvdash q$, then either $p \notin K \div q$ or $q \notin K \div p$. (expulsiveness).

However, Lindström and Rabinowicz have proposed that it can be used as a lower limit for contraction. Transitively relational partial meet contraction would be the upper limit, and a reasonable operation of contraction should be somewhere between these two extremes [231, pp. 115]. This has been called Lindström's and Rabinowicz's interpolation thesis [292].

Fermé [96] and Rott [295] showed that a transitively relational partial meet contraction satisfies *fullness*[2] if and only if it also satisfies:

If $p \in K$ and $\nvdash p$, then $p < p \vee q$ or $p < p \vee \neg q$.

Entrenchment-based revision can be defined from entrenchment-based contraction via the Levi identity. However, it is also possible to define entrenchment-based revision directly from an entrenchment ordering, by means of the following equivalence [231, 287, 97]:

$(*_{EBR})$ $\quad q \in K * p$ if and only if either $(p \rightarrow \neg q) < (p \rightarrow q)$ or $p \vdash \bot$.

Conversely, an entrenchment relation can be obtained from an operation of revision as follows:

$(C_{\leq *})$ $\quad p \leq q$ if and only if: If $p \in K * \neg(p \wedge q)$ then $q \in K * \neg(p \wedge q)$.

In order to clarify $(C_{\leq *})$, assume that $p \in K * \neg(p \wedge q)$ and $q \in K * \neg(p \wedge q)$; then $p \wedge q \in K * \neg(p \wedge q)$, which implies $\vdash p \wedge q$. This means that if we are not in the limiting case, then $(C_{\leq *})$ implies $p \notin K * \neg(p \wedge q)$. Due to the Levi identity, this is equivalent with $p \notin (K - (p \wedge q)) + \neg(p \wedge q)$, or equivalently, $p \notin K - (p \wedge q)$.

Theorem 4.27. [97] Let $\leq$ be a standard entrenchment ordering on a consistent belief set K. Furthermore let $*$ be the revision on K defined by condition $*_{EBR}$ from $\leq$. Then $*$ is a transitively relational partial meet revision, and $C_{\leq *}$ also holds.

[2] In the presence of the AGM contraction postulates fullness characterizes maxichoice contraction (see Section 3.8).

Theorem 4.28. [97] Let $*$ be a transitively relational partial meet revision on a consistent belief set K. Furthermore let $\leq$ be the relation defined from $*$ by condition $C_{\leq *}$. Then $\leq$ satisfies the standard entrenchment postulates and $*_{EBR}$ also holds.

Rott [296] provided a characterization of partial meet revision in terms of a weakened version of epistemic entrenchment:

Theorem 4.29. [296] Let $\leq$ be an ordering on a consistent belief set K and let $\leq$ satisfy reflexivity, intersubstitutivity, choice, maximality 2, K-minimality and K-representation.[3] Furthermore let $*$ be an entrenchment-based revision on K defined by condition $*_{EBR}$ from $\leq$. Then $*$ is a partial meet revision.

Theorem 4.30. [296] Let $*$ be be a partial meet revision on a consistent belief set K. Furthermore let $\leq$ be the relation defined from $*$ by condition $C_{\leq *}$. Then $\leq$ satisfies the entrenchment postulates reflexivity, intersubstitutivity, choice, maximality 2, K-minimality and K-representation.

The belief change literature contains several approaches related with epistemic entrenchment. In *quantitative formalisms*, Dubois and Prade [83, Sec. 2.5] remarked that a possibility or necessity measure provides a preorder on sentences. Spohn [315] defined a model of ordinal conditional functions where the order of sentences represents the degree of implausibility of their negations. In *qualitative formalisms*, Rott [288] introduced a preference relation that provides a canonical representation of an epistemic entrenchment and a belief base. Williams has introduced operations based on ensconcement relations on belief bases (see Section 6.2). She has shown that Rott's proposal is equivalent to the ensconcement approach. In the framework of descriptor revision, epistemic entrenchment has been generalized to relations of proximity between success conditions (see Chapter 14). For a discussion on quantitative vs. qualitative formalisms see [107, Sec. 1].

4.3 Specified Meet Contraction

As we saw in Chapter 3, partial meet contraction was introduced as a modification of a simpler (and purely logical) operation, namely full meet contraction:

$$K \sim p = \bigcap (K \perp p)$$

Full meet contraction was first described by Alchourrón and Makinson [6]. It was soon realized that it removes too much from the belief set. Partial meet contraction was constructed to solve this problem. As we have already seen, it operates by selecting a subset of the elements of $K \perp p$ for intersection, and its defining formula is:

[3] These properties are defined in Observation 4.5.

$$K - p = \bigcap \gamma(K \perp p)$$

Specified meet contraction, introduced in [169], is an alternative modification of full meet contraction to make it less restrictive. In specified meet contraction, the selection mechanism does not operate on the remainder set but on the sentence to be contracted. The underlying intuition is that in order to contract by a complex sentence, we often only remove parts of its contents:

> I realized that $p \wedge q$ cannot be true, so I had to give up either my belief in p or my belief in q, or perhaps both. After some deliberation I chose to give up only p.

Specified meet contraction employs a *sentential selector* f that takes us from the sentence that is the contraction input to the sentence that will be eliminated by full meet contraction. The formal definition is as follows:

Definition 4.5. [169] An operation $-$ on K is an operation of *specified meet contraction* if and only if there is a sentential selector f such that for all p, $K - p = K \sim f(p)$.

The two constructions are compared in Figure 4.9. The following theorem identifies, in the form of an axiomatic characterization, the operations that can be constructed as specified meet contractions:

Theorem 4.31. [168] Let $-$ be a sentential operation on a finite-based belief set K. The following two conditions are equivalent:

(1) $-$ satisfies:
$K - p = \mathrm{Cn}(K - p)$ *(closure)*,
$K - p \subseteq K$ *(inclusion)*, and
$K - p$ is finite-based. *(finite-based outcome)*

(2) There is a sentential selector f such that for all p: $f(p) \in K$ and $K - p = K \sim f(p)$.

It can be seen from this theorem that specified meet contraction covers many types of operations that are not partial meet contractions. In particular, neither the *success* nor the *recovery* postulate needs to be satisfied for an operation to be a specified meet contraction. But on the other hand, one of the required three postulates, namely *finite-based outcome*, is not satisfied in general by partial meet contraction. Therefore, only some partial meet contractions are also specified meet contractions.

Observation 4.6. [177] Let K be a finite-based belief set and $-$ a partial meet contraction on K. Then $-$ satisfies *finite-based outcome* if and only if $-$ is (also) a specified meet contraction.

It has often been argued that *finite-based outcome* is a desirable property for an operation of contraction. For instance, all three AGM authors have endorsed it [6, pp. 21–22], [138, p. 90], [239, p. 384]. This gives us reason to investigate those partial

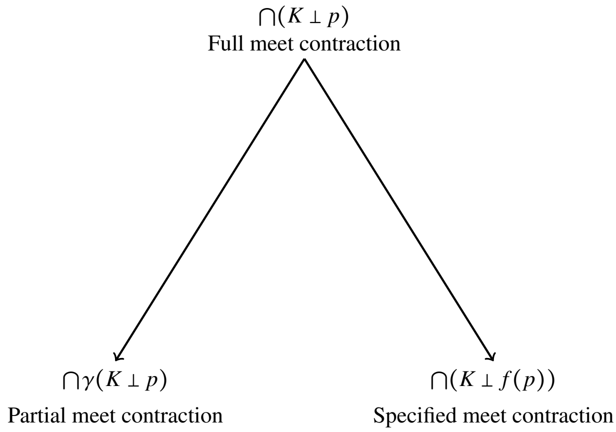

Fig. 4.9: A comparison between three forms of contraction: full meet contraction, partial meet contraction, and specified meet contraction.

meet contractions that are also specified meet contractions. They can of course be specified in terms of the four basic AGM postulates for contraction that are not mentioned in the above axiomatic characterization of specified meet contraction, namely *extensionality, vacuity, success*, and *recovery*. Alternatively, they can be characterized in terms of properties of the sentential selector. This is fairly easily done since each of these postulates corresponds exactly to a property of the sentential selector f. For instance, a specified meet contraction $-$ satisfies *recovery* if and only if its sentential selector f satisfies the property:

$$\text{If } p \in K, \text{ then } \vdash p \rightarrow f(p).$$

The properties corresponding to the remaining three basic AGM postulates can be found in Table 4.1. The reader is referred to [168, 169, 171, 172, 174] for additional formal results on specified meet contraction, including further properties of the sentential selector f and extensions to multiple and iterated variants of specified meet contraction.

Properties of sentential selectors	Contraction postulates
If $\vdash p \leftrightarrow q$ then $\vdash f(p) \leftrightarrow f(q)$	*extensionality*
$f(p) \in K$ if and only if $p \in K$	*vacuity*
If $\vdash p$ then $\vdash f(p)$	*failure*
If $p \in K \setminus \mathrm{Cn}(\varnothing)$, then $\nvdash p \vee f(p)$	*success*
If $p \in K$, then $\vdash p \rightarrow f(p)$	*recovery*

Table 4.1: Some one-to-one correspondences between the properties of a sentential selector and the properties of the specified meet contraction that it gives rise to.

4.4 Kernel Contraction

Whereas partial meet contraction is based on a selection of what to retain, kernel contraction [154] and safe contraction [7] are based on a selection of sentences to remove. Safe contraction was proposed by Alchourrón and Makinson in 1985, the same year that the seminal AGM paper was published. Kernel contraction was introduced nine years later, but for expository reasons we begin with the latter operation.

We define the *kernel* set for a sentence p and a belief set K as follows:

Definition 4.6. [154] Let K be a belief set and p a sentence. Then $K \perp\!\!\!\perp p$ is the set such that $A \in K \perp\!\!\!\perp p$ if and only if:

$$\begin{cases} A \subseteq K \\ A \vdash p \\ \text{If } B \subset A \text{ then } B \nvdash p. \end{cases}$$

$K \perp\!\!\!\perp p$ is called the *kernel set of K with respect to p* and its elements are the *p-kernels* of K.

Basically, at least one element of each p-kernel of K must be removed in the contraction process, since otherwise the sentence p would still be implied. On the other hand, due to the minimality criterion, only sentences that are included in one or more elements of the kernel set may have to be discarded. The remaining problem is how to choose the sentences to remove. The most general case, i.e., without additional criteria for the selection, is an incision function defined as follows:

Definition 4.7. [154] An *incision function* σ for K is a function such that for all sentences p:

$$\begin{cases} \sigma(K \perp\!\!\!\perp p) \subseteq \bigcup(K \perp\!\!\!\perp p) \\ \text{If } \varnothing \neq A \in K \perp\!\!\!\perp p, \text{ then } A \cap \sigma(K \perp\!\!\!\perp p) \neq \varnothing. \end{cases}$$

Thus, incision functions cut into each p-kernel, removing at least one sentence. A subset of K implies p if and only if it contains at least one p-kernel. The set obtained from K by removing all sentences selected by the incision function does not contain any p-kernel. This is why incision functions can be used to derive contraction operations.

Definition 4.8. [154] Given a belief set K, a sentence p and an incision function σ for K, the *kernel contraction* of K by p, denoted by $K -_\sigma p$, is defined as:

$$K -_\sigma p = K \setminus \sigma(K \perp\!\!\!\perp p)$$

However, the contraction obtained in this definition can yield an outcome $K -_\sigma p$ that is not closed under logical consequence. In order to ensure that the *closure* postulate is satisfied we have to impose a further condition on the incision function:

Definition 4.9. [154] An incision function σ for a belief set K is *smooth* if and only if it holds for all subsets A of K that: if $A \vdash q$ and $q \in \sigma(K \perp\!\!\!\perp p)$, then $A \cap \sigma(K \perp\!\!\!\perp p) \neq \emptyset$. A kernel contraction is smooth if and only if it is based on a smooth incision function.

Smoothness is what is required for a kernel contraction on a belief set to satisfy *closure*. Importantly, the smooth kernel contractions on belief sets coincide with the partial meet contractions [154].[4] This connection can be further specified in terms of the relation between selection and incision functions:

Observation 4.7. [87] Let γ a selection function for a set K. The function σ_γ defined as follows:

$$\sigma_\gamma(K \perp\!\!\!\perp p) = K \setminus \bigcap \gamma(K \perp p)$$

is an incision function for K. We call it the *associated* incision function of γ.

Observation 4.8. [87] Let σ be an incision function for a set K. The function γ_σ on $K \perp p$ defined as follows:

$$\gamma_\sigma(K \perp p) = \begin{cases} \{H : H \in K \perp p \text{ and } (K \setminus \sigma(K \perp\!\!\!\perp p)) \subseteq H\} & \text{if } K \perp p \neq \emptyset \\ \{K\} & \text{otherwise} \end{cases}$$

is a selection function on K.

Observation 4.9. [87] Let γ be a selection function for a set K, σ_γ the incision function defined from γ using Observation 4.7, and γ' the selection function defined from σ_γ using Observation 4.8, i.e., $\gamma' = \gamma_{\sigma_\gamma}$. Then:

1. $\gamma(K \perp p) \subseteq \gamma'(K \perp p)$
2. $\bigcap \gamma'(K \perp p) = \bigcap \gamma(K \perp p)$
3. $\gamma(K \perp p) = \gamma'(K \perp p)$ if whenever $\bigcap \gamma(K \perp p) \subseteq H$ for $H \in K \perp p$, then $H \in \gamma(K \perp p)$.

We can conclude from this that the seemingly different strategies of selecting sets to retain (partial meet contraction) and selecting sentences to remove (kernel contraction) yield essentially the same result. As we will see in Chapter 6, this sameness depends essentially on the logical closure of the set K.

4.5 Safe Contraction

In safe contraction, a relation on sentences is used to determine which sentences to discard from each p-kernel [7, 8]. In this operation, the elements of the belief set K are ordered according to a relation $<$ (traditionally called a hierarchy). Intuitively,

[4] For belief bases, however, the smooth kernel contractions form a proper subset of the partial meet contractions.

$q \prec r$ means that r should be retained rather than q if we have to give up one of them, and we say that "q is less safe that r". The ordering $\prec$ helps us to choose which element to remove from each kernel. The remaining beliefs are safe and form the outcome of the safe contraction of a belief set K by p (modulo $\prec$).

Definition 4.10. [7] Let $\prec$ be a relation on K. Any sentence q in K is *safe* with respect to p if and only if q is not minimal under $\prec$ among the elements of any $A \in K \perp\!\!\!\perp p$. The set of all safe sentences of K with respect to p is denoted by K/p.

Definition 4.11. [7] Let K be a belief set and $\prec$ a relation on K. The *safe contraction* on K associated with $\prec$ is equal to the set of consequences of the safe sentences:

$$K - p = Cn(K/p)$$

Since safe contraction is a special case of kernel contraction, the following theorem should not come as a surprise:

Theorem 4.32. [7] Any safe contraction on K is a partial meet contraction.

However, the converse of that theorem does not hold. In other words, there are partial meet contractions that are not safe contractions. In order to relate safe contraction to transitively relational partial meet contraction, additional properties on $\prec$ are needed:

Definition 4.12. [8] Let K be a belief set and $\prec$ a relation over K. Then $\prec$ is a *virtually connected hierarchy* over K if and only it satisfies:

If $p_1 \prec \cdots \prec p_n$, then not $p_n \prec p_1$. (acyclicity)
If $\vdash p \leftrightarrow p'$ and $\vdash q \leftrightarrow q'$, then: $p \prec q$ if and only if $p' \prec q'$. (intersubstitutivity)
If $p \prec q$ then either $p \prec r$ or $r \prec q$. (virtual connectivity)

Definition 4.13. [8] Let K be a belief set and $\prec$ a relation over K. Then, for all p, q and $r \in K$, $\prec$ is a *regular hierarchy* over K if and only it satisfies:

If $p \prec q$ and $q \vdash r$, then $p \prec r$. (continuing-up)
If $p \vdash q$ and $q \prec r$, then $p \prec r$. (continuing-down)

Theorem 4.33. [8, 289] Let K be a belief set and $-$ an operation on K. Then $-$ is a safe contraction, based on a regular and virtually connected hierarchy $\prec$ over K, if and only if $-$ is a transitively relational partial meet contraction.

Additional results on safe contraction, including characterizations of the effects on $-$ of various properties of $\prec$, have been obtained by Rott [289] and Rott and Hansson [301].

Chapter 5
Criticism of the Model

Although the AGM model is commonly considered to be the standard model of belief change, it has been subject to extensive criticism. In this chapter we make a résumé of this criticism. Much of the critical discussion has referred either to the postulates for partial meet contraction and revision, or to various aspects of the use of belief sets to represent belief states.

For additional elaborations on philosophical issues relating to the AGM model, see also [165, 190, 294].

5.1 The Recovery Postulate and Minimal Change

One of the basic principles of the AGM theory [4] is that belief changes should take place with minimal loss of previous beliefs. In the opinion of the AGM trio, the postulate of *recovery* guarantees minimal losses of contents in the contraction process [133, p. 65] [238, p. 352] [241, p. 478]. *Recovery* is based on the intuition that "it is reasonable to require that we get all of the beliefs ... back again after first contracting and then expanding with respect to the same belief" [130]. This can be exemplified as follows:

Example 5.1. [163] I believed that I had my latchkey on me (p). Then I felt in my left pocket, where I usually keep it, and did not find it. I lost my belief in p (but without starting to believe in $\neg p$ instead). Half a second later, I found the key, and regained my belief in p.

However, counter-examples have been constructed in which *recovery* seems to give rise to implausible results:

Example 5.2. [148] I believed that Cleopatra had a son (s). Therefore I also believed that Cleopatra had a child (c or equivalently $s \vee d$ where d denotes that Cleopatra had a daughter). Then I received information that made me give up my belief in c, and I contracted my belief set accordingly, forming $K - c$. Soon afterwards I learned from

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_5

a reliable source that Cleopatra had a child. It seems perfectly reasonable for me to then add c (i.e., $s \vee d$) to my set of beliefs without also reintroducing s.

This pattern is incompatible with recovery, and therefore also with partial meet contraction. The problem (if we see it as such) is that every element of $K \perp (s \vee d)$ has $s \vee d \rightarrow s$ as one of its elements, and therefore $s \vee d \rightarrow s$ is also an element of $K - (s \vee d)$. Consequently, $(K - (s \vee d)) + (s \vee d)$ implies s.

Example 5.3. [157][1] I previously entertained the two beliefs, "*x is divisible by 2*" (p) and "*x is divisible by 6*" (q). When I received new information that induced me to give up the first of these beliefs (p), the second (q) had to go as well (since p would otherwise follow from q).
I then received new information that made me accept the belief "*x is divisible by 8.*" (r). Since p follows from r, $(K - p) + p$ is a subset of $(K - p) + r$; thus by *recovery* I obtain that "*x is divisible by 24*" (s), contrary to intuition.

This example shows that retaining the sentence $p \rightarrow q$ after contraction of K by p gives rise to unintuitive results. There seem to be cases when this sentence has to be removed. Due to recovery, AGM contraction cannot eliminate it.

Makinson [241, p. 478] noted that "as soon as contraction makes use of the notion 'y is believed only because of x', we run into counterexamples [to *recovery*]". He argued that this is provoked by the use of a justificatory structure that is not represented in the belief set and that, without this structure, *recovery* can be accepted; or, in Makinson's words, it can be accepted in a "naked" theory. In [161], Hansson replied: "Actual human beliefs always have such a justificatory structure (...). It is difficult if not impossible to find examples about which we can have intuitions, and in which the belief set is not associated with a justificatory structure that guides our intuitions. Against this background, it is not surprising that, as Makinson says, *recovery* 'appears to be free of intuitive counterexamples' in which the belief set is not associated with a justificatory structure. It also seems to be free of confirming examples of this kind". Glaister argued that the problem exhibited in the counterexamples dissolves if we pay sufficient attention to exactly what is to be contracted. In the Cleopatra case, he claims, the contraction is more accurately represented by a multiple contraction by the set $\{s \vee d, s \vee \neg d, \neg s \vee d\}$ than by $s \vee d$ [141].

Niederée [259] pointed out several implausible formal properties that follow from recovery:

Observation 5.1. [259] Let K be a belief set and $p \in K$. Then, regardless of whether or not q is in K, *recovery* together with *closure* implies that:

1. $q \rightarrow p \in K - (p \vee q)$,
2. $p \in (K - (p \vee q)) + q$,
3. $\neg q \in (K - (p \vee q)) + \neg p$.

---

[1] We use here the modified version introduced in [105].

As we mentioned in Section 3.5, operations that satisfy the other five basic contraction postulates but not *recovery* are called withdrawals [239]. In the literature several withdrawals have been proposed [58, 95, 105, 223, 224, 250, 302].

An obvious reaction to these difficulties would be to replace *recovery* by some other postulate that puts a limit to how much can be removed in a contraction. However, no plausible such alternative postulate has been presented. In particular, although the following postulates are more intuitively appealing, they are equivalent to *recovery* in the presence of the other basic AGM postulates:

If $q \in K$ and $q \notin K - p$ then there is some set K' such that $K - p \subseteq K' \subseteq K$ and $p \notin Cn(K')$ but $p \in Cn(K' \cup \{q\})$. (relevance) [146]

If $q \in K$ and $q \notin K - p$ then there is some set K' such that $K' \subseteq K$ and $p \notin Cn(K')$ but $p \in Cn(K' \cup \{q\})$. (core-retainment) [148]

If $q \in K$ and $q \notin K - p$ then $K - p \nvdash p \vee q$. (disjunctive elimination) [91]

5.2 The Success Postulates

Partial meet revision satisfies the following postulate:

Revision success: $p \in K * p$

Several authors have found this to be an implausible feature of belief revision, even if p is not a contradiction. Hence Cross and Thomason pointed out that a system obeying this postulate

> "is totally trusting at each stage about the input information; it is willing to give up whatever elements of the background theory must be abandoned to render it consistent with the new information. Once this information has been incorporated, however, it is at once as susceptible to revision as anything else in the current theory.
>
> Such a rule of revision seems to place an inordinate value on novelty, and its behaviour towards what it learns seems capricious." [65]

Similarly, one of the AGM postulates for partial meet contraction:

Contraction success: If $\nvdash p$, then $p \notin K - p$,

has been contested on the grounds that we should "allow a reasoner to refuse the withdrawal of p not only in the case where p is a logical truth. There may well be other sentences ('necessary truths') which are of topmost importance for him" [290, p. 54]. Both with respect to *revision success* and *contraction success*, a common strategy among critics has been to construct AGM-style operations that do not always give primacy to the new information. (See Chapter 8.)

5.3 Remainder Sets: Information vs. Informational Value

Some researchers have argued that remainder sets retain "too much information". According to Levi [222, 223], some of the information in the belief set may have no value for the inquiring agent; consequently, the agent tries to retain as much of the valuable information as possible, instead of as much of the information as possible. He argued that measures of information should be replaced by measures of informational value,[2] and proposed an alternative construction, based on saturatable sets: H is a saturatable set with respect to p if and only if $H = Cn(H)$ and $Cn(H \cup \{\neg p\}) \in \mathcal{L} \perp\!\perp$. Furthermore, $S(K,p)$ is the set of saturatable sets with respect to p that are subsets of K. As Alchourrón and Makinson proved, $K \perp p \subseteq S(K,p)$ [5]. *Partial meet Levi contraction*, based on a selection among all the saturatable subsets of K with respect to p, is defined as $K - p = \bigcap \gamma(S(K,p))$, where γ is a selection function defined in the same way as in the AGM account. Hansson and Olsson [193] proved that an operation $-$ on K is a partial meet Levi contraction if and only if it satisfies *closure, inclusion, vacuity, success, extensionality*, and *failure*.

The main problem with this construction is that it allows for very drastic contractions. For instance, the following rather extreme operation is a partial meet Levi contraction: [193, pp. 111–112]

$$K - p = \begin{cases} K & \text{if } p \notin K \text{ or } \vdash p \\ Cn(\varnothing) & \text{otherwise} \end{cases}$$

5.4 The Expansion Property

It follows from the basic AGM postulates that the revision operation satisfies the following property:

If $K \nvdash \neg p$ then $K * p = Cn(K \cup \{p\})$. (expansion property of revision)

The *expansion property* of revision is just as tenacious as the *recovery* property of contraction. In the spheres model, it follows from the assumption that the possible worlds that are compatible with the present belief set form the innermost sphere (as in Figure 4.7). In the original AGM formulation of partial meet revision it follows from the Levi identity in combination with the contraction postulate *vacuity*.[3] However, it is easy to find examples in which it does not seem plausible:

[2] "... when seeking to answer a question, not all new information is relevant to the question being asked. This is, perhaps, the chief of several reasons why measures of informational value ought to be carefully distinguished from measures of information" [222, p. 123].

[3] According to the Levi identity, $K * p = Cn((K - \neg p) \cup \{p\})$. When $K \nvdash \neg p$, *vacuity* yields $K - \neg p = K$.

Example 5.4. [151] John is a neighbour about whom I initially know next to nothing.

Case 1: I am told that he goes home from work by taxi every day (t). This makes me believe that he is a rich man (r).

Case 2: When told t, I am also told that John is a driver by profession (d). In this case I am not made to believe that he is a rich man (r).

In case 1 we have $r \in K * t$, and due to the expansion property $K * t = Cn(K \cup \{t\})$. Since K is logically closed it follows that $t \rightarrow r \in K$. In case 2, the expansion property yields $K * (t\&d) = Cn(K \cup \{t\&d\})$. Combining this with $t \rightarrow r \in K$ we obtain $r \in K * (t\&d)$, contrary to the description of case 2.

This example exemplifies a quite common pattern of belief change: When we acquire a new belief that does not contradict our previous beliefs (such as t in the example), we often include in the outcome some additional belief (such as r in the example) that does not follow deductively but nevertheless serves to make the belief set more complete and/or more coherent.

The expansion property can also go wrong in the opposite direction, as illustrated in the following example:

Example 5.5. Valentina was uncertain whether or not her husband is unfaithful to her (u), but she still believed that her husband loves her (l). However, when she learnt that he is unfaithful to her, she lost her belief that he loves her.

In this case we have $l \in K$ and $K \nvdash \neg u$. The *expansion property* of revision requires that $K * u = Cn(K \cup \{u\})$; thus $l \in K * u$, contradicting the plausible pattern of belief change exhibited in the example.

The expansion property of revision has been much less discussed than the recovery property of contraction, but it is no less problematic and no less difficult to remove from the AGM framework. Both properties have provided impetus for the development of alternative frameworks.

In [268] and [269], Pagnucco and coworkers introduced a new belief change operation, *abductive expansion*. Unlike AGM expansion (consisting in $K+p = Cn(K \cup \{p\})$), in abductive expansion the agent also incorporates a justification or explanation of the new belief. The justification is the "abduction" of a formula p and it can be defined as follows:

Definition 5.1. [269] Let K be a belief set. An *abduction function* for K is a function f such that for each sentence p:

1. If $K \cup \{p\}) \nvdash \bot$, then $K \cup \{f(p)\} \vdash p$ and $K \cup \{f(p)\} \nvdash \bot$
2. If $K \cup \{p\}) \vdash \bot$, then $f(p) = \top$
3. If $\vdash p \leftrightarrow p'$, then $\vdash f(p) \leftrightarrow f(p')$

Definition 5.2. $\oplus$ is an *abductive expansion* for K if and only if there is an abduction function f fot K such that $K \oplus p = K + \{f(p)\}$ for all p.

The postulates that characterize abductive expansion are:

$K \oplus p$ is a belief set. (closure)
If $\neg p \notin K$, then $p \in K \oplus p$. (limited success)
$K \subseteq K \oplus p$. (inclusion)
If $\neg p \in K$, then $K \oplus p = K$. (vacuity)
If $\neg p \notin K$ then $\neg p \notin K \oplus p$. (consistency)
If $\vdash p \leftrightarrow q$, then $K \oplus p = K \oplus q$. (extensionality)

For the supplementary level, disjunctive factoring is added:

Either $K \oplus (p \vee q) = K \oplus p$, or $K \oplus (p \vee q) = K \oplus q$, or $K \oplus (p \vee q) = (K \oplus p) \cap (K \oplus q)$. (disjunctive factoring)

A semantic account of abduction can be based on Grove's systems of spheres. The construction is very similar to the one used in AGM belief revision. However, whereas $\|K\|$ is the innermost sphere in AGM revision, in abductive expansion there may be spheres that are proper subsets of $\|K\|$.

5.5 Are Belief Sets Too Large?

Belief sets have been criticized for being too extensive in two important respects that are both problematic from the viewpoint of cognitive realism: their logical closure and their infinite structure.

The use of a logically closed belief set to represent the belief state has important implications. In particular it means that all beliefs are treated as if they have independent status. Suppose you believe that you have your keys in your pocket (p). It follows that you also believe that either you have your keys in your pocket or the Archbishop of York is a Quranist Muslim ($p \vee q$). However, $p \vee q$ has no independent standing; it is in the belief set only because p is there. Therefore, if you give up your belief in p we should expect $p \vee q$ to be lost directly, without the need for any mechanism to select it for removal. In the AGM framework, however, "merely derived" beliefs such as $p \vee q$ have the same status as independently justified beliefs such as p. Belief base models (to be discussed in Chapter 6) have largely been constructed in order to distinguish between these two types of beliefs.

The logical closure of belief sets is also problematic from another point of view. In a study of the philosophical foundations of AGM, Hans Rott pointed out that the theory is unrealistic in its assumption that epistemic agents are "ideally competent regarding matters of logic. They should accept all the consequences of the beliefs they hold (that is, their set of beliefs should be logically closed), and they should rigorously see to it that their beliefs are consistent" [294].[4] In the same article he argued that the AGM model is not based on a principle of minimal change, something that has often been taken for granted.

[4] See Section 10.3 for resource-bounded agents.

However, as we noted in Section 3.2, logical closure only requires that the agent be able to draw the inferences that have been incorporated into the consequence operation *Cn*, and in the minimal case this does not go beyond classical sentential truth-functional logic. Furthermore, Isaac Levi has proposed that the belief set K should be interpreted as containing the statements that the agent is *committed to* believing, rather than those that she actually believes in [220, 222]. Such an interpretation may have other problems, but it defuses problems created by the high demands on inferential competence that seem to follow from logical closure.

Since actual human agents have finite minds, a good case can be made that a cognitively realistic model of belief change should be finitistic, and this in two senses. First, both the original belief set and the belief sets that result from a contraction should be finite-based, i.e., obtainable as the logical closure of some finite set. Secondly, the outcome set, i.e., the class of belief sets obtainable by contraction from the original belief set ($\{X \mid (\exists p)(X = K - p)\}$), should be finite [153, 179]. Partial meet contraction does not in general satisfy either of these two finitistic criteria. To the contrary, even if the original belief set K is finite-based, the standard AGM axioms do not ensure that the outcome $K - p$ of contracting it by a single sentence p is also finite-based [169].[5] And even if both K and $K - p$ are finite-based, the procedure that takes us from K to $K - p$ involves a choice among infinitely many sets (the elements of $K \perp p$), none of which is finite-based. Such a procedure does not seem to satisfy reasonable criteria of cognitive realism. This has led to the development of finitistic models such as belief base models (Chapter 6) and specified meet contraction (Section 4.3).

5.6 Lack of Information in the Belief Set

Belief sets have been criticized not only for being too large but also for lacking important information.

Most importantly, AGM contraction or revision in its original form is a "one shot" operation. After partial meet contraction of K by p we obtain a new belief set $K - p$ but we do not obtain a new selection function to be used in further operations on this new belief set. In other words, the original AGM framework does not satisfy the principle of *categorial matching*, according to which the representation of a belief state after a change should have the same format (and contain the same types of information) as the representation of the belief state before the change[6] [139]. In studies of iterated revision, various ways to extend the belief state representation to solve this problem have been investigated. (See Chapter 7.)

[5] Unless, of course, the language $\mathcal{L}$ is logically finite, by which is meant that it does not contain an infinite set of logically non-equivalent sentences. Arguably, this is a strong and implausible condition, since it excludes the possibility of expressing natural numbers of unlimited size.

[6] However, a fairly small rearrangement of the AGM definition is sufficient to make iterated partial meet contraction possible. We will return to that in Section 7.3.

Chapter 6
Belief Bases

It was understood from the beginning that the use of logically closed sets of sentences to represent belief states is not cognitively realistic. In an article published in 1985 Makinson pointed out that "in real life, when we perform a contraction or derogation, we never do it to the theory itself (in the sense of a set of propositions closed under consequence) but rather on some finite or recursive or at least recursively enumerable base for the theory" [238, p. 357]. The use of belief bases rather than (logically closed) belief sets has turned out to increase the expressive power of the belief change framework in important ways.

6.1 Representing Belief States with Belief Bases

Belief sets are very large entities. For instance, if p is in the belief set then so are both $p \vee q$ and $p \vee \neg q$ for any q in the language. Even for a very simple belief set such as $Cn(p)$, to contract by p involves making a choice between (excluding one or both of) $p \vee q$ and $p \vee \neg q$ for any q in the language. For computer implementations, the use of finite belief bases is a necessity.[1]

A belief base is a set A of sentences such that a sentence p is believed if and only if $A \vdash p$. Obviously, to each belief base A corresponds a belief set $Cn(A)$. Important distinctions can be introduced through the use of belief bases [66, 94, 114, 116, 146, 147, 150, 155, 255, 293, 327]. Although belief bases need not be finite, most work on them has focused on finite bases.

There are two opposing trends in the use of belief bases to represent the beliefs of an epistemic agent. One approach, supported by Dalal [66], is associated with a coherentist epistemic representation in which all elements of the belief set have equal status, and belief bases are a merely expressive resource. This interpretation requires that the outcome of a belief change operation is the same for different belief base representations of the same belief set. This principle is known as *irrelevance of*

[1] For an overview of the computational costs of performing changes on belief bases see [257].

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_6

syntax. For instance, if $Cn(A) = Cn(B)$, then $Cn(A-p) = Cn(B-p)$ and $Cn(A*p) = Cn(B*p)$.

In the other approach, the elements of the belief base A have another (and typically higher) epistemic status than those elements of $Cn(A)$ that are not in A. Consider the following example [167]: Suppose that the belief set contains the sentence p, "Shakespeare wrote Hamlet". Due to logical closure it then also contains the sentence $p \vee q$, "Either Shakespeare wrote Hamlet or Charles Dickens wrote Hamlet". The latter sentence is a "mere logical consequence" that should have no standing of its own, and it will therefore not be an element of the belief base. In this model, changes are made on the belief base. The intuition is that a sentence that is merely derived from others cannot survive if the beliefs that support it are removed. For another example, Susan believes that the moon rotates around the earth (p). She also believes that her copy of *Don Quijote* is on her bedside table (q). Therefore, she believes that the moon rotates around the earth if and only if her copy of *Don Quijote* is on her bedside table ($p \leftrightarrow q$). When entering the bedroom, she sees that the book is not on the nightstand as she had expected. She acquires the new belief $\neg q$ and consequently she cannot retain both her previous beliefs p and $p \leftrightarrow q$. According to the AGM approach (and other belief set approaches), both p and $p \leftrightarrow q$ are elements of the belief set on which choices are made, and she must make a choice which (if any) of them to retain. In the belief base approach p is in the belief base but $p \leftrightarrow q$ is a merely derived belief that disappears automatically when q is removed. It is not subject to any choice. This use of belief bases makes it possible to distinguish between different ways to hold the same beliefs. For instance $A_1 = \{p, q\}$ and $A_2 = \{p \wedge q\}$ have the same logical consequences, and consequently, they generate the same belief set. However, the difference between A_1 and A_2 is not just a "notational bondage" that should be straightened out by some process of "articulation" (Belnap [30], cited by Rott in [293]). In this case A and B represent genuinely different belief states.

Another important feature of belief bases is that they allow us to express the difference between different inconsistent belief states. This is not possible in the belief set approach, since there is only one inconsistent belief set. If K is inconsistent, then $K = \mathcal{L}$. In contrast, the following two belief bases:

$A_1 = \{p, \neg p, q_1, q_2, q_3, q_4, q_5, q_6\}$ and
$A_2 = \{p, \neg p, \neg q_1, \neg q_2, \neg q_3, \neg q_4, \neg q_5, \neg q_6\}$

are different although they have the same logical closure. We can expect $A_1 - \neg p$ and $A_2 - \neg p$ to be different consistent belief bases. This is of course a more sensible way to deal with inconsistencies than the conflation of all inconsistencies that is necessitated in operations on belief sets. Actual epistemic agents can find themselves in an inconsistent state, and they are able to extricate themselves from it without collapsing all distinctions. On local inconsistencies in belief bases, see [195].

6.2 Change Operations on Belief Bases

In order to define operations of change on belief bases, largely the same constructions can be applied as for belief sets. The expansion operation for a belief base A is defined as $A + p = A \cup \{p\}$.

Just as for belief sets, *partial meet contraction* can be defined in terms of remainder sets, such that for all p: $A - p = \bigcap \gamma(A \perp p)$. Hansson [147] axiomatically characterized this operation:

Theorem 6.1. Let A be a belief base. A function $-$ on A is a *partial meet base contraction* for A if and only if it satisfies:

If $\nvdash p$, then $A - p \nvdash p$. (success)
$A - p \subseteq A$ (inclusion)
If $q \in A$ and $q \notin A - p$, then there exists some A' such that $A - p \subseteq A' \subseteq A$ and $p \notin Cn(A')$ but $p \in Cn(A' \cup \{q\})$. (relevance)
If $p \in Cn(A')$ if and only if $q \in Cn(A')$ for all subset A' of A, then $A - p = A - q$. (uniformity)

Success and *inclusion* are just adapted from belief sets to belief bases. *Relevance* has much the same function as *recovery* has for belief sets, namely to prevent unnecessary losses of beliefs. It requires of an excluded sentence q that it contributes in some way to the fact that A implies p. *Uniformity* postulates that if two sentences p and q have the same behaviour in a belief base A, i.e., if they are implied by exactly the same subsets of A, then the outcomes of contracting by p and by q should be identical.

Observation 6.1. [163, p. 71]

1. If $-$ satisfies inclusion and relevance, then it satisfies

   If $\vdash p$, then $A - p = A$. (failure)
   If $A \nvdash p$, then $A - p = A$. (vacuity)

2. If $-$ satisfies uniformity, then it satisfies

   If $\vdash p \leftrightarrow q$, then $A - p = A - q$. (extensionality)

Kernel contraction can also be applied to belief bases. Thus, we can define a *kernel base contraction* for a base A based on an incision function σ such that for all p: $A - p = A \setminus \sigma(A \perp\!\!\!\perp p)$. Hansson [154] axiomatically characterized this operation:

Theorem 6.2. Let A be a belief base. A function $-$ on A is a *kernel base contraction* for A if and only if it satisfies: *success, inclusion, uniformity* and

If $q \in A$ and $q \notin A - p$, then there exists some A' such that $A' \subseteq A$ and $p \notin Cn(A')$ but $p \in Cn(A' \cup \{q\})$. (core-retainment)

Just as for belief sets, a kernel base contraction is *smooth* if and only if for all subsets A' of A: If $A' \vdash q$ and $q \in \sigma(A \perp\!\!\!\perp p)$, then $A' \cap \sigma(A \perp\!\!\!\perp p) \neq \varnothing$ [154]. A kernel base contraction is *relevant* if and only if for all $q \in \sigma(A \perp\!\!\!\perp p)$, there exists some X such that $(A \setminus \sigma(A \perp\!\!\!\perp p)) \subseteq X \subseteq A$, $X \nvdash p$ and $X \cup \{q\} \vdash p$ [87].

Theorem 6.3. Let A be a belief base and $-$ a kernel base contraction for A. Then:

1. [154] $-$ is smooth if and only if it satisfies

   $A \cap Cn(A - p) \subseteq A - p$. (relative closure)

2. [87] $-$ is relevant if and only if it satisfies *relevance*.

Since *relevance* implies *core-retainment*, it follows from Theorem 6.2 that all partial meet base contractions are kernel base contractions. However, the converse implication does not hold. The following example shows that even if a kernel contraction is smooth it need not be (reconstructible as) a partial meet contraction [163, pp. 91–92]. Let p, q, and r be logically independent sentences, and let $A = \{p, q, r\}$. We then have

$$A \perp\!\!\!\perp (p \wedge (q \vee r)) = \{\{p, q\}, \{p, r\}\}.$$

Now let σ be such that

$$\sigma(A \perp\!\!\!\perp (p \wedge (q \vee r))) = \{p, r\}.$$

It can straightforwardly be verified that a function σ with this property can be a smooth incision function for A. It follows that

$$A -_{\sigma} (p \wedge (q \vee r)) = \{q\}.$$

To see that $-_{\sigma}$ cannot be reconstructed as a partial meet contraction, consider the relevant remainder set:

$$A \perp (p \wedge (q \vee r)) = \{\{p\}, \{q, r\}\}.$$

A selection function γ for A must select a non-empty subset of $A \perp (p \wedge (q \vee r))$, i.e. $\gamma(A \perp (p \wedge (q \vee r)))$ is either $\{\{p\}\}, \{\{q, r\}\}$, or $\{\{p\}, \{q, r\}\}$. It follows that the contraction outcome is either $\{p\}, \{q, r\}$, or $\varnothing$. It cannot be $\{q\}$.

Contrary to partial meet and kernel contraction, epistemic entrenchment cannot be straightforwardly transferred to a belief base framework. However, Williams [332, 334, 335] defined an *ensconcement* relation on a belief base A as a transitive and connected relation $\preceq$ that satisfies the following three conditions:[2]

[2] $p \prec q$ means $p \preceq q$ and $q \npreceq p$.

(≤1) If $q \in A \setminus Cn(\varnothing)$, then $\{p \in A : q \prec p\} \nvdash q$.
(≤2) If $\nvdash p$ and $\vdash q$, then $p \prec q$, for all $p,q \in A$.
(≤3) If $\vdash p$ and $\vdash q$, then $p \leq q$, for all $p,q \in A$.

(≤1) says that the formulae that are strictly more ensconced than p do not (even conjointly) imply p. Conditions (≤2) and (≤3) say that tautologies are the most ensconced formulae. Given an ensconcement relation, a cut operation is defined by:

$$cut_A(p) = \{q \in A : \{r \in A : q \leq r\} \nvdash p\}$$

Observation 6.2. [333]

$$\text{If } p \in A, \text{ then } cut_A(p) = \{q \in A : p \prec q\}.$$

We can define a base contraction operation − using the *cut* operation:

Definition 6.1. [333] Let A be a belief base and ≤ an ensconcement relation. Then − is an ensconcement-based contraction if and only if:

$q \in A - p$ if and only if $q \in A$ and either (i) $p \in Cn(\varnothing)$ or (ii) $cut_A(p) \vdash p \vee q$.

Definition 6.2. [333] Let A be a belief base and ≤ an ensconcement relation. Then − is a brutal base contraction if and only if:

$q \in A - p$ if and only if $q \in A$ and either (i) $p \in Cn(\varnothing)$ or (ii) $cut_A(p) \vdash q$.

Ensconcement-based contraction is closely related to AGM contraction, whereas brutal base contraction is closely related to severe withdrawal [288][3].

Ensconcement-based contraction and brutal base contraction have been axiomatically characterized:

Theorem 6.4. [91, 100] Let A be a belief base. A function − on A is an *ensconcement-based contraction* for A if and only if it satisfies: *success, inclusion, vacuity, extensionality, disjunctive elimination*, [4] and:

$$A-(p \wedge q) = \begin{cases} A-p, \text{ or} \\ A-q, \text{ or} \\ A-p \cap K-q \end{cases} \quad \text{(conjunctive factoring)}$$

If $q \in A$, $p \notin A-(p \wedge q)$ and $q \notin A-(q \wedge r)$, then $p \notin A-(p \wedge r)$. (transitivity)
If $q \in A$ and $\{r \in A : q \notin A-(q \wedge r)\} \nvdash p$, then $q \in A-p$. (EB1)
If $q \in A-p$ then $\{r \in A : r \in A-(r \wedge p)\} \vdash p \vee q$. (EB2)

Theorem 6.5. [127] Let A be a belief base. A function − on A is a *brutal base contraction* for A if and only if it satisfies: *success, inclusion, vacuity, failure, relative closure*, and:

[3] For a thorough study of brutal base contraction, see [127].
[4] See page 43.

If $A - q \nvdash p$, then $A - q \subseteq A - p$. (strong inclusion)
If $q \in A$, $A \vdash p$ and $A - p = A - q$, then $p \in Cn(A - q \cup \{r \in A : A - q = A - r\})$. (uniform behaviour)

Williams [333] and Nebel [255] proposed the following way to define belief base contractions, where $\div$ is an operation for belief sets:

$$A - p = (Cn(A) \div p) \cap A$$

Fermé, Krevneris and Reis [91] proved that, in the previous formula, $-$ satisfies *success, inclusion, vacuity, extensionality* and *disjunctive elimination* if and only if $\div$ is a partial meet contraction. They call these operations *basic AGM-related base contractions*.

6.3 Belief Base Revision from Belief Base Contraction

As we mentioned above, one of the most important advantages of belief bases is that they make it possible to distinguish between different inconsistent belief states. This feature can be used to construct two substantially different types of revision operations based on contraction, depending on whether the negation of the added sentence is contracted before or after its addition:

$A * p = A - \neg p + p$ (internal revision, Levi identity) [4]
$A * p = A + p - \neg p$ (external revision, reversed Levi identity) [152]

The second of these options is not viable for a belief set, since if $\neg p \in K$ then the first of its two suboperations results in the set $K + p$ that contains the whole language and therefore removes all traces of the original belief set. Internal and external revision have both been axiomatized:

Theorem 6.6. [152] The operation $*$ is an operation of *internal partial meet revision* for a belief base A if and only if it satisfies:

$A * p \vdash p$ (success)
$A * p \subseteq A \cup \{p\}$ (inclusion)
$A * p$ is consistent if p is consistent. (consistency)
If $q \in A$ and $q \notin A * p$, then there is some A' such that $A * p \subseteq A' \subseteq A$, A' is consistent but $A' \cup \{q\}$ is inconsistent. (relevance)
If for all $A' \subseteq A$, $A' \cup \{p\}$ is inconsistent if and only if $A' \cup \{q\}$ is inconsistent, then $A \cap (A * p) = A \cap (A * q)$. (uniformity)

Theorem 6.7. [152] The operation $*$ is an operation of external partial meet revision if and only if it satisfies *success*, *inclusion*, *consistency*, *relevance*, and:

If p and q are elements of A and it holds that for all $A' \subseteq A$, $A' \cup \{p\}$ is inconsistent if and only if $A' \cup \{q\}$ is inconsistent, then $A \cap (A * p) = A \cap (A * q)$. (weak uniformity)
$A \cup \{p\} * p = A * p$. (pre-expansion)

It can easily be shown with examples that pre-expansion does not hold for internal partial meet revision, and that uniformity does not hold for external partial meet revision [152]. Therefore, neither of these operations is a special case of the other.

The distinction between different inconsistent belief bases also makes it possible to construct meaningful operations of consolidation, i.e., removal of inconsistency. (See Chapter 8.)

6.4 Other Belief Base Approaches

Nebel has proposed belief base operations in which a complete, reflexive and transitive relation over the elements of the belief base is used to prioritize among its elements [256]. This approach was further developed by Weydert who also related it to the AGM postulates [329].

Di Giusto and Governatori have developed an approach in which the elements of the belief base are divided into two categories, facts and rules. Facts are removed if that is necessary to accommodate new facts. Rules are not removed but can instead be changed. Hence, suppose that the belief base contains the fact $a \wedge b$ and the two rules $a \rightarrowtail c$ and $b \rightarrowtail c$. After revision by the new fact $\neg c$, a new belief base can be obtained that contains the facts $a \wedge b$ and $\neg c$ and the two rules $(a \wedge \neg b) \rightarrowtail c$ and $(b \wedge \neg a) \rightarrowtail c$ [73].

Bochman has developed a theory of belief revision in which an epistemic state is represented by a triple $\langle S, <, l \rangle$, where S is a set of objects called admissible belief states, $<$ a strict preference relation on these states, and l a function that assigns a (logically closed) belief set to each element of S. One and the same belief set may be assigned to several elements of S. This structure shares many features with belief bases [36].

It is commonly assumed that the belief base approach corresponds to a foundationalist epistemology, whereas the original AGM framework that applies operations of change directly to the belief set represents a coherentist view of belief change. Gärdenfors has provided the most thorough justification of this interpretation [134]. Del Val claimed that the two approaches are equivalent [68]. Doyle accepted Gärdenfors's analysis of the relationship between the belief set/belief base and coherentism/foundationalism distinctions. However, he argued that the fundamental concern for conservatism that Gärdenfors appealed to in his defence of coherentism applies equally to the foundationalist approach [79]. A more radical criticism was ventured in [194] where it was argued that the original AGM approach is incompatible with important characteristics of coherentism. In [164] Hansson claimed that the application of partial meet contraction to belief bases comes much closer to expressing coherentist intuitions than their application directly to belief sets.

6.5 Base-Generated Operations on Belief Sets

Operations on a belief set can be constructed by assigning to it a belief base and an operation on that belief base. Such operations will be called base-generated, and they can be introduced in the following style:

Definition 6.3. [153] An operation $-$ on a belief set K is a *base-generated partial meet contraction* if and only if there is a belief base A for K and an operation $\div$ of partial meet contraction for A such that for all sentences p: $K - p = Cn(A \div p)$.

Other types of base-generated contraction, such as base-generated kernel contraction, are defined in the same way.

Let $-$ be a base-generated operation on a belief set K, and let A be the base and $\div$ the operation on A from which $-$ is generated. It follows directly that if $\div$ satisfies *success* (i.e., $A \div p \nvdash p$ if $\nvdash p$), then so does $-$ (i.e., $K - p \nvdash p$ if $\nvdash p$). We can say that *success* is a *closure-invariant* postulate. The same applies to several other common contraction postulates:

Observation 6.3. [158] The following contraction postulates are closure-invariant: *inclusion, vacuity, success, extensionality* and *failure*.

The first four of these postulates coincide with four of the five postulates for a withdrawal in Makinson's sense [239]. The fifth postulate for withdrawals is *closure* ($K \div p = Cn(K \div p)$) that is satisfied by all base-generated contractions. Therefore, all base-generated partial meet contractions are withdrawals. However, they do not in general satisfy *recovery*, the sixth of the basic Gärdenfors postulates.

In order to fully characterize the properties of base-generated contractions, some new postulates are needed:

Observation 6.4. [154] An operation $-$ on a consistent belief set K is generated by an operation of kernel contraction for a finite base for K if and only if $-$ satisfies *closure*, *inclusion*, *vacuity*, *success*, and:

> There is a finite set A such that for every sentence p: $K - p = Cn(A')$ for some $A' \subseteq A$. (*finitude*)
> If it holds for all r that $p \in K - r$ if and only if $q \in K - r$, then $K - p = K - q$. (*symmetry*)
> If $K - q \nsubseteq K - p$, then there is some r such that $K - r \nvdash p$ and $(K - r) \cup (K - q) \vdash p$. (*weak conservativity*)

By strengthening weak conservativity to conservativity, we obtain a characterization of base-generated partial meet contraction:

Observation 6.5. [153] An operation $-$ on a consistent belief set K is generated by an operation of partial meet contraction for a finite base for K if and only if $-$ satisfies *closure*, *inclusion*, *vacuity*, *success*, *symmetry*, *finitude*, and:

> If $K-q \nsubseteq K-p$, then there is some r such that $K-p \subseteq K-r \nvdash p$ and $(K-r) \cup (K-q) \vdash p$. (*conservativity*)

Finitude is a fairly strong postulate that implies *finite-based outcome* and also ensures that there is only a finite set of possible contraction outcomes. *Symmetry* says that if two beliefs stand or fall together, then they yield the same contraction outcome. *Conservativity* and *weak conservativity* are two ways to express the requirement that in order for beliefs to be given up in the contraction by p they must in some way contribute to the fact that the belief set implies p. The relation between *conservativity* and *weak conservativity* is similar to that between *relevance* and *core-retainment*.

Chapter 7
Iterated Change

An AGM contraction or revision takes us from a belief set to a new belief set. In doing this, it makes use of a selection mechanism such as a selection function or an entrenchment relation. However, it does not provide a new selection mechanism to be used for further changes of the new belief set. The problem of constructing models that allow for iterated change is one of the most studied problems in the literature on belief change.

7.1 Revising Epistemic States

In order to represent iterated (repeated) belief change we need models in which the outcome of a belief contraction or a belief revision can itself be contracted or revised. This is not possible if the outcome of a contraction or revision consists only of a new belief set. It also has to contain information on how that new belief set will be changed in response to new inputs. Whereas standard AGM operations take us from a complete belief state (belief set + change mechanism) to an incomplete belief state (belief set only), for iterated change we need operations that take us from a complete belief state to another complete belief state.

Accounts of iterated belief change have used different representations of the change mechanism that is part of such a complete belief state. The most common of these is a preorder on the set of possible worlds, or equivalently a complete sphere system (cf. Section 4.1). The current belief set can be inferred from this preorder; it is simply the intersection of the worlds in the highest equivalence class (innermost sphere). A standard operation of change such as those presented in Section 4.1 provides us with a new belief set, or equivalently, with the innermost sphere of a new sphere system. To make additional changes possible, we need an operation of change that gives rise to a new complete sphere system, from which the new belief set can be inferred, and which can in its turn be subject to further changes, etc.

The most influential formulation of this approach is due to Darwiche and Pearl [67]. They propose that the epistemic state contains in addition to the belief set a

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_7

structure that commands the belief changes. The formal definition of an epistemic state used by Darwiche and Pearl is:

Definition 7.1. [151, 67] Let there be a set $\mathcal{E}$ of objects called *belief states*. s is a support function from $\mathcal{E}$ to $\mathcal{P}(\mathcal{L})$ such that for all $\Psi \in \mathcal{E}$, $s(\Psi)$ is a belief set.

If the elements of $\mathcal{E}$ are finite-based, then s can be replaced by a function B from $\mathcal{E}$ to $\mathcal{L}$ such that $B(\Psi)$ is a single sentence that is equivalent with the conjunction of $s(\Psi)$. Intuitively, $s(\Psi)$ is the belief set associated with Ψ, and if $s(\Psi)$ is finite-based, then $B(\Psi)$ is a single sentence representing it.

Darwiche and Pearl modified the Katsuno and Mendelzon postulates (R1)-(R6) for revision to work for epistemic states whose associated belief sets are represented by a single sentence.

(CR1) $B(\Psi * p) \vdash p$
(CR2) If $B(\Psi) \wedge p \nvdash \bot$ then $\vdash B(\Psi * p) \leftrightarrow B(\Psi) \wedge p$.
(CR3) If $p \nvdash \bot$, then $B(\Psi * p) \nvdash \bot$.
(CR4) If $\Psi_1 = \Psi_2$ and $\vdash p_1 \leftrightarrow p_2$ then $\vdash B(\Psi_1 * p_1) \leftrightarrow B(\Psi_2 * p_2)$.
(CR5) $B(\Psi * p) \wedge q \vdash B(\Psi * (p \wedge q))$
(CR6) If $B(\Psi * p) \wedge q \nvdash \bot$ then $B(\Psi * (p \wedge q)) \vdash B(\Psi * p) \wedge q$.

This modification involves a weakening of the extensionality postulate to (CR4) that allows the outcome of a belief revision to be different for different belief states although they have the same belief set.

Darwiche and Pearl introduced four additional postulates that are known as the DP-postulates for iterated revision:

(DP1) If $q \vdash p$ then $\vdash B((\Psi * p) * q) \leftrightarrow B(\Psi * q)$.
(DP2) If $q \vdash \neg p$, then $\vdash B((\Psi * p) * q) \leftrightarrow B(\Psi * q)$.
(DP3) If $B(\Psi * q) \vdash p$, then $B((\Psi * p) * q) \vdash p$.
(DP4) If $B(\Psi * q) \nvdash \neg p$, then $B((\Psi * p) * q) \nvdash \neg p$.

Postulate (DP1) states that the later evidence q cannot discredit the previous evidence p if q implies p (thus, in a sense, making p redundant). Postulate (DP2) states that if q contradicts the previous evidence p, then q completely eradicates the effect of p on the belief set. Postulate (DP3), on the other hand, ensures that p is retained after accommodating the more recent evidence q, given that p would be believed after revision by q. Lastly, postulate (DP4) stipulates that if p is not contradicted after revision by q, then it should not be contradicted after revision by first p, then q.

Darwiche and Pearl relate (DP1)-(DP4) with properties of total preorders in the possible world context. A global faithful assignment is defined as in Definition 4.4:

Definition 7.2. [205, 67] Let Ψ be a belief state. A total preorder $\leq_\Psi$ on possible worlds, with the strict part $<_\Psi$ and the symmetric part $\simeq_\Psi$, is a *global faithful assignment* associated with the belief state Ψ if and only if the following conditions hold:

1. If $B(\Psi) \subseteq \omega$ and $B(\Psi) \subseteq \omega'$, then $\omega \simeq_\Psi \omega'$.
2. If $B(\Psi) \subseteq \omega$ and $B(\Psi) \nsubseteq \omega'$, then $\omega <_\Psi \omega'$.

Theorem 7.1. [67] Let Ψ be a belief state:

1. An operation $*$ on Ψ satisfies (CR1)-(CR6) if and only if there is a global faithful assignment $\leq_\Psi$ for Ψ such that $\Psi * p = \bigcap(min(\|p\|, \leq_\Psi))$.
2. $*$ also satisfies (DP1)-(DP4) if and only if $\leq_\Psi$ satisfies:

(DPR1) If $p \in \omega_1$ and $p \in \omega_2$, then $\omega_1 \leq_\Psi \omega_2$ if and only if $\omega_1 \leq_{\Psi * p} \omega_2$.
(DPR2) If $\neg p \in \omega_1$ and $\neg p \in \omega_2$, then $\omega_1 \leq_\Psi \omega_2$ if and only if $\omega_1 \leq_{\Psi * p} \omega_2$.
(DPR3) If $p \in \omega_1$, $\neg p \in \omega_2$ and $\omega_1 <_\Psi \omega_2$, then $\omega_1 <_{\Psi * p} \omega_2$.
(DPR4) If $p \in \omega_1$, $\neg p \in \omega_2$ and $\omega_1 \leq_\Psi \omega_2$, then $\omega_1 \leq_{\Psi * p} \omega_2$.

According to (DPR1), the order among the p-worlds remains unchanged after revision by p. According to (DPR2), the order among the $\neg p$-worlds remains unchanged after revision by p. (DPR3) says that if a p-world is higher ranked than a $\neg p$-world, then it remains so after revision by p. (DPR4) says in addition that if a p-world is equally ranked with a $\neg p$-world, then after revising by p it either remains so or improves its relative position.

(DP1)-(DP4) have become the benchmark for iterated revision, and new proposals are almost invariably compared to them. However, convincing counterexamples can be given against each of the DP postulates. Stalnaker [319, pp. 205–206] has presented a counterexample to DP1. Counterexamples to DP2 have been proposed by Konieczny and Pino Pérez [212] and by Stalnaker [319, pp. 205–206]. (See also the discussions of DP2 by Lehmann [219] and by Jin and Thielscher [201].)

Jin and Thielscher [201] and Booth and Meyer [47] have pointed out that DP3 and DP4 are too permissive since they do not rule out operations in which all newly acquired information is given up as soon as an agent learns a fact that contradicts some of its current beliefs. Counterexamples to DP3 and DP4 have been given in [189]. The following is a counterexample to DP3:

Example 7.2. [189] During my time as a clerk at the headquarters of Destination Paradise airlines, service technicians discovered cracks in the petrol tanks of some of the older planes. These airplanes were always immediately grounded until the tank had been replaced by a new one. But some of the pilots were worried that deep cracks could develop in a few days. Such an occurrence in the interval between two service inspections could potentially cause a severe accident.

Case 1: At 9 a.m. I overheard a conversation in the coffee room. One of the secretaries said: "I have been told that airplane DP3 has a crack in the petrol tank." This made me believe that DP3 had, at 9 a.m., a crack in the petrol tank (p). It also made me believe that the plane was in for repair. One hour later my boss told me: "I have been called to a meeting. There was a terrible accident just a few minutes ago. DP3 has caught fire in the air and crashed, and apparently there are few chances that anyone has survived." This made me believe that DP3 had crashed (q). It also made me reverse my previous belief that it had a crack in its tank one hour ago, since in that case it would not have been in service. Hence, $K * p * q \nvdash p$.

Case 2: I did not overhear the coffee table conversation. But at ten o'clock my boss told me that DP3 had just caught fire in the air and crashed (q). Since a crack in the petrol tank was the only plausible cause of such an accident that I was aware of, I now also believed that there had been a crack in the petrol tank at 9 a.m. that day (p). Hence, $K * q \vdash p$.

Jin and Thielscher [201] and Booth and Meyer [47] have independently proposed the following condition (known as the *independence postulate*) instead of DP3 and DP4:

(Ind) If $B(\Psi * q) \nvdash \neg p$, then $B((\Psi * p) * q) \vdash p$.

Given the revision postulates (CR1)-(CR6), (Ind) is stronger than (DP3) and (DP4). (DP1), (DP2) and (Ind) are characteristic of a family of operations called admissible revision operations [47]. In global faithful assignment, (Ind) corresponds to the following postulate [201, 47]:

(R-Ind) If $p \in \omega_1$, $\neg p \in \omega_2$, and $\omega_1 \leq_\Psi \omega_2$, then $\omega_1 <_{\Psi * p} \omega_2$.

7.2 Major Classes of Iterable Operations

In a model of iterated belief revision there may be more than one way to arrive at one and the same belief set. Does it make any difference for further changes how we arrive at it? We can divide iterable operations into three classes according to their ability to remember the revision history and take it into account:

Operations without memory: In this case, each belief set is revised in a predetermined way, independently of how it was obtained, i.e.:

If $B(\Psi * p) = B(\Upsilon * p)$, then $B((\Psi * p) * q) = B((\Upsilon * p) * q)$.

Full meet revision is a trivial example of an iterable operation without memory. Areces and Becher have analyzed this class of operations [10].

Operations with full memory: In this case the full history of changes is conserved, so that rollbacks of previous changes are possible. Operations with full memory have been proposed, for example, by Brewka [56], Lehmann [219], and Konieczny and Pérez [212]. Falappa, García, Kern-Isberner and Simari proposed another type of revision in which discarded beliefs can be reused [89].

Operations with partial memory: In this case it makes a difference for future revisions how a belief set was arrived at, but the information remembered is not sufficient to reconstruct the previous states. Most of the proposed iterable revision operations are of this type.

In a recent review Rott recognized three major types of iterable revision operations on belief sets. They all have partial memory [299]:

Conservative revision, originally called natural revision, has been studied by Boutilier [54, 55] and Rott [297]. This operation is conservative in the sense that it only makes the minimal changes of the preorder that are needed to accept the input. In revision by p, the maximal p-worlds are moved to the bottom of the preorder which is otherwise left unchanged. The main characteristics of this operation are:

(Nat) If $B(\Psi * p) \vdash \neg q$, then $B((\Psi * p) * q) = B(\Psi * q)$.
(CRNat1) If $B(\Psi * p) \subseteq \omega_1$ and $B(\Psi * p) \nsubseteq \omega_2$, then $\omega_1 <_{\Psi * p} \omega_2$.
(CRNat2) If $B(\Psi * p) \nsubseteq \omega_1$ and $B(\Psi * p) \nsubseteq \omega_2$, then $\omega_1 \leq_{\Psi * p} \omega_2$ if and only if $\omega_1 \leq_{\Psi} \omega_2$.

Moderate revision, also called lexicographic revision, was originally studied by Nayak [253, 254]. When revising by p, this operation rearranges the preorder by putting the p-worlds at the bottom (but preserving their relative order) and the $\neg p$-worlds at the top (but preserving their relative order). It has the following properties:

(Lex) If $q \nvdash \neg p$, then $B((\Psi * p) * q) \vdash p$.
(CRLex) If $p \in \omega_1$ and $\neg p \in \omega_2$, then $\omega_1 <_{\Psi * p} \omega_2$

Radical revision is similar to moderate revision, but it differs in making the new belief irrevocable, i.e., impossible to remove. Segerberg characterized a class of such operations axiomatically [309]. It is further investigated in [98]. In radical revision by p based on a preorder of possible worlds, the relative order of the p-worlds is retained whereas the $\neg p$-worlds are removed from the preorder, thus becoming inaccessible. The main characteristic of this operation is:

(Irr) $B((\Psi * p) * \neg p) \vdash \bot$.
(CRIrr) For all $\omega_i \in \mathfrak{W}$, $\omega_i \subset B((\Psi * p) * \neg p)$.

Delgrande, Dubois and Lang argue that since revision assumes a static world, there is no reason why the outcome of an iterated revision should depend on the order of the inputs. Therefore, they propose that iterated revision should take the form of *prioritized merging*, a special case of multiple revision [70]. Several other types of iterable operations have been proposed; see for instance [51, 43, 57, 203, 210, 260].

7.3 Making AGM Contraction Iterable

We noted in Section 7.1 that any AGM contraction, standardly defined, can only be used for one particular belief set. Let us look more closely at the reason why this is so:

Original definition of partial meet contraction: [4]
(1) If $K \perp p \neq \varnothing$ then $\varnothing \neq \gamma(K \perp p) \subseteq K \perp p$.
(2) If $K \perp p = \varnothing$ then $\gamma(K \perp p) = \{K\}$.

(3) $K - p = \bigcap \gamma(K \perp p)$

Suppose that we wish to use the same selection function γ for two different belief sets K_1 and K_2. Let $\top$ be a tautology. It follows from clause (2) that $\gamma(K_1 \perp \top) = \{K_1\}$ and $\gamma(K_2 \perp \top) = \{K_2\}$, thus $\gamma(K_1 \perp \top) \neq \gamma(K_2 \perp \top)$. But we also have $K_1 \perp \top = K_2 \perp \top = \varnothing$, so for γ to be a function it must be the case that $\gamma(K_1 \perp \top) = \gamma(K_2 \perp \top)$. We can conclude from this contradiction that in this framework, each selection function can only be used for one belief set.

However, this formal problem has a surprisingly simple solution. We can rearrange the definition of partial meet contraction as follows:

Alternative definition of partial meet contraction:
(1') $\gamma(K \perp p) \subseteq K \perp p$ and if $K \perp p \neq \varnothing$ then $\gamma(K \perp p) \neq \varnothing$.
(2') $K - p = \bigcap \gamma(K \perp p)$ unless $\gamma(K \perp p) = \varnothing$ in which case $K - p = K$.

As applied to a single belief set K, this definition is equivalent with the original definition. The major difference is that it allows us to use one and the same selection function for all belief sets. With this simple reconstruction, partial meet contraction can be turned into a global operation, applicable to all belief sets. Since partial meet revision is definable from partial meet contraction via the Levi identity, this means that we have global AGM operations for both contraction and revision.

An obvious question to ask about this construction is: What properties does it give rise to in global (iterated) change, in addition to the properties of the classical operation that only works for a single belief set? The answer to that question may be surprising:

Observation 7.1. [172][1] Let $\mathcal{L}$ be infinite, let K_1 and K_2 be logically closed, and let $p_1 \in K_1 \setminus \mathrm{Cn}(\varnothing)$ and $p_2 \in K_2 \setminus \mathrm{Cn}(\varnothing)$. If $K_1 \perp p_1 = K_2 \perp p_2$ then $K_1 = K_2$ and $\vdash p_1 \leftrightarrow p_2$.

Corollary 7.1. [172] Let $\mathbb{X}$ be a set of belief sets, and for each $K \in \mathbb{X}$ let γ_K be a selection function for the set of all K-remainders.[2] Then there is a selection function γ for the set of remainders of elements of $\mathbb{X}$[3] that gives rise to a partial meet contraction γ such that $\gamma_K(K \perp p) = \gamma(K \perp p)$ for all $K \in \mathbb{X}$ and all sentences p.

The corollary tells us that any combination of (local) selection functions for each belief set K can be unified into a single, global selection function. Furthermore, it tells us what properties this construction gives rise to in addition to the properties of the local operation. The answer is: none.

This means that with the alternative definition, a single AGM-style selection function can be used to perform both contractions and revisions on all possible belief sets, without imposing any constraints on the properties of these operations other than the constraints already imposed by standard AGM operations on a single belief set.

---

[1] This does not hold for belief bases [152].

[2] This is the set $\{X \mid (\exists p \in \mathcal{L})(X \in K \perp p)\}$.

[3] This is the set $\{X \mid (\exists K \in \mathbb{X})(\exists p \in \mathcal{L})(X \in K \perp p)\}$.

Chapter 8
Non-Prioritized Change

In AGM revision, new information has primacy. This is mirrored in the success postulate for revision. At each stage the system has total trust in the input information, and previous beliefs are discarded whenever that is needed to consistently incorporate the new information. This is an unrealistic feature since in real life, cognitive agents sometimes do not accept the new information that they receive. There may be different reasons for this: the new information contradicts highly entrenched previous beliefs, the reliability of the source of the information is doubtful, or the information is too complex to be accepted more than partially. Belief revision that violates the success postulate is called non-prioritized belief revision.

8.1 Classification of Revision Operations

Revision operations can be classified with a focus on the revision process [162, 97]:

Integrated Revision/Choice. Choose among the originally believed sentences and the input in one single step.
Examples: entrenchment-based AGM revision [4], update [206], entrenchment-based and possible world variants of credibility-limited revision [191], Schlechta's and Rabinowicz's revision [306, 274], revision by comparison [107] and improvement operations [213].

Decision + Revision. (1) Decide whether to fully accept, partially accept, or reject the input. (2) Revise when appropriate.
Examples: screened revision [242] and selective revision [101].

Contraction + Expansion. (1) Remove the negation of the input. (2) Add the input.
Examples: AGM revision via the Levi identity and internal revision of belief bases [152].

Expansion + Contraction. (1) Add the input. (2) Remove its negation.
Example: external revision of belief bases [152].

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_8

Expansion + Consolidation. (1) Add the input. (2) Regain consistency by contracting by $\perp$.
Example: semi-revision [159, 261].

8.2 Some Constructions of Non-Prioritized Revision

In this section we will present some major proposals for the construction of non-prioritized revision.

Makinson [242] proposed *screened revision*, a simple model of non-prioritized belief revision. He introduced a set A of sentences that are immune to revision. The outcome of revising by sentences that contradict $K \cap A$ is identical to the original belief set. If the input sentence is compatible with $K \cap A$, then the belief set is revised essentially in the AGM way. Formally, screened revision for a belief set K is defined as follows:

$$K\#_A p = \begin{cases} K * p & \text{if } p \text{ is consistent with } K \cap A. \\ K & \text{otherwise} \end{cases}$$

where $*$ is an AGM revision function with the additional constraint that for all p, $K \cap A \subseteq K * p$.

A more general approach, called *generalized screened revision*, was proposed in [160]:

$$K\#_f p = \begin{cases} K * p & \text{if } p \text{ is consistent with } K \cap f(p). \\ K & \text{otherwise} \end{cases}$$

where f is a function such that for each sentence p, $f(p)$ is a set of sentences. $*$ is a (modified) AGM revision function such that for all p, $K \cap f(p) \subseteq K * p$. Different properties can be added to f. Makinson [242] proposed, for example, $f(p) = \{q : p < q\}$, where $<$ is a binary relation on the language.

Credibility-limited revision [191] is a revision model that is based on the assumption that the new information must stay within our limit of credibility in order to be accepted. The following example illustrates the underlying intuitions:

Example 8.1.

1. Carlos tells me: "Today I have lunch with Miguel." I believe him.
2. Santiago tells me: "Today I have lunch with the King of Sweden and the President of the USA." I don't believe him.

In the first case, we are disposed to accept the new information, but in the second case, our reaction is to reject it. The reason is that in the second case, the new belief exceeds our *credibility limit* for new information. The assertion that the King of Sweden and the President of the USA have lunch with Santiago is "too distant" from our corpus of beliefs. Formally, credibility-limited revision is defined by a set $\mathcal{C}$, consisting of all the credible sentences of the language, and a standard operation

$*$ of revision:

$$K \odot p = \begin{cases} K * p & \text{if } p \in \mathcal{C} \\ K & \text{otherwise} \end{cases}$$

where $*$ is an AGM revision function and $\odot$ is the credibility-limited revision induced by $*$ and $\mathcal{C}$.

This model has been axiomatically characterized, and it has also been developed in terms of epistemic entrenchment and possible world models [191]. In [103] the model was extended to belief bases and in [45] to iterated revision.

If a belief base is inconsistent, then it can be made consistent by removing enough of its more dispensable elements. This operation is called *consolidation*. The consolidation of a belief base A is denoted by $A!$. A plausible way to perform consolidation is to contract by falsum (contradiction), i.e., $A! = A - \bot$ [147].

Unfortunately, this recipe for the consolidation of inconsistent belief bases does not have a plausible counterpart for inconsistent belief sets. The reason is that since belief revision operates within classical logic, there is only one inconsistent belief set. Once an inconsistent belief set has been obtained, all distinctions have been lost, and consolidation cannot restore them.

Consolidation can be combined with expansion to construct *semi-revision*, an operation of non-prioritized revision for belief bases [159]:

$$A \odot p = (A + p)!$$

As a result of the consolidation, the input sentence may be discarded. Furthermore, contrary to most other forms of non-prioritized revision, the consolidation process may discard both p and $\neg p$. Therefore this model violates the "non-indifference principle" ($p \in K \odot p$ or $\neg p \in K \odot p$). Fuhrmann [118] extended this construction to inputs that are belief bases or belief sets, $A \odot B = (A \cup B)!$. Olsson proposed a coherentist version of semi-revision, where coherence takes the place of consistency [261]. Hansson and Wassermann proposed an operation that regains consistency only in a local part of the belief base, the part that is relevant for p and $\neg p$ [195].

Selective revision [101] introduces another way to deal with an input sentence: Only a part of it can be accepted. The following example illustrates the practical relevance of this option:

Example 8.2. [101] One day when you return back from work, your son tells you, as soon as you see him: "A dinosaur has broken grandma's vase in the living-room." You probably accept one part of the information, namely that the vase has been broken, while rejecting the part of it that refers to a dinosaur.

An operation $\odot$ of selective revision can be constructed from a standard revision operation $*$ and a transformation function f from and to sentences:

$$K \odot p = K * f(p)$$

In the intended cases, $f(p)$ does not contain any information that is not contained in p (in other words, $\vdash p \rightarrow f(p)$). By adding further conditions on f, various additional properties can be obtained for the operation of selective revision. Some plausible such properties are:

$\vdash f(f(p)) \leftrightarrow f(p)$ (idempotence)
If $\vdash p \rightarrow q$ then $\vdash f(p) \rightarrow f(q)$. (monotony)
If $\vdash p \leftrightarrow q$ then $\vdash f(p) \leftrightarrow f(q)$. (extensionality)
If $\nvdash \neg p$, then $\nvdash \neg f(p)$. (consistency preservation)
$\nvdash \neg f(p)$ (consistency)
If $K \nvdash \neg p$, then $\vdash f(p) \leftrightarrow p$. (weak maximality)
Either $\vdash f(p) \leftrightarrow p$ or $\vdash f(\neg p) \leftrightarrow \neg p$. (disjunctive maximality)

Improvement operations [213] are another family of non-prioritized belief revision operations. We present them in Section 10.2.

8.3 Non-Prioritized Contraction

The *success* postulate for contraction implies that all non-tautological beliefs are retractable. As was observed by Rott [290, p. 54], this is not a fully realistic requirement, since actual doxastic agents are known to have beliefs of a non-logical nature that nothing can bring them to give up. In non-prioritized contraction some non-tautological beliefs may be shielded from contraction.

An operation $\div$ of *shielded contraction* [102] can be based on an ordinary contraction operation $-$ and a set $\mathcal{R}$ of retractable sentences, so that:

$$K \div p = \begin{cases} K - p & \text{if } p \in \mathcal{R} \\ K & \text{otherwise} \end{cases}$$

This construction can be further specified by adding various requirements on the structure of $\mathcal{R}$. It has close connections with credibility-limited revision [102, 242]. In [103] shielded contraction was characterized for belief bases.

Chapter 9
Multiple Change

In the original AGM model the input is a single sentence. This is a limitation of the framework, since agents often simultaneously receive more than one piece of information. In models of multiple change, the input is a (possibly infinite) set of sentences. There are fundamental differences between iterated and (simultaneous) multiple belief changes, i.e., $K * p * q$ is not in general identical to $K * \{p, q\}$. However, the generalization of revision to the multiple case is simple if the input is finite: Revision by the set $\{p_1, ..., p_n\}$ corresponds to revision by the conjunction of its elements, $p_1 \wedge ... \wedge p_n$. It is much less obvious how to generalize contraction.

9.1 Choice and Package Contraction

Fuhrmann and Hansson identified two types of multiple contraction [121]. In *choice contraction*, a set of sentences is considered to be removed from the belief set K if and only if it is not a subset of the contraction outcome. Choice contraction can be denoted by $-_\exists$, and its success condition is

$B \nsubseteq K -_\exists B$, unless $B \subseteq Cn(\varnothing)$.

In *package contraction*, B is considered to be removed from K if and only if all of its elements have been removed. Package contraction can be denoted by $-_\forall$, and its success condition is

$B \cap (K -_\forall B) = \varnothing$, unless $B \cap Cn(\varnothing) \neq \varnothing$.

If B is finite, then choice contraction by B can plausibly be equated with standard single-sentence contraction by the conjunction of all its elements. Package contraction cannot be reduced to contraction by a single sentence, and it is therefore the most interesting of the two operations.

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_9

The following postulates are plausible properties for package multiple contraction:

$K -_\forall B = Cn(K -_\forall B)$ (p-closure)

$K -_\forall B \subseteq K$ (p-inclusion)

If $B \cap K = \varnothing$, then $K -_\forall B = K$. (p-vacuity)

If $B \cap Cn(\varnothing) = \varnothing$, then $B \cap (K -_\forall B) = \varnothing$. (p-success)

If for every sentence p in B there is a sentence q in C such that $\vdash p \leftrightarrow q$, and vice versa, then $K -_\forall B = K -_\forall C$. (p-extensionality)

$K \subseteq Cn((K -_\forall B) \cup B)$ (p-recovery)

If B is finite, then $K \subseteq Cn((K -_\forall B) \cup B)$. (p-finite recovery)

If every subset X of K implies some element of B if and only if X implies some element of C, then $K -_\forall B = K -_\forall C$. (p-uniformity)

If $p \in K$ and $p \notin K -_\forall B$, then there is a set K' such that $K -_\forall B \subseteq K' \subseteq K$ and $B \cap K' = \varnothing$ but $B \cap Cn(K' \cup \{p\}) \neq \varnothing$. (p-relevance)

If $p \in K$ and $p \notin K -_\forall B$, then there is a set K' such that $K' \subseteq K$ and $B \cap K' = \varnothing$ but $B \cap Cn(K' \cup \{p\}) \neq \varnothing$. (p-core-retainment)

The postulates *p-closure*, *p-inclusion*, *p-vacuity*, *p-success*, *p-extensionality*, *p-recovery* and *p-finite recovery* are generalizations of the basic AGM postulates (for contraction) that were originally presented in [146]. Analogously, the postulates *p-uniformity* and *p-relevance* are generalizations of the postulates *uniformity* and *relevance* for single sentence contraction [147, 149]. The postulate *p-core-retainment* generalizes the single sentence contraction postulate *core-retainment* [108].

Observation 9.1. Let K be a belief set and $-_\forall$ a multiple contraction on K. Then:

- If $-_\forall$ satisfies *p-inclusion* and *p-relevance*, then it satisfies *p-closure* [121].
- If $-_\forall$ satisfies *p-relevance*, then it satisfies *finite p-recovery* [121].
- If $-_\forall$ satisfies *p-inclusion*, *p-uniformity* and *p-relevance*, then it satisfies *p-vacuity* and *p-extensionality* [275].

9.2 Multiple Partial Meet Contraction

Most of the major AGM-related contraction operations have been generalized to multiple package contraction. We therefore have a wide variety of such operations, including multiple partial meet contraction [146, 121, 225], multiple kernel contraction [108], multiple specified meet contraction [171], and a multiple version of Grove's sphere system [276, 92].

For partial meet contraction, the first step is to generalize the definition of remainder sets. For any belief set K and set of sentences B, the *remainders* of K by B are the maximal subsets of K that do not imply any element of B. $K \perp B$ denotes the *remainder set* of K by B.

The definition of a *package selection function* is a straightforward generalization of the notion of a selection function (presented in Definition 3.6).

Definition 9.1. [146, 121] Let K be a belief set. A *package selection function* for K is a function γ such that for all sets of sentences B:

1. If $K \perp B$ is non-empty, then $\gamma(K \perp B)$ is a non-empty subset of $K \perp B$, and
2. If $K \perp B$ is empty, then $\gamma(K \perp B) = \{K\}$.

Finally we are in a position to introduce the definition of a *partial meet multiple contraction*:

Definition 9.2. [146, 121] Let K be a belief set and γ a package selection function for K. The *partial meet package contraction* on K that is generated by γ is the operation $-_{\forall\gamma}$ such that for all sets of sentences B:

$$K -_{\forall\gamma} B = \bigcap \gamma(K \perp B).$$

Partial meet multiple contraction has been axiomatically characterized as follows:

Theorem 9.1. [121] Let K be a belief set and $-_{\forall}$ a multiple contraction on K. Then $-_{\forall}$ is a partial meet package contraction if and only if it satisfies the postulates *p-inclusion*, *p-success*, *p-uniformity* and *p-relevance*.

A partial meet package contraction on a belief set also satisfies the postulates *p-closure*, *finite p-recovery*, *p-vacuity* and *p-extensionality* [121]. Li [225] has proved that partial meet package contraction does not in general satisfy *p-recovery* in the infinite case.

9.3 Multiple Kernel Contraction

Package kernel contraction was introduced in [108]. It is defined by means of kernel sets, i.e., minimal subsets of a belief base A that imply some element of the set B:

Definition 9.3. [108] Let A and B be two sets of sentences. The *package kernel set of A with respect to B*, denoted by $A \perp\!\!\!\perp_P B$, is the set such that $X \in A \perp\!\!\!\perp_P B$ if and only if:

1. $X \subseteq A$

2. $B \cap Cn(X) \neq \varnothing$
3. If $Y \subset X$, then $B \cap Cn(Y) = \varnothing$.

The *package incision function* for (a set) A is, roughly speaking, an operation that selects at least one element from each of the sets in $A \perp\!\!\!\perp_P B$, for any set B:

Definition 9.4. [108] A function σ is an *incision function* for A if and only if, for all sets B:

1. $\sigma(A \perp\!\!\!\perp_P B) \subseteq \bigcup(A \perp\!\!\!\perp_P B)$
2. If $\varnothing \neq X \in A \perp\!\!\!\perp_P B$, then $X \cap \sigma(A \perp\!\!\!\perp_P B) \neq \varnothing$.

We can now, finally, present the definition of package kernel contraction.

Definition 9.5. [108] Let σ be an incision function for A. The *package kernel contraction* $\approx_\sigma$ for A based on σ is defined as follows:

$$A \approx_\sigma B = A \setminus \sigma(A \perp\!\!\!\perp_P B).$$

Package kernel contraction has been axiomatically characterized:

Theorem 9.2. [108] A multiple contraction $-_\forall$ for a set of sentences A is a package kernel contraction if and only if it satisfies *p-inclusion*, *p-success*, *p-uniformity* and *p-core-retainment*.

The last two theorems show that every package partial meet contraction is a package kernel contraction.

9.4 Sphere-Based Multiple Contraction

Reis and Fermé [276] have presented the possible worlds counterpart of package partial meet contraction. More precisely, they have shown how the remainders can be defined in terms of possible worlds and, making use of that way of defining the remainders, how package partial meet contractions can be defined as intersections of (appropriate) sets of possible worlds.

We start by the definition of the set $\mathbb{X}_{K \perp B}$ that allows us to define a remainder set in terms of possible worlds.

Definition 9.6. [276] Let K be a belief set and B be a set of sentences. We denote by $\mathbb{X}_{K \perp B}$ the subset of $\mathcal{P}(\bigcup\{\|\neg\alpha_i\| : \alpha_i \in B \cap K\})$ such that a set $\mathfrak{X}$ of possible worlds is an element of $\mathbb{X}_{K \perp B}$ if and only if:

1. $\mathfrak{X} \cap \|\neg\alpha_i\| \neq \varnothing$, for all $\alpha_i \in B \cap K$.
2. If $\omega \in \mathfrak{X}$ then there is some $\alpha_j \in B \cap K$ such that $\mathfrak{X} \cap \|\neg\alpha_j\| = \{\omega\}$.

The relation between the set $\mathbb{X}_{K \perp B}$ and the remainder set $K \perp B$ is the following:

Observation 9.2. [276] Let K be a belief set and B a finite set of sentences. Then:

1. If $\mathfrak{X} \in \mathbb{X}_{K\perp B}$ then $\bigcap(\|K\| \cup \mathfrak{X}) \in K\perp B$.
2. If $X \in K\perp B$ then there is some $\mathfrak{X} \in \mathbb{X}_{K\perp B}$ such that $X = \bigcap(\|K\| \cup \mathfrak{X})$.
3. $K\perp B = \{\bigcap(\|K\| \cup \mathfrak{X}) : \mathfrak{X} \in \mathbb{X}_{K\perp B}\}$.

The first statement of the observation says that the addition of an element of $\mathbb{X}_{K\perp B}$ to $\|K\|$ gives rise to a set of worlds whose intersection is a maximal subset of K that does not imply any element of B. The second statement says that each of the remainders can be constructed in this way by means of an element of $\mathbb{X}_{K\perp B}$. The third shows how to construct the remainder set $K\perp B$ from $\mathbb{X}_{K\perp B}$.

$\mathbb{X}_{K\perp B}$ can be used to define package partial meet contraction in terms of possible worlds:

Theorem 9.3. [276] Let K be a belief set. An operation $-_\forall$ is a partial meet package contraction on K if and only if for any set of sentences B:

$$K -_\forall B = \bigcap(\|K\| \cup (\bigcup f(\mathbb{X}_{K\perp B}))),$$

where f is a propositional package selection function for $\|K\|$ such that for all sets of sentences B: (1) $f(\mathbb{X}_{K\perp B}) \subseteq \mathbb{X}_{K\perp B}$, and (2) if $\mathbb{X}_{K\perp B} \neq \varnothing$ then $f(\mathbb{X}_{K\perp B}) \neq \varnothing$.

Reis, Peppas and Fermé [277] proved that any method for constructing multiple contractions which is based on systems of spheres fails to fully characterize the class of transitively relational partial meet multiple contractions. They explain this informally as follows: Let K be a belief set and $-_\forall$ and $-'_\forall$ two package contractions on K induced (by the same method) from the systems of spheres $\mathbb{S}$ and $\mathbb{S}'$, respectively. Assume additionally that $-_\forall$ and $-'_\forall$ are such that for any sentence p of $\mathcal{L}$, $K -_\forall \{p\} = K -_\mathbb{S} p$ and $K -'_\forall \{p\} = K -_{\mathbb{S}'} p$, where $-_\mathbb{S}$ is the $\mathbb{S}$-based contraction on K and $-_{\mathbb{S}'}$ is the $\mathbb{S}'$-based contraction on K. It follows that if $-_\forall$ and $-'_\forall$ agree on all contractions by singleton sets, then it must be the case that $\mathbb{S} = \mathbb{S}'$, and therefore $-_\forall$ and $-'_\forall$ have to agree on contractions by sets of arbitrary size (i.e., $K -_\forall A = K -'_\forall A$ for any set of sentences A), no matter what the method of producing $-_\forall$ and $-'_\forall$ from $\mathbb{S}$ (and $\mathbb{S}'$) might be. On the other hand, there exist partial orders $\sqsubset$ and $\sqsubset'$ on the set of remainders of K that agree when restricted to remainders of K by singleton sets, but differ otherwise. Consequently, they induce transitively relational partial meet package contractions that are different from each other but, nevertheless, agree on all contractions by singleton sets.

Nevertheless, systems of spheres can be used to define various classes of partial meet package contraction. Different such classes of operations have been presented in [275, 92, 278]. Some of these methods have been translated to epistemic entrenchment and axiomatically characterized [93].

Other kinds of multiple change have been proposed in the literature. Spohn has proposed a ranking-theoretic account of multiple package contraction [317]. Fuhrmann has investigated operations such as the subtraction $p - q$ that asserts p with the exception of what q says, and the merge $p \circ q$ that extracts the maximal consistent content from p and q jointly [118, 119]. Finally, Zhang has investigated the combination of iterated and multiple contraction, represented by series of contractions such as $K - A_1 - ... - A_n$, where each A_k is a set of sentences [339].

Chapter 10
Alternative Operations of Change

In the original AGM model there are three major types of operations: contraction, revision, and expansion. Subsequently a large number of additional types of operations have been proposed. In this section we summarize some of them.

10.1 Update

In 1992, Katsuno and Mendelzon presented a type of operation of change that they called update [206]. Whereas revision operations are suited to capture changes that reflect evolving knowledge about a static situation, update operations are intended to represent changes in beliefs that result from changes in the objects of belief. The difference was pointed out for the first time by Keller and Winslett [207] (in the context of relational databases) and is captured in the following example [338]:

Example 10.1. Initially the agent knows that there is either a book on the table (p) or a magazine on the table (q), but not both.

Case 1: The agent is told that there is a book on the table. She concludes that there is no magazine on the table. This is revision.

Case 2: The agent is told that subsequently a book has been put on the table. In this case she should not conclude that there is no magazine on the table. This is update.

This difference is evident in the possible worlds approach. Katsuno and Mendelzon proposed that when the world changes, the agent changes each of the worlds that (s)he considers to be possible in order to accommodate the input while changing as little else as possible. They constructed update as follows (see Figure 10.1 for a graphical representation):

Definition 10.1. A *local faithful assignment* is a function mapping each possible world ω to a total preorder $\leq_\omega$ such that if $\omega \neq \omega'$, then $\omega <_\omega \omega'$.

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_10

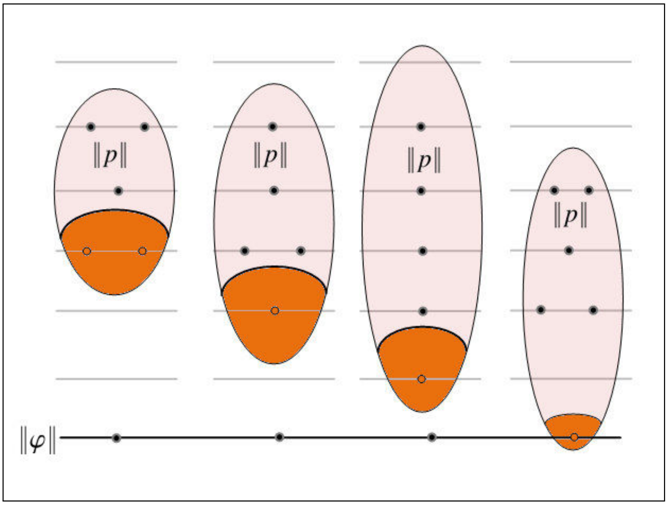

Fig. 10.1: Dots represent possible worlds and the line at the bottom represents $\|\varphi\|$. Each column represents the distribution and ordering of the possible worlds regarding the *local faithful assignment* of one world ω_i in $\|\varphi\|$. $\|\varphi \diamond p\|$ is the union of all the dark orange regions.

Definition 10.2. Let K be a finite-based belief set. Let $\varphi \in \mathcal{L}$ be such that $Cn(\varphi) = K$. An operation $\diamond$ on φ is an *update* if and only if there is a local faithful assigment such that:

$$\|\varphi \diamond p\| = \bigcup\{\min(\|p\|, \leq_\omega) : \varphi \in \omega\}$$

Update on finite-based belief sets has been axiomatically characterized as follows:

Theorem 10.2. [206] An operation $\diamond$ is an update operation if and only if it satisfies:

(U1) $\varphi \diamond p \vdash p$
(U2) If $\varphi \vdash p$, then $\vdash \varphi \diamond p \leftrightarrow \varphi$.
(U3) If $\varphi \nvdash \bot$ and $p \nvdash \bot$, then $\varphi \diamond p \nvdash \bot$.
(U4) If $\vdash \varphi_1 \leftrightarrow \varphi_2$ and $\vdash p_1 \leftrightarrow p_2$, then $\vdash \varphi_1 \diamond p_1 \leftrightarrow \varphi_2 \diamond p_2$.
(U5) $(\varphi \diamond p_1) \wedge p_2$ implies $\varphi \diamond (p_1 \wedge p_2)$.
(U6) If $\varphi_1 \diamond p_1 \vdash p_2$ and $\varphi_2 \diamond p_2 \vdash p_1$, then $\vdash \varphi_1 \diamond p_1 \leftrightarrow \varphi_2 \diamond p_2$.
(U7) If φ is complete,[1] then $(\varphi \diamond p) \wedge (\varphi \diamond q)$ implies $\varphi \diamond (p \vee q)$.
(U8) $\vdash (\varphi_1 \vee \varphi_2 \diamond p) \leftrightarrow (\varphi_1 \diamond p) \vee (\varphi_2 \diamond p)$.

It follows from (U2) that if $\varphi \vdash \bot$, then $\varphi \diamond p \vdash \bot$ for all p. In other words, if a belief set is inconsistent, then consistency cannot be regained with an update.[2]

[1] φ is complete if and only if for all $p \in \mathcal{L}$, $\varphi \vdash p$ or $\varphi \vdash \neg p$.

[2] With respect to this property, Katsuno and Mendelzon said [206, p. 190]: "We can never repair an inconsistent theory using update, because update specifies a change in the world. If there is no

If we compare update with AGM revision (see postulates (R1)-(R6) in Section 3.6), we can note some interesting formal differences. In particular, postulate R2 (vacuity) does not hold for update. Update and its relation with revision have been further studied by Becher [28] and others.

10.2 Changes in the Strengths of Beliefs

Sometimes when a statement is presented to us, this makes us consider it to be more credible than before, but we still do not believe it. Such a change may not affect the belief set, but it will affect how the belief state responds to new inputs. This kind of belief change was studied by Cantwell. He introduced the operations of *raising* and *lowering*, whereby the degree of plausibility required for a sentence to be included into the belief set is changed in either direction [57].

One important class of such operations is the improvement operations investigated by Konieczny and Pérez [213]. These operations do not (necessarily) satisfy the success postulate, although they increase ("improve") the agent's estimate of the plausibility of the new information [211, 213]. For instance, if you see what looks like wolf tracks in your garden, this makes it more plausible than before that a wolf has visited your garden, but you will presumably still not believe it.

In the construction of improvement operations it has often been assumed that if the agent receives the same new information sufficiently many times, then (s)he will finally believe it. For an epistemic state Ψ, an operation $\circ$ and a natural number n, $\circ^n$ is defined by recursion in the following way:

$$\begin{aligned} \Psi \circ^0 p &= \Psi \\ \Psi \circ^{n+1} p &= (\Psi \circ^n p) \circ p \end{aligned}$$

and the operation $\star$ is defined as $\Psi \star p = \Psi \circ^n p$, where n is the first integer such that $B(\Psi \circ^n p) \vdash p$. B is a function that takes us from a belief state to its associated belief set. The key property of this operation is:

> There exists an integer n such that $B(\Psi \circ^n p) \vdash p$. (iterative success)

Another interesting property is

> If $B(\Psi) \nvdash p$, then there is some q such that $\nvdash B((\Psi \circ p) \star q) \leftrightarrow B(\Psi \star q)$. (non-triviality)

This property says that any revision by a formula p that is not a consequence of the epistemic state modifies the epistemic state of the agent. In a possible worlds model, improvement by p means that some or all of the p-worlds are moved to a

---

set of worlds that fits our current description, we have no way of recording the change in the real world."

higher position in the preorder, but this does not necessarily lead to a change in the belief set.

Improvement operations have been combined with credibility-limited revision in the following way: A new piece of information is accepted if it is judged credible by the agent. However, if it is not considered credible, then its epistemic status is nevertheless raised, in the manner that this would be done by an improvement operation [46].

In quantitative theories of belief change, such as probabilistic and ranking theories (Sections 11.1 and 11.2), the degree of acceptance of each sentence is represented by a numerical value. Changes in the strength of beliefs can then be represented as changes in these values. However, the meaning of these numbers is not entirely clear (especially not for non-probabilistic functions), and real agents are notoriously bad at reasoning with them [324]. These difficulties are largely avoided by using a preorder instead of a numerical representation. Comparative degrees of belief can then be specified by taking certain beliefs as points of reference. The operation of change will adjust the position of an input sentence in an ordering to be the same as that of a reference sentence. Such an operation requires two sentences as inputs: the sentence to be adjusted and the reference sentence to which it will be adjusted. Since two sentences are involved, Rott called such operations two-dimensional [299].

Fermé and Rott proposed the operation of *revision by comparison*. In the intended case, the input sentence p is accepted to the same degree as a previously believed sentence q. However, if the negation of the input sentence p is more plausible than the reference sentence q, then q will be removed from the outcome [107]. Therefore, revision by comparison violates the DP postulates for iterated change that we mentioned in Section 7.1 (in particular DP2 since it collapses distinctions between the positions in the ordering of some $\neg p$-worlds). Rott has proposed a variant, *bounded revision*, that captures the spirit of revision by comparison and also satisfies the DP postulates [300]. As Rott pointed out, revision by comparison reduces the number of equivalence classes in the preorder, whereas bounded revision increases it.

10.3 Resource-Bounded and Local Change

AGM is a theory of changes of beliefs undertaken by highly idealized reasoners with unlimited cognitive capacities. In contrast, real reasoners such as humans, computers, and robots have limited resources. As was noted by Wassermann, it is important to distinguish between a limited implementation of a theory for ideal reasoning, and a theory for reasoners with limited resources [327].

Harman has put forward a highly useful list of principles that should be valid for any resource-bounded agent [196]:

Clutter avoidance: One should not clutter one's mind with trivialities.
Recognized implication: One has a reason to believe p if one recognizes that p is implied by one's views.

Recognized inconsistency: One has a reason to avoid believing things that one recognizes to be inconsistent.
Positive undermining: One should stop believing p whenever one positively believes that one's reasons for believing p are no good.
Conservatism: One is justified in continuing to fully accept something in the absence of a special reason not to.
Interest condition: One should add a new proposition p to one's beliefs only if one is interested in whether p is true (and it is otherwise reasonable for one to believe p).
Get back principle: One should not give up a belief one can easily (and rationally) get right back.

Doyle investigated characteristics of real agents such as mental inertia and constitutional elasticity [80]. He proposed a formal structure, a *reason maintenance system* (RMS), to capture these characteristics. Alechina, Jago, and Logan used RMS to construct a resource-bounded operation of contraction [9].

The two features of resource-boundedness that have attracted most attention among researchers are finitude and inconsistency tolerance. Both belief bases (Chapter 6) and specified meet contraction (Section 4.3) have been constructed largely in order to avoid the infinite structures of the standard AGM model.

Gabbay and Hunter maintain that there is a fundamental difference between how inconsistencies are handled by real agents and how they have usually been treated in formal logical systems. For a real agent it need not be necessary to restore consistency; it may be sufficient to have rules that specify how to act when an inconsistency arises [122]. What makes inconsistencies devastating in the AGM model is that there is only one inconsistent belief set, namely the whole language. This is an unsatisfactory feature of belief set representation, since two agents can both have inconsistent beliefs without having the same beliefs. As we saw in Section 6.1, belief bases fare much better in this respect. There are many different inconsistent belief bases, and they can reasonably be taken to represent different inconsistent belief states [147]. This feature of belief bases was employed in Hansson and Wassermann's model of *local change* [195]. Given a belief base B and a sentence r, the r-compartment of B is the subset of B that is relevant for r. In local change, revision of B by r involves changes only of the r-compartment; hence a part of the belief base can be made consistent while the belief base as a whole remains inconsistent. Wassermann has shown how these principles can be used to provide a model of change that satisfies Harman's principles. This can be accomplished with a construction involving a short-term memory in which recently computed results are temporarily stored [326]. She also showed how local change can be used for diagnosis [328], i.e., the process of finding a compartment that may have caused an abnormal behaviour of the system [279].

In a similar vein, Parikh [270] proposed a principle for relevance-sensitive change. Its basic principle is that if a belief set can be split into two independent parts (expressed in different sublanguages), then a revision of one part does not affect the other. This principle is not satisfied by AGM revision. Peppas provided a semantics for Parikh's relevance-sensitive condition in terms of systems of spheres

[271], and Kourousias and Makinson have investigated the conditions under which Parikh's relevance-sensitive condition is satisfied [216, 243].

10.4 Paraconsistent Belief Change

Consistency preservation is a central requirement in AGM revision. The reason for this is that the underlying logic is supraclassical and therefore satisfies the *explosion principle*, namely that anything follows from a contradiction (ex contradictione quodlibet, $\{p, \neg p\} \vdash q$). Consequently there is, as we just noted, only one inconsistent belief set, namely the whole language. If we arrive at an inconsistent belief set, then we have lost all distinctions. To avoid this we have to steer clear of contradictions in all operations on belief sets in a supraclassical logic.

However, this does not seem to be how cognitive agents behave in practice. Real agents can believe in contradictory statements without believing everything and losing all distinctions. In order to model that feature of actual reasoning, we can weaken the consequence relation and make it paraconsistent (which means that the explosion principle does not hold). Relatively little work has been performed on paraconsistent belief revision, but important contributions have been made for instance by Restall and Slaney [280], Priest [273], Mares [248], Tanaka [320], and Testa, Coniglio and Ribeiro [322].

The underlying logic used by Restall and Slaney [280] avoids the explosion principle by demanding a connection between the premises and the conclusion of an inference. In a valid inference the premises have to be *relevant* to the conclusion. Mares [248] developed a model in which an agent's belief state is represented by a pair of sets. One of these is the belief set, and the other consists of the sentences that the agent rejects. A belief state is coherent if and only if the intersection of these two sets is empty, i.e., if and only if there is no statement that the agent both accepts and rejects. In this model, belief revision preserves coherence but does not necessarily preserve consistency.

Priest [273] and Tanaka [320] suggested that in a paraconsistent logic, revision can be performed by just adding sentences without removing anything. In other words, if the logic tolerates inconsistencies, then expansion can serve the function usually assigned to revision. Furthermore, Priest [273] pointed out that in a paraconsistent framework, revision on belief sets can be performed as external revision, i.e., with the reversed Levi identity. In a supraclassical framework, external revision can only be used on belief bases. (See further Section 6.3.) Testa, Coniglio and Ribeiro [322] showed that this holds for semi-revision as well. In a supraclassical system, semi-revision (defined in Section 8.2) can only be used for belief bases, but in a paraconsistent system it can also be used for belief sets. The reason for this difference is that the intermediate inconsistent belief set that arises in external revision and semi-revision extinguishes all distinctions if the underlying logic is supraclassical but not if it is paraconsistent.

10.5 Some Other Operations of Change

Indeterministic change: The AGM model and most other models of belief change are deterministic in the sense that given a belief set and an input, the resulting belief set is well-determined. There is no scope for chance in determining the outcome of the change. Although this may not be a realistic feature, it substantially simplifies the formal structure. In indeterministic belief change, an operation can have more than one admissible outcome. Indeterministic belief change has been studied by Gallier [126] and by Lindström and Rabinowicz [231]. The latter authors gave up the assumption that epistemic entrenchment satisfies connectedness. This resulted in Grove-style sphere systems with spheres ("fallbacks") that are not linearly ordered but still all include the original belief set.

Replacement is an operation that replaces one sentence by another in a belief set. An operation of replacement has two variables, such that in $K|^p_q$, p has been replaced by q. Hence, the outcome is a belief set that contains q but not p. This operation can have outcomes that are not obtainable through either partial meet contraction or partial meet revision. Replacement can also be used as a kind of Sheffer stroke for belief change, i.e., an operation in terms of which the other operations can be defined. Contraction by p can be defined as the replacement $|^p_\top$ of p by a tautology, revision by p as the replacement $|^\bot_p$ of falsum by p, and expansion by p as the replacement $|^\top_p$ of a tautology by p. (Tautologies are as usual taken to be unremovable.) Partial meet replacement has been axiomatically characterized, and a semantic account in terms of possible worlds has been provided [170].

Reconsideration, introduced by Johnson and Shapiro [204, 203], is a non-prioritized operation on belief bases. It represents changes that are performed in hindsight in order to eliminate negative effects of previously performed changes. Previously removed beliefs can be reintroduced if there are no longer any valid reasons for their removal. This operation can be seen as an optimization that eliminates the negative effects of the order in which inputs have been received. It can involve an examination of all current and previous beliefs, but the same result can also be produced by an algorithm that examines a subset of the retracted basic beliefs, using dependency relationships.

Chapter 11
Extended Representations of Belief States

The AGM model is a simple and elegant representation of quite complex phenomena. Obviously, the trade-off between simplicity and relevance can be made differently. Many of the modifications of the framework that have been proposed consist in extensions of the belief state representation so that it contains more information than what is contained in a belief set or belief base. We have already studied one such extension in Sections 7.1 and 10.2. This chapter is devoted to additional extensions.

11.1 Probability and Plausibility

The AGM model and other logical approaches to belief revision represent features of doxastic behaviour that differ from those represented by probabilistic models. The degrees of belief represented for instance by entrenchment relations do not coincide with probabilities [298]. It seems difficult to construct a reasonably manageable model that covers both the logic-related and the probabilistic properties of belief change. (Problems connected with the lottery and preface paradoxes have a major role in creating this difficulty [217, 237].)

However, some authors have explored the interrelations between the two types of models. Lindström and Rabinowicz showed how belief revision can be connected with accounts of conditional probability that allow the condition to have probability zero [230]. Makinson further investigated this and other connections between the two frameworks [244]. Insights from the AGM model can be used as an impetus for considering "revisionary" accounts of conditional probability, i.e., accounts in which $\mathfrak{p}(q,r)$, the probability of q given r, is not defined in the standard way. (According to the standard definition, $\mathfrak{p}(q,r)$ is equal to $\mathfrak{p}(q\&r)/\mathfrak{p}(r)$ when $\mathfrak{p}(r) \neq 0$, and otherwise undefined.) Furthermore, the notion of non-prioritized revision that has been developed in the AGM tradition (revision not satisfying $p \in K * p$ for all p; see Chapter 8) can be usefully transferred to a probabilistic context. There it corresponds to "vacuous" conditionalizing when the condition is too unbelievable

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_11

to be taken seriously, i.e., $\mathfrak{p}(q,r) = \mathfrak{p}(q)$ when r is highly unlikely. Makinson also discussed "hyper-revisionary" probabilistic conditionalization, in which the fact that something believed to be very improbable actually happens is taken as a reason to believe that the probability was underestimated. There is an analogy between hyper-revisionary conditionalization and belief revision that violates the AGM property that if $K + p$ is logically consistent, then $K * p = K + p$. Such violations would be justified if K and p are epistemically but not logically incompatible.

In order to investigate the relationship between AGM and Bayesian conditionalization, Bonanno introduced what he called the qualitative Bayes rule, namely that

> "... if at a state the information received is consistent with the initial beliefs–in the sense that there are states that were considered possible initially and are compatible with the information–then the states that are considered possible according to the revised beliefs are precisely those states." [39]

Bonanno constructed and characterized a model of belief revision that satisfies this condition. It complies with the AGM postulates for partial meet revision.

Friedman and Halpern have developed a model based on a notion of plausibility that is a generalization of probability. Instead of assigning to each set A of sentences a number $\mathfrak{p}(A)$ in $[0,1]$, representing its probability, they assign to it an element $Pl(A)$ of a partially ordered set. $Pl(A)$ is called the "plausiblity" of A. If $Pl(A) \leq Pl(B)$ then B is at least as plausible as A. A sentence p is believed if and only if p is more plausible than $\neg p$. Changes in belief take the form of changes in the plausibility ordering. Conditions on such changes have been identified that produce a revision operation that is essentially equivalent with partial meet revision [111, 113].

Several other authors have presented probability-based and plausibility-based belief revision models that have close connections with the AGM model [20, 82, 13, 16].

11.2 Ranking Models

In Wolfgang Spohn's ranking theory of belief change, a belief state is represented by a ranking function κ that assigns a non-negative real number to each possible world W, representing the agent's degree of disbelief in W [314, 315, 316]. A sentence p is assigned the value $\kappa(p) = min\{\kappa(W) \mid p \in W\}$. Furthermore, p is believed if and only if $\kappa(\neg p) > 0$, i.e., if and only if every $\neg p$-world is disbelieved to a non-zero degree. The conditional rank of q given p is $\kappa(q \mid p) = \kappa(p\&q) - \kappa(p)$. For any sentence p and number x, the $p \mapsto x$-conditionalization of κ is defined by:

$$\kappa_{p\mapsto x}(q) = min\{\kappa(q \mid p), \kappa(q \mid \neg p) + x\}$$

Contractions, expansions and revisions can all be represented as conditionalizations. The outcome depends on the numerical values involved. In addition, other operations such as the strengthening or weakening of beliefs already held are straightforwardly representable in this framework. Important results on belief revision based on ranking functions, including an axiomatic representation that clarifies their relationship to AGM operations, have been reported by Hild and Spohn [198]. A generalization of Spohn's ranking functions has been proposed by Weydert [330].

11.3 Conditionals and the Ramsey Test

Belief revision theory has primarily been concerned with belief states and inputs expressed in terms of classical sentential (truth-functional) logic. The inclusion of non-truth-functional expressions into the language has interesting and often surprisingly drastic effects.

Among the several formal interpretations of non-truth-functional *conditionals*, such as counterfactuals, one is particularly well suited to the formal framework of belief revision, namely the so-called Ramsey test. It is based on a suggestion by Frank Ramsey (1903–1930) that has been further developed by Robert Stalnaker and others [318, pp. 98–112]. The basic idea is that "if p then q" is taken to be believed if and only if q would be believed after revising the present belief state by p. Let $p \mathrel{\Box\!\!\rightarrow} q$ denote "if p then q", or more precisely: "if p were the case, then q would be the case". The Ramsey test says:

$$p \mathrel{\Box\!\!\rightarrow} q \text{ holds if and only if } q \in K * p.$$

If we wish to treat conditional statements like $p \mathrel{\Box\!\!\rightarrow} q$ on par with statements about actual facts, then they will have to be included in the belief set when they are assented to by the agent; thus:

$$p \mathrel{\Box\!\!\rightarrow} q \in K \text{ if and only if } q \in K * p.$$

However, inclusion in the belief set of conditionals that satisfy the Ramsey test will require radical changes in the logic of belief change. As one example of this, contraction cannot then satisfy the inclusion postulate ($K - p \subseteq K$). The reason for this is that contraction typically provides support for conditional sentences that were not supported by the original belief state. Hence, if I give up my belief that John is severely mentally retarded, then I gain support for the conditional sentence "If John has lived 30 years in London, then he understands the English language" [151].

A famous impossibility theorem by Gärdenfors shows that the Ramsey test is incompatible with a set of plausible postulates for revision [132]. The crucial part of the proof consists in showing that the Ramsey test implies the following monotonicity condition:[1]

$$\text{If } K \subseteq K', \text{ then } K * p \subseteq K' * p.$$

[1] The proof is straightforward: Let $K \subseteq K'$ and $q \in K * p$. The Ramsey test yields $p \mathrel{\Box\!\!\rightarrow} q \in K$, then $K \subseteq K'$ yields $p \mathrel{\Box\!\!\rightarrow} q \in K'$, and finally one more application of the Ramsey test yields $q \in K' * p$.

This condition is incompatible with the AGM postulates for revision, and it is also easily shown to be implausible. Let K be a belief set in which you know nothing specific about Ellen and K' one in which you know that she is a lesbian. Let p denote that she is married and q that she has a husband. Then we can have $K \subseteq K'$ but $q \in K * p$ and $q \notin K' * p$.

Several solutions to the impossibility theorem have been put forward. One option investigated by Rott and others is to reject the Ramsey test as a criterion for the validity of conditional sentences [186, 285]. Levi accepts the test as a criterion of validity but denies that such conditional sentences should be included in the belief set when they are valid. In his view, they lack truth values and should therefore not be included in belief sets [221]. This, of course, blocks the impossibility result. Levi and Arló-Costa have investigated a weaker version of the Ramsey test that is not blocked by Gärdenfors's result and is also compatible with the AGM model [12, 15].

In a somewhat similar vein, Lindström and Rabinowicz have proposed that a conditional sentence expresses a determinate proposition about the world only relative to the subject's belief state. Given a conditional statement $p \mathrel{\Box\!\!\rightarrow} q$ and a belief set K, there is some sentence $r^K_{p \Box\!\!\rightarrow q}$ such that $p \mathrel{\Box\!\!\rightarrow} q$ holds in the belief state represented by K if and only if $r^K_{p \Box\!\!\rightarrow q} \in K$. In this way we can have the Ramsey test in the form

$$r^K_{p \Box\!\!\rightarrow q} \in K \text{ if and only if } q \in K * p,$$

which is not blocked by the impossibility result [233, 232].

Yet another option is to accept both the Ramsey test and the inclusion of conditional sentences into the belief set. Then belief sets containing $\Box\!\!\rightarrow$ will behave very differently under operations of change than the common AGM belief sets, and the standard AGM postulates will not hold [151, 286]. Not even the simple operation of expansion can be retained. Reusing an example from Section 5.4, we can suppose that you have no idea about John's profession, but then "expand" your belief set by the belief that he is a taxi driver. You will then lose the conditional belief that if John goes home by taxi every day, then he is a rich man–hence this is not an expansion after all [151]. As was noted by Rott, "[e]xpansions are not the right method to 'add' new sentences if the underlying language contains conditionals which are interpreted by the Ramsey test" [286].

Ryan and Schobbens have related the Ramsey test to update rather than revision (cf. Section 10.1) and found the test to be compatible and indeed closely connected with update operations [303].

Kern-Isberner has proposed a framework for revision that is based on a conditional valuation function that assigns (numerical) values to both non-conditional and conditional sentences. In this framework–which differs from AGM in important respects–conditional sentences can be elements of belief sets, and revisions can be performed with conditional sentences as inputs [209]. A partly similar approach has been developed by Weydert [331].

11.4 Modal, Doxastic, and Temporal Sentences

The inclusion of *modal sentences* in belief sets has been investigated by Fuhrmann. Let $\Diamond p$ denote that p is possible. The following, seemingly reasonable definition:

$$\Diamond p \in K \text{ if and only if } \neg p \notin K,$$

gives rise to problems similar to those exhibited in Gärdenfors's theorem, and essentially the same types of solutions have been discussed [115].

Lindström and Rabinowicz have investigated the inclusion into a belief revision framework of *introspective beliefs*, i.e., allowing for $\mathbb{B}p \in K$, where $\mathbb{B}p$ denotes "I believe p". Paradoxical results not unsimilar to those for conditionals are obtained in this case as well [234]. Similar results were obtained by Friedman and Halpern [112].

Dupin de Saint-Cyr and Lang introduced *temporally labelled sentences* into belief revision and proposed a belief change operation, called belief extrapolation, in which predictions are based on initial observations and a principle of minimal change [84]. Bonanno has developed logics that contain both a next-time temporal operation and a belief operation. The basic postulates of AGM revision are satisfied, and a strong version of the logic also satisfies the supplementary postulates [40, 41].

Booth and Richter have developed a model of fuzzy revision on belief bases. In this model, both the elements of the belief base and the input formulas come attached with a numerical degree (whose precise interpretation is left open). They showed that partial meet operations on belief bases can be faithfully extended to this fuzzy framework [52].

Finally, Fuhrmann has generalized partial meet operations to *arbitrary collections of (not necessarily linguistic) items* that have a dependency structure satisfying the Armstrong axioms for dependency structures in database relationships [117, 120].

11.5 Changes in Norms, Preferences, Goals, and Desires

Norms: The AGM model was partly the outcome of attempts to formalize changes in norms [5]. In spite of this, authors who tried to apply the AGM model to normative change have found it to be in need of rather extensive modifications to make it suitable for that purpose.

Boella, Pigozzi, and van der Torre analyzed normative change in a framework where a norm system is represented by a set of pairs of formulas. The pair $\langle p,q \rangle$ should be read "if p, then it is obligatory that q". In this framework, however, postulates for norm contraction and revision that are closely analogous to the AGM postulates give rise to inconsistency [38].

Governatori and Rotolo proposed a model for changes in legislation that, among several other aspects, also includes an explicit representation of time. Such a model can account for phenomena such as retroactivity that are difficult to deal with in an input-assimilating framework such as AGM [142].

Hansson and Makinson investigated the relationship between changes and applications of a norm system. In order to apply a norm system with conflicting norms to a particular situation, some of the norms may have to be ignored. Although these norms will remain intact for future situations, the problem of how to prioritize among conflicting items is similar to selecting sentences for removal in belief contraction [192].

Common law systems (case law systems), such as those of the United Kingdom, the United States, Canada, Australia, and New Zealand, have a structure that differs significantly from the civil law systems (Roman law systems) that dominate on the European continent and in Latin America. Most studies of changes in legal systems are primarily applicable to civil law systems in which the legal system changes through the modification of statutes. In common law systems, legal change takes place to a large extent through court decisions that are constrained by previous decisions made in other courts. John Horty [200] has investigated the nature of norm change in common law systems. It turns out to have interesting logical properties that differ from those of changes in statutory codes.

Preferences: A model of changes in preferences can be obtained by replacing the standard propositional language in AGM by a language consisting of sentences of the form $p \geq q$ ("p is at least as good as q") and their truth-functional combinations. The acquisition of a new preference takes the form of revision by such a preference sentence. The adjustments of the original preference state that are needed to maintain consistency in such revisions can be modelled by partial meet contraction. However, some modifications of the AGM model seem to be necessary in order to obtain a realistic model of preference change [144, 156, 218].

Goals and desires: Desires are often allowed to be contradictory, since they need not be actively pursued by the agent. Likewise, we often have goals that are difficult or impossible to combine. Intention selection is a process aimed at removing such contradictions, and ending up with a consistent set of intentions [265]. Paglieri and Castelfranchi have proposed a model of Data-oriented Belief Revision (DBR) in which significant attention is paid to the mutual influences between beliefs and goals [264]. In one direction, beliefs support and regulate goal processing and the transition from desires to intentions [62, 60, 61, 63]. In the other direction, goals can affect belief revision by determining the relevance (usefulness to the agent's current purposes) and likability (capability of fulfilling the agent's goal) of potential beliefs. Relevance increases the likelihood that a potential belief will be considered a candidate for belief, whereas likeability increases its chances of being actually believed, once considered. Boella, da Costa Pereira, Pigozzi, Tettamanzi, and van der Torre have also analyzed the role of goals in belief revision [37]. Their model is similar to DBR in its selection criteria, but it puts more emphasis on preventing wishful thinking from having any influence on the agent's beliefs.

Chapter 12
Applications, Connections, and Implementations

The AGM model has turned out to have a surprising number of connections with other areas of research.

12.1 Non-Monotonic and Defeasible Logic

In [135], Gärdenfors pointed out that belief revision and non-monotonic logic are motivated by quite different ideas. Belief revision deals with the dynamics of belief states, whereas non-monotonic logic is concerned with how we jump to conclusions from what we believe. But it is nevertheless possible to translate concepts, models, and results from one of these areas to the other. The first step towards such a translation was taken by Makinson and Gärdenfors [245]. Non-monotonic reasoning can be expressed by an inference operation $\mid\!\sim$ such that $A \mid\!\sim p$ denotes that A is a good enough reason to believe p, or in other words that p is a plausible consequence of A. We can also define an operation C that stands in the same relation to $\mid\!\sim$ as Cn to $\vdash$, i.e., $\mathrm{C}(A) = \{p \mid A \mid\!\sim p\}$. The crucial difference is that whereas Cn satisfies:

$$\text{If } A \subseteq B, \text{ then } \mathrm{Cn}(A) \subseteq \mathrm{Cn}(B) \text{ (monotony)}$$

C does not satisfy the corresponding property. Given a belief set K representing the background beliefs we can translate formulas between the two frameworks as follows [245]:

$$p \mid\!\sim q \text{ if and only if } q \in K * p.$$

Due to this connection, discussions in belief revision and non-monotonic logic have become increasingly interconnected. Lindström [229] and Rott [291] have shown that belief revision and non-monotonic logic are closely connected through their reconstructibility in terms of choice functions satisfying various rationality postulates. Other contributions in this tradition are [240, 287]. (See Section 12.7.)

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_12

In the last years of his life, Carlos Alchourrón published a series of articles on the logic of defeasible conditionals [1, 2, 3]. He proposed that conditional constructions in ordinary language can often be understood as saying that an antecedent p together with a set of assumptions is a sufficient condition for the consequent q. Such conditionals can be represented by a formula $\Box(f(p) \rightarrow q)$, where $f(p)$ is a function that takes us from p to the conjunction of p and these presuppositions. The connection between this approach and AGM was initially somewhat unclear [29], but in [106] axioms were given that relate it to a generalized version of AGM revision for an implicit underlying belief set K (not necessarily satisfying the postulates inclusion and vacuity). This result shows that there are close connections between Alchourrón's approach and Pagnucco's concept of abductive expansion, which is also closely associated with non-monotonic inference [268]. Billington, Antoniou, Governatori and Maher [35] have clarified the relationships between AGM revision and non-monotonic inference, showing that some of the AGM postulates do not hold or need to be modified for revision in a non-monotonic framework.

12.2 Modal and Dynamic Logics

We can distinguish between three ways to integrate belief revision with a modal logic.[1] First, one can add modal operators to the logical language (cf. Section 11.4). Secondly, one can formalize the execution of expansion, revision, and contraction with dynamic modal operators, similar to those for program execution. Thirdly, one can add both epistemic and dynamic modal operators, thereby integrating belief and belief change in one language. This has been done in the dynamic doxastic logic by Segerberg and collaborators, and later, but largely independently, in dynamic epistemic logic.

A reason to investigate belief revision in modal logic is that the theories of belief change developed within the AGM tradition are not logics in a strict sense, but rather informal axiomatic theories of belief change. Instead of characterizing the models of belief and belief change in a formalized object language, the AGM approach uses a natural language (ordinary mathematical English) to characterize the mathematical structures under study. The approaches to be mentioned here "internalize" the operations of belief change into the object language.

Explicit belief operators: An early approach was closely connected to non-prioritized belief revision. Let Bp denote that p is believed by the agent. Then $\Box Bp$ can denote that p is necessarily believed, i.e., it is believed and no amount of epistemic input can change this. $\Diamond\neg Bp$ means that it is possible to arrive at some state of belief in which p is not believed. $\Diamond\Box Bp$ signifies that it is possible to arrive at some state in which p is believed and can after that no longer be disbelieved, etc. Depending on the details of the revision process, this can be shown to give rise to either an S4.2 or an S4 logic for the modal operator [155].

[1] This section relies heavily on personal communication from Hans van Ditmarsch.

Dynamic modalities for belief revision: In various publications, van Benthem, de Rijke, and Fuhrmann [32, 116, 283, 33] introduced an "update logic" including the following notation:

$[\div p]q$ (q holds after contraction by p)
$[*p]q$ (q holds after revision by p)
$[+p]q$ (q holds after expansion by p)

This update logic can be seen as a precursor of subsequent treatments of belief change in dynamic logic.

Dynamic doxastic logic (DDL) extends propositional logical theories of belief sets with both Hintikka-style doxastic operators [199] and dynamic modal operators for belief change [307, 308, 310]. DDL was defined by Krister Segerberg as a logical framework for reasoning about doxastic change. The basic DDL represents an agent with opinions about the external world and an ability to change these opinions in the light of new information. Such an agent is non-introspective in the sense that it lacks opinions about its own belief states, for example $B[*p]q$ is not a well-formed formula. Lindström and Rabinowicz extended the model to include such formulas [235]. The result allows not only introspective agents but also iterated change, which has been studied by John Cantwell [57].

Dynamic epistemic logic (DEL) also studies changes in information and investigates actions with epistemic impact on agents [272, 24, 75]. Like DDL it has epistemic operators for belief or knowledge and also dynamic operators for changes in belief or knowledge. The best-studied dynamics is that of the *public announcement* of a formula ϕ. This can—with many reservations—be seen as a kind of belief expansion with ϕ. One such reservation is that the postulate of success is not necessarily satisfied: after public announcement of ϕ it need not be the case that $B\phi$ is true, i.e., that ϕ is believed. The standard counterexample is the Moore-sentence $p \wedge \neg Bp$ [252, 199]. Clearly, $B(p \wedge \neg Bp)$ is inconsistent for standard notions of knowledge and belief. In DEL, the Moore-sentence is just one example of an unsuccessful update; see [76].

To model belief revision in DEL, we need Kripke models where knowledge, belief, and degrees of belief (or conditional belief) can all be encoded. This can be done by adding *plausibility relations* to the Kripke models, and identifying belief in ϕ with truth of ϕ in the most plausible of the epistemically accessible states. For a simplified example, suppose that only the two states s and t are considered possible by an agent, but (s)he considers s to be more plausible than t. Furthermore suppose that p holds in s but not in t. Then the agent believes that p is true. Belief revision with $\neg p$ revises the plausibilities such that t becomes more plausible than s. Now, the agent believes that p is false: $B\neg p$. So we have $Bp \wedge [*\neg p]B\neg p$, where $[*\neg p]$ is not a 'hard' (i.e., truthful) public announcement but a tentative or 'soft' (i.e., preference changing) public announcement. An alternative belief revision mechanism in the DEL setting would eliminate the state s from consideration (a 'hard' update, as for the execution of public announcements), after which t is the only remaining state. Again, the agent believes that p is false. These issues were addressed by Aucher

[19], van Benthem [34], van Ditmarsch and Labuschagne [74, 77], and Baltag and Smets [25] (with many follow-up papers). Conditional reasoning and reasoning with different degrees of belief can also be modelled in such settings.

12.3 Horn Clause Contraction

In a formal language based on atomic sentences, a sentence is called a *literal* if it is either an atomic sentence (also called a positive literal) or the negation of an atomic sentence (also called a negative literal). A *clause* is a disjunction of literals. A *Horn clause* is a clause with at most one positive literal. (This makes p_1, $\neg p_1$ and $\neg p_1 \vee \neg p_2 \vee \neg p_3 \vee \neg p_4$ Horn clauses.)

The logic of Horn clauses has found extensive use in artificial intelligence, database theory, and logic programming, in particular in applications to truth maintenance systems and deductive databases. With the exception of a few early contributions [94, 284], studies of Horn clause belief change started only recently.

Delgrande [69] and Delgrande and Wassermann [72] investigated the contraction of theories expressed in propositional Horn logic. The main conclusions were the following:

- There are at least two plausible types of contraction functions for a Horn clause framework. In *entailment-based contraction (e-contraction)* the formula to be contracted is entailed by a belief set and the desired outcome is a belief set which does not entail that formula. In *inconsistency-based contraction (i-contraction)* the addition of a formula would lead to inconsistency and the desired outcome is a smaller belief set to which that formula can be added.
- Recovery is not a desirable property for contraction of Horn clauses. Such contractions have some features that are usually associated with the contraction of belief bases.
- In Horn clause contraction, maxichoice contraction appears to constitute an appropriate approach.

Regarding the third point, Delgrande and Wassermann [72] found that the application of contraction operations to remainder sets (as in AGM) has undesirable properties in Horn clause logic. Instead they developed an account of maxichoice Horn contraction in terms of what they call weak remainder sets (defined semantically instead of syntactically).

Booth, Meyer, and Varzinczak provided a generalization of Delgrande's partial meet constructions for Horn clauses (infra-contraction) [49] (cf. [121]). Booth, Meyer, Varzinczak, and Wassermann [50] showed that this construction coincides with standard kernel contraction, as applied to Horn clauses. Fotinopoulos and Papadopoulos obtained a characterization of Horn contraction in terms of a system of spheres [110], and Zhuang and Pagnucco characterized Horn contraction in terms of epistemic entrenchment [342].

12.4 Description Logic

Description logics (DL) is a family of formal knowledge representation languages.[2] Description logics are decidable fragments of first-order logic. These fragments are more expressive than propositional logic. Originally descended from semantic networks, they describe domains in terms of classes, properties, relationships and individuals.

In order to represent the domain, description logics use two structures: the TBox (terminological box) and the ABox (assertional box). The TBox contains sentences describing concept hierarchies (i.e., relations between concepts) while the ABox contains sentences stating where in the hierarchy individuals belong (i.e., relations between individuals and concepts). For example, the statement "all Portuguese are Europeans" belongs in the TBox, while "João is Portuguese" belongs in the ABox.

Description logics are used for knowledge representation, in particular as the basis for widely used languages that represent formal ontologies. Description logics have been successful in detecting incoherences in databases (ontology debugging), but provide little support for resolving these incoherences (ontology repair). Methodologies for belief change can be used to improve their performance in that respect. Techniques for ontology debugging are closely related to the identification of the kernel set in kernel (and safe) contraction. Furthermore, methods for combining information from different sources (merging) can be useful for this task. Using proposals by Benferhat, Kaci, Le Berre and Williams [31], Meyer, Lee and Booth [251] proposed basic strategies for solving incoherences in a framework based on description logic. Ribeiro and Wassermann have proposed a belief revision approach to finding and repairing inconsistencies in ontologies represented in description logics [282].

12.5 Belief Change by Translation Between Logics

The AGM model has been axiomatized and described for classical logic. One strategy to extend it to other logics is to translate the source logic into classical logic, perform the change and then translate back the outcome into the source logic. This method was proposed independently by two groups. Gabbay, Rodrigues and Russo [123, 124] proposed to translate a non-classical logic L into first-order classical logic, perform the revision there and then translate the results back. For example, let τ be a translation function from L into classical logic, and let T^τ be a classical logic encoding the basic properties of L. If the axiomatization of $T^\tau$is sound and complete, then for any belief set K and sentence α in the logic L we have: $K \vdash_L \alpha$ if and only if $T^\tau \cup K^\tau \vdash \alpha^\tau$, where K^τ and α^τ are the translation in classical logic of K and α respectively. Therefore, a revision operation $*_L$ in the logic L can be defined as $K *_L \alpha = \{\beta : K^\tau * (\alpha^\tau \wedge T^\tau) \vdash \beta^\tau\}$.

[2] For a detailed overview see [23].

Coniglio and Carnielli [64] define logics as two-sorted first-order structures, and they argue that this broad definition encompasses a wide class of logics that are interesting theoretically as well as from the point of view of applications. The language, concepts and methods of model theory can be used to describe the relationships between logics through morphisms of structures called transfers. This leads to a formal framework for studying several properties of abstract logics and their attributes such as consequence operators, syntactical structures, and internal transformations. They define a model of belief change, called Wide Belief Revision Systems, to define belief revision for non-standard logics.

12.6 Truth and Learning

Formal learning theory uses mathematical models to investigate the assimilation of new information. However, it differs from belief revision theory in attending to different types of problems. In learning theory, the main focus is on inductive problems. These are problems in which data is accumulated and processed in order to answer some empirical research question. The research question is usually represented by a partition of the set of possible worlds, i.e., a collection of non-empty sets of possible worlds such that each possible world is an element of exactly one of them. The question has been fully answered when we know in which of the elements of the partition the actual world is located. The cognitive agent receives a series of data that successively eliminates some of the possible answers. The central issue for the learning theorist is to construct inductive strategies, strategies for drawing conclusions from these data as they arrive [208]. Such a strategy should answer the following question:

> (LT) What–if any–conclusion concerning our research question should we draw from some sequence $p_1, p_2, \ldots, p_n$ of data (observations)?

A central requirement on an inductive strategy is that as more and more data is received, it should converge to the truth, i.e., to a true answer to the research question. In many problems we cannot know for sure when we have reached the truth. It is characteristic of an inductive process that we have to interrupt it at some stage when we have an answer to the research question that is safe enough for our purposes. Otherwise we might have to go on for ever. (LT) is closely related to a question in belief revision, namely:

> (BR) What is the belief set that we obtain by subjecting the original belief set K to a series of revisions $K * p_1 * p_2 * \cdots * p_n$?

In order to answer this question we need operations for iterated belief revision. (See Chapter 7.) But importantly, belief revision is performed with specifications (goals) that differ from those of learning theory. As we have just noted, the dominant concern of learning theory is to find the truth or, more specifically, a true answer to a given question. In contrast, belief revision is not concerned with finding the truth.

Its dominant concern is to perform the assimilation of new information in a way that changes the original belief set as little as possible. Not surprisingly, this fundamental difference has implications for the properties of the assimilation process. Genin and Kelly have shown that an operation of revision cannot provide a credible account of the inductive processes studied in learning theory unless several of the standard AGM postulates are dispensed with [140]. The relationship between the two types of operations is a promising but largely unexplored research topic.

12.7 Connections with Social Choice

The classical theory of rational choice was developed by mathematical economists [18, 312, 313]. It has occupied a central role in the social sciences for more than half a century.

Choice rationality is concerned with how to choose rationally among a set of alternatives. The standard representation of choice is a *choice function*. A choice function is defined over a set $\mathcal{A}$ of alternatives, and for each subset of that set it chooses, intuitively speaking, the most choiceworthy (or "best") alternatives. Formally, C is a choice function for $\mathcal{A}$ if and only if it is a function such that for each subset $\mathcal{B}$ of $\mathcal{A}$: (1) $C(\mathcal{B})$ is a subset of $\mathcal{B}$ and (2) if $\mathcal{B}$ is non-empty then so is $C(\mathcal{B})$. Notably, $C(\mathcal{B})$ can have more than one element. However, this does not mean that the agent chooses more than one alternative, only that there is more than one alternative that she is willing to choose. Among the rationality properties that have been proposed for choice functions, the following two are the most important ones [311, 312]:

Chernoff (property α)
If $\mathcal{B}_1 \subseteq \mathcal{B}_2$ then $\mathcal{B}_1 \cap C(\mathcal{B}_2) \subseteq C(\mathcal{B}_1)$.

Property β
If $\mathcal{B}_1 \subseteq \mathcal{B}_2$ and $X, Y \in C(\mathcal{B}_1)$, then: $X \in C(\mathcal{B}_2)$ if and only if $Y \in C(\mathcal{B}_2)$

Our choices are to a large extent guided by our preferences. In economics and social choice theory it is often assumed that a rational agent's choices are completely determined by his or her preferences in the following way:

$$C(\mathcal{B}) = \{X \in \mathcal{B} \mid (\forall Y \in \mathcal{B})(X \geq Y)\}$$

A choice function C is *relational* if and only if it is based on some preference relation $\geq$ in this way. There are tight formal connections between the properties of choices and those of preferences. It is possible to base a choice function on a given preference relation $\geq$ if and only if $\geq$ satisfies completeness and acyclicity. All such choice functions (i.e., all relational choice functions) satisfy Chernoff. A relational choice function satisfies property β if and only if the underlying preference relation is transitive [311].

The structures studied in rational choice theory have a close connection with the selection functions employed in belief revision. Rott relates belief revision, non-monotonic reasoning and rational choice, and shows how standard postulates of belief change and non-monotonic reasoning correspond to the constraints of classical theories of rational choice. According to Rott, these connections constitute an important bridge between practical and theoretical rationality [295].

Olsson [262] concedes that Rott's work is indisputable as a formal achievement, but puts his philosophical conclusions in question. According to Olsson, Rott has not discovered any surprising connections between belief revision and choice. Instead he has reconnected the AGM theory with its roots in counterfactual reasoning and rational choice.

Further studies on the relations between belief change and social choice have been reported by Arló-Costa and Pedersen [17] and by Bonanno [42].

12.8 Implementations

Belief revision has usually been studied from the viewpoint of ideal agents (resource bounded or not). The performance of belief change operations in a computer or robot gives rise to new challenges. From a practical point of view we have to deal with actual limitations in memory, reasoning, accuracy, etc. From a theoretical point of view the computational tractability of the proposed algorithms is a major challenge.

In the 1980s several algorithms for the implementation of belief change operations were proposed. Most of them were constructed to determine which beliefs are supported and how, and to perform changes while minimizing the number of (usually atomic) sentences to be changed. Major examples are the algorithms proposed by Doyle, [78], Borgida [53], Winslett [338], Dalal [66] and Satoh [305]. Katsuno and Mendelzon [205] provided an overview of several of these approaches.

Two logics suitable for supporting belief revision systems have been proposed. The first of these was due to Martins and Shapiro in an early paper that also described an actual implementation [249], and the second to Gabbay, Pigozzi, and Woods [125]. Recent works on implementation also include proposals by Williams [336], Williams and Sims [337], and Delgrande and Schaub [71].

A core aspect in implementation is the space and time required for computation. One of the first studies of the computational costs of belief change algorithms was performed by Eiter and Gottlob [85]. In a survey written in 1998 Nebel said:

> The general revision problem for propositional logic appears to be hopelessly infeasible from a computational point of view because [the sources of complexity] are located on the second level of the polynomial hierarchy [257].

However, specific algorithms can reduce this computational cost. The determination of what properties have to be given up in order to gain efficiency is a major challenge for future studies. An interesting step was taken by Jin and Thielscher when they proposed a model called Reinforcement Belief Revision. This model combines two important desiderata for belief change implementations: It satisfies the standard rationality postulates, and the time and space required for its implementation can be assessed [202].

Chapter 13
Multiagent Belief Change

Classical AGM operations model the belief changes of a single agent. They can be extended so that more than one agent is involved. In many situations coherent beliefs based on several sources are needed. However, these sources may contradict each other. Consequently, agents need to combine their beliefs in order to obtain a common, consistent set of beliefs.

The first approaches to multiagent belief change were proposed by Baral, Kraus, Minker, and Venkataram. They combined different belief bases to obtain a consistent combination that satisfies a set of integrity constraints [26, 27]. Later, Revesz [281] and Lin [228] showed how to obtain a common belief set based on distances between models. Liberatore and Schaerf proposed *arbitration* [226, 227], a method for *mutual belief revision* involving two agents. These proposals have been followed by several other approaches to multiagent change [81, 246, 247, 236, 21].

13.1 Merging

Konieczny and Pino Perez have proposed operations of *merging* that generalize several previously proposed methods for conflict-solving through the combination of information from several agents [215]. As in classical belief revision, rationality postulates are proposed. However, as the authors pointed out [214], there is an important difference, namely the social aspect of merging. Postulates are needed on how conflicts between sources of information should be resolved. Consequently, it is possible to distinguish between different families of merging operations, depending on their behaviour towards the sources, for instance majority-respecting behaviour (for a detailed discussion, see [214]).

In order to introduce merging formally we have to introduce some notation:

A belief base A is a finite set of propositional formulae $\{p_1,\ldots,p_n\}$. We write $\bigwedge A$ for the conjunction of the formulae of A, i.e., $\bigwedge A = p_1 \wedge \ldots \wedge p_n$. $\mathcal{A}$ denotes the set of belief bases. Let $A_1,\ldots,A_n$ be n (not necessarily different) belief bases. A *profile* E is a finite and non-empty multiset of belief bases $E = \{A_1,\ldots,A_n\}$ (hence

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_13

different agents are allowed to exhibit identical bases). It represents the beliefs of a group of n agents. $\mathcal{E}$ denotes the set of profiles. $\bigwedge E$ denotes the conjunction of the bases of E, i.e., $\bigwedge E = \bigwedge A_1 \wedge \ldots \wedge \bigwedge A_n$. $\bigvee E$ denotes the disjunction of the bases of E, i.e., $\bigvee E = \bigwedge A_1 \vee \ldots \vee \bigwedge A_n$. A profile E is said to be consistent if and only if $\bigwedge E$ is consistent. $\sqcup$ denotes the union of profiles. Two profiles are equivalent, denoted by $E_1 \equiv E_2$, if there is a bijective function f from E_1 onto E_2 such that for any $A \in E_1$, $\vdash f(A) \leftrightarrow \bigwedge A$.[1]

Merging operations can be constructed as functions from the set of profiles and the set of propositional formulae (which will represent integrity constraints) to the set of belief bases, i.e., $\triangle : \mathcal{E} \times \mathcal{L} \mapsto \mathcal{A}$. We will use the notation $\triangle_p(E)$ instead of $\triangle(E,p)$. $\triangle_p(E)$ is the merged belief base obtained from the profile E under the integrity constraint p (which corresponds to the input sentence in traditional belief change).

The following axioms have been proposed for merging operations [214]:

(IC0) $\triangle_p(E) \vdash p$
(IC1) If p is consistent, then $\triangle_p(E)$ is consistent
(IC2) If $\bigwedge E$ is consistent with p, then $\triangle_p(E) \equiv (\bigwedge E) \wedge p$
(IC3) If $E_1 \equiv E_2$ and $\vdash p_1 \leftrightarrow p_2$, then $\triangle_{p_1}(E_1) \equiv \triangle_{p_2}(E_2)$
(IC4) If $A_1 \vdash p$ and $A_2 \vdash p$, then $\triangle_p(\{A_1,A_2\}) \wedge A_1$ is consistent if and only if $\triangle_p(\{A_1,A_2\}) \wedge A_2$ is consistent
(IC5) $\triangle_p(E_1) \wedge \triangle_p(E_2) \vdash \triangle_p(E_1 \sqcup E_2)$
(IC6) If $\triangle_p(E_1) \wedge \triangle_p(E_2)$ is consistent, then $\triangle_p(E_1 \sqcup E_2) \vdash \triangle_p(E_1) \wedge \triangle_p(E_2)$
(IC7) $\triangle_{p_1}(E) \wedge p_2 \vdash \triangle_{p_1 \wedge p_2}(E)$
(IC8) If $\triangle_{p_1}(E) \wedge p_2$ is consistent, then $\triangle_{p_1 \wedge p_2}(E) \vdash \triangle_{p_1}(E)$

We follow [214] in explaining the postulates. (IC0) to (IC2) correspond to the AGM postulates success, consistency and vacuity. (IC3) corresponds to Katsuno and Mendelzon's postulate (R4) for belief revision (see Section 3.6). (IC4) states that when two belief bases are merged, the merging operation must not give preference to one of them. (IC5) and (IC6) together state that if the merge of E_1 is consistent with the merge of E_2 then the conjunction of these two merges is equivalent with the merge of $E_1 \sqcup E_2$. (IC7) and (IC8) are closely related with the AGM supplementary postulates. They imply that an interpretation that is preferred among the possible interpretations will remain preferred if one restricts the possible choices.

Merging is an extension of revision. If a profile is composed of just a single finite base A, then merging and revision coincide, as we see in the following theorem:

Theorem 13.1. [214] Let A be a finite belief base. Let φ such that $\vdash \varphi \leftrightarrow \bigwedge A$ (i.e., φ represents the conjunction of all the sentences in A). Let $\triangle$ be a merging operation for $\{A\}$. Then $*$ defined as $\varphi * p = \triangle_p(\{A\})$ satisfies the revision postulates (R1) - (R6).

[1] Note that this is different from claiming that E_1 and E_2 are equivalent if and only if for each element of A_i of E_1 there exists an element of A_j of E_2 such that $\vdash \bigwedge A_i \leftrightarrow \bigwedge A_j$ and vice-versa. For example, $E_1 = \{A_1,A_1,A_1,A_2\}$ and $E_2 = \{A_2,A_2,A_2,A_1\}$ satisfy this condition; however, we expect different outcomes from merging operations on these two profiles.

Merging operations can be characterized in terms of possible worlds. We need to define first the notion of a *syncretic assignment*:

Definition 13.1. [214] A *syncretic assignment* is a function mapping each profile E to a total preorder $\leq_E$ ($<_E$, $\simeq_E$) on possible worlds such that for any profiles E, E_1, E_2 and for any belief bases A, A', the following conditions hold:

1. If $\omega_1 \vdash \bigwedge E$ and $\omega_2 \vdash \bigwedge E$, then $\omega_1 \simeq_E \omega_2$.
2. If $\omega_1 \vdash \bigwedge E$ and $\omega_2 \nvdash \bigwedge E$, then $\omega_1 <_E \omega_2$.
3. If $E_1 \equiv E_2$, then $\leq_{E_1} = \leq_{E_2}$
4. For any belief bases A and A' and any possible world ω_1 with $\omega_1 \vdash A$ there is some possible world ω_2 such that $\omega_2 \vdash A'$ and $\omega_2 \leq_{\{A\} \sqcup \{A'\}} \omega_1$.
5. If $\omega_1 \leq_{E_1} \omega_2$ and $\omega_1 \leq_{E_2} \omega_2$, then $\omega_1 \leq_{E_1 \sqcup E_2} \omega_2$.
6. If $\omega_1 <_{E_1} \omega_2$ and $\omega_1 \leq_{E_2} \omega_2$, then $\omega_1 <_{E_1 \sqcup E_2} \omega_2$.

Theorem 13.2. [215] An operation $\triangle$ is an IC merging operation if and only if there exists a syncretic assignment that maps each profile E to a total preorder $\leq_E$ such that $\triangle_p(E) = \bigcup(min(\|p\|, \leq_E))$

The (IC) postulates do not tell us what properties the outcome should have when there is no consensus among the agents. Konieczny and Pino Perez provided additional properties to characterize merging operations with different behaviours in such cases. For instance, a majority merging operation satisfies (IC0)-(IC8) and

(Maj) For each E_1 and E_2 there is some n such that $\triangle(E_1 \sqcup E_2^n) \vdash \triangle(E_2)$.

This postulate says that if a subgroup is repeated sufficiently many times, then the opinion of that subgroup will prevail. For other merging operations and their semantics, see [214].

13.2 Argumentation

Argumentation theory is concerned primarily with the evaluation of claims based on premises in order to reach conclusions. The work by Toulmin [323] published in the 1950s can perhaps be recognized as a point of departure, but full attention to his proposal was delayed until the beginning of the 1980s. The work of Jon Doyle on Truth Maintenance Systems (TMS) [78] could be considered as a starting point for this research, as it attracted attention to the problem of how something comes to be believed. This work can also be regarded as a first effort in the overlapping area where belief revision and argumentation meet. The research on argumentation during the past few decades led to many developments that include both theoretical and practical contributions. The two areas grew along mostly separate lines of inquiry until recently; but, in the last decade there has been a renewed interest in the interplay between belief revision and general forms of reasoning. Argumentation represents a reasoning mechanism that is particularly apt for obtaining the consequences of a potentially inconsistent repository of beliefs, in ways that connect closely with studies of belief change.

Falappa, Kern-Isberner and Simari [90] have argued that belief revision and argumentation theory are complementary approaches. Belief revision describes the way in which an agent is supposed to change her beliefs when new information arrives or she observes changes in the world, whereas argumentation theory deals with strategies agents employ for their own reasoning, or to change the beliefs of other agents, by providing reasons for such changes. By combining the two, the variety and complexity of reasoning processes is better accounted for than if only one of them is used.

In [88], these authors propose a system that uses argumentative structures in the form of explanations for non-prioritized revisions of a belief base B. In this model, an epistemic input is composed of a sentence p and a set of reasons to believe in p, i.e., rules and prerequisites A from which p can be deductively derived. A partial acceptance revision operation is constructed such that A is initially accepted, which leads to the creation of $B \cup A$ as a (possibly inconsistent) intermediate belief base. Then inconsistencies are removed from $B \cup A$, giving rise to a consistent revised belief base $B * A$. The operation $*$ is an operation of external revision in the sense explained in Section 6.3.

Paglieri and Castelfranchi proposed Data-oriented Belief Revision (DBR) as an alternative to AGM [61, 265]. This model combines belief revision with argumentation, following Toulmin's account of argumentation [323]. The application of DBR to argumentation is primarily intended to highlight structural communalities between arguments and belief-supporting networks. The suggestion is that arguments should be studied also as attempts to change the audience's beliefs [266, 267]. The model contains two basic informational categories, data and beliefs, in order to account for the distinction between pieces of information that are simply gathered and stored by the agent (data), and pieces of information that the agent considers to be (possibly up to a certain degree) truthful representations of the state of the world (beliefs). Contrary to beliefs, data are allowed to be contradictory. When a belief is abandoned, this does not entail removal of the corresponding information from the agent's memory, i.e., disbelieving is not forgetting.

13.3 Game Theory

Belief change and game theory are related in several ways. Booth and Meyer [48] investigated equilibria in belief merging. They assume a *Principle of Equilibrium* according to which all agents simultaneously make the appropriate response to what all the other agents do. The removal functions of the individual agents are combined to construct a social belief removal function. The key idea is that the social belief removal function can perform a minimal change in the AGM sense. Two classes of removal functions for agents have been studied: basic and hyperregular removal. The former has been axiomatically characterized by Booth [44].

Zhang studied bargaining situations from another viewpoint that is also related to belief change. Traditionally, a bargaining situation is abstracted as a game whose set of potential outcomes can be represented by $\langle S, d\rangle$ where S represents the set of feasible alternatives and $d \in S$ the disagreement point, i.e., the outcome to be expected if the negotiations break down and no bargain can be reached. A bargaining solution is a function that assigns to $\langle S, d\rangle$ a unique element of S. Bargaining theory explores the relationship between the bargaining situation and the solution. Zhang proposed an axiomatic characterization of bargaining solutions that is based on postulates from AGM and game theory [340].

One of the major ways to solve a game is to employ backwards induction, by which is meant that one reasons backwards in a chain of successive decisions approximately as follows: What will the last player do in the different situations that (s)he might find herself in? Given that, what will the next last player do in the situations that she can find herself in, and given that the third last player, etc., all the way back to the first player. Models of backwards induction in games will have to include representations of conditional reasoning and therefore also belief-contravening suppositions.

Arló-Costa and Bicchieri [14] proposed an operation for hypothetical reasoning in games, based on a proposal by Samet [304], that is intuitively well-behaved and satisfies some of the classical AGM properties. These authors developed models in which the condition that all players are disposed to behave rationally at all nodes is both necessary and sufficient for them to play the backward induction solution in centipede games (i.e., games in which two players take turns choosing either to take a slightly larger share of an increasing pot, or to pass the pot to the other player). This result was obtained without assuming that rationality is commonly known (as it is in [22]) or commonly hypothesized by the players (as it is in [304]).

Chapter 14
Descriptor Revision

This chapter introduces a new approach to belief revision that was recently presented in [182] and has been further developed in [190, 183, 185, 189, 187, 184, 341]. It provides us with a mechanism for belief change that is in important respects more general than previous approaches. Descriptor revision is based on two new formal constructions: A generalized notation for success conditions and the application of choice mechanisms directly to the set of potential outcomes of the belief change operation.

14.1 Belief Descriptors

Standard operations of belief change are defined in terms of their success conditions, such as $p \in K * p$ for revision and $p \notin K - p$ for contraction. These are statements about what is believed in the belief state that the operation results in. We have encountered some operations with other success conditions, such as $\perp \notin K!$ for consolidation, $A \cap (K - A) = \varnothing$ for package contraction, $A \nsubseteq K - A$ for choice contraction, etc. All these are conditions on which sentences should be included in the outcome of the operation.

Descriptor revision makes use of a general notation for success conditions. This notation is based on a metalinguistic belief predicate $\mathfrak{B}$ that is applied to sentences in the object language $\mathcal{L}$ in which beliefs are expressed. For instance, $\mathfrak{B}p$ denotes that p is believed in the belief state in question, and $\mathfrak{B}(r \vee s)$ that $r \vee s$ is believed. We can combine $\mathfrak{B}$-sentences with the usual truth-functional operations. Hence, $\neg\mathfrak{B}p$ denotes that p is not believed (which is different from $\mathfrak{B}\neg p$, which means that $\neg p$ is believed), and $\mathfrak{B}p \vee \mathfrak{B}q$ means that either p or q is believed (which is different from $\mathfrak{B}(p \vee q)$, which denotes that $p \vee q$ is believed). We can also use a set of such sentences as a success condition. All elements of such a set have to be satisfied for the set to be satisfied; for instance $\{\mathfrak{B}p, \neg\mathfrak{B}q\}$ is satisfied if and only if p is believed and q is not.

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7_14

Upper-case Greek letters such as $\Psi, \Xi, \ldots$ will be used to denote non-empty sets consisting of truth-functional combinations of $\mathfrak{B}$-sentences. Obviously, single such sentences can also be presented as sets; we just write, for instance, $\{\neg\mathfrak{B}q\}$ instead of $\neg\mathfrak{B}q$. Sets consisting of truth-functional combinations of $\mathfrak{B}$-sentences will be called (belief) descriptors. The formal definition is as follows:

Definition 14.1. [182] An *atomic belief descriptor* is a sentence $\mathfrak{B}p$ with $p \in \mathcal{L}$. It is *satisfied* by a belief set K if and only if $p \in K$.

A *molecular belief descriptor* is a truth-functional combination of atomic descriptors. Conditions of satisfaction are defined inductively, such that K satisfies $\neg\alpha$ if and only if it does not satisfy α, it satisfies $\alpha \vee \beta$ if and only if it satisfies either α or β, etc.

A *composite belief descriptor* (in short: descriptor; denoted by upper-case Greek letters $\Psi, \Xi, \ldots$) is a non-empty set of molecular descriptors. A belief set K satisfies a composite descriptor Ψ if and only if it satisfies all its elements.

The symbol $\mathfrak{B}$ is not part of the object language, and therefore it cannot be used to express an agent's beliefs about her own beliefs.[1]

With descriptors we can express a wide variety of success conditions, including but not restricted to those that have been discussed in the literature on AGM-style belief change. To mention just one example, the operation of "making up one's mind" aims at either belief or disbelief in a specified sentence p. Its success condition is $\mathfrak{B}p \vee \mathfrak{B}\neg p$ [341].

The following notation is used to express the logical relations among descriptors:

Definition 14.2. [182] Let K be a belief set and let Ψ and Ξ be descriptors.

(1) $K \Vdash \Psi$ means that K satisfies Ψ, and
(2) $\Psi \Vdash \Xi$ means that all belief sets satisfying Ψ also satisfy Ξ.

14.2 Descriptor Revision

Descriptors are so versatile that when we use them, we only need one operation of belief change. That operation is called *descriptor revision* and is denoted by $\circ$. For any belief set K and descriptor Ψ, $K \circ \Psi$ is, intuitively speaking, the outcome of an adjustment of K to make it satisfy Ψ.

We will now turn to the second basic idea of descriptor revision. (The first was the use of descriptors to represent success conditions.) That idea is to perform belief changes by choosing directly among the belief sets that are potential outcomes of the operation. This is an entirely different approach from that applied in the AGM framework.

[1] It is possible to include an autoepistemic belief predicate into the language. It may or may not coincide with $\mathfrak{B}$, depending on whether or not the agent's autoepistemic beliefs accord with her epistemic conduct. See [190].

In the original AGM model, belief change is based on a choice among the elements of $K \perp p$, the set of inclusion-maximal subsets of K that do not imply p. In the sphere-based model, it is based on a selection among possible worlds. Possible worlds are very large entities, far beyond human cognition, and so are the elements of the remainder set.[2] It is cognitively unrealistic to assume that we perform our belief changes by choosing among an infinite collection of remainders or possible worlds, and then forming the intersection of those that we have selected.

In descriptor revision we choose directly among the belief sets that are potential outcomes of the operation. The framework contains an *outcome set* or set of potential outcomes, denoted by $\mathbb{X}$. Intuitively speaking, $\mathbb{X}$ consists of all those belief sets that a belief change can result in [179]. We can interpret it as consisting of all those belief sets that are sufficiently plausible, sufficiently close at hand, or sufficiently coherent to be the outcome of a belief change. When we perform the belief change $K \circ \Psi$, we employ a choice function that picks out one among those elements of $\mathbb{X}$ that satisfy Ψ. If none of the elements of $\mathbb{X}$ satisfies Ψ, then we let $K \circ \Psi = K$.

Descriptor revision in this general, choice-based form has been axiomatically characterized with the following four postulates [190]:

$K \circ \Psi = \mathrm{Cn}(K \circ \Psi)$ (closure)
$K \circ \Psi \Vdash \Psi$ or $K \circ \Psi = K$ (relative success)
If $K \circ \Xi \Vdash \Psi$ then $K \circ \Psi \Vdash \Psi$. (regularity)
If $K \circ \Xi \Vdash \Psi$ if and only if $K \circ \Xi \Vdash \Psi'$ for all Ξ, then $K \circ \Psi = K \circ \Psi'$. (uniformity)

The closure property comes with our use of belief sets as (idealized) representations of belief states. Both relative success and regularity can be seen as weaker and more reasonable versions of the implausible success condition $K \circ \Psi \Vdash \Psi$. Uniformity says essentially that if two success conditions are satisfied by exactly the same elements of the outcome set, then the changes that have these two conditions as inputs yield the same outcome.

Several more specified forms of descriptor revision have been developed. Arguably the most important of these is *centrolinear revision* [190]. It is based on a relation $\leqq$ on the outcome set $\mathbb{X}$, such that $K \leqq X$ for all $X \in \mathbb{X}$. For all descriptors Ψ we construct $K \circ \Psi$ as the unique $\leqq$-minimal element of $\mathbb{X}$ that satisfies Ψ, unless Ψ is unsatisfiable within $\mathbb{X}$, in which case $K \circ \Psi = K$. (The construction only works if $\leqq$ satisfies certain conditions; see [190] for details.) In order to characterize centrolinear revision axiomatically we can use the above four postulates that hold for all choice-based descriptor revisions, and in addition the following two:[3]

If $K \Vdash \Psi$ then $K \circ \Psi = K$. (confirmation)
If $K \circ \Psi \Vdash \Xi$ then $K \circ \Psi = K \circ (\Psi \cup \Xi)$. (cumulativity)

---

[2] If $\mathcal{L}$ is logically infinite (contains infinitely many logically non-equivalent sentences) and $K \neq \mathrm{Cn}(\varnothing)$, then $K \perp p$ has an infinite number of elements, and each of these elements is logically infinite. This applies even if there is a sentence q such that $K = \mathrm{Cn}(\{q\})$ [169]. See Section 5.5.

[3] However, uniformity is not needed for the characterization since it follows from relative success, regularity and cumulativity [190].

We can think of centrolinear revision as an operation whose outcome is as close to the original belief set as the success condition allows it to be. This can be illustrated in a spatial model by positioning the elements of $\mathbb{X}$, i.e., the potential belief sets, at different distances from the original belief set K, for instance around K on a surface or on a straight line with K at one end.

Alternatively, centrolinear revision can be based on a relation of *epistemic proximity* on descriptors that has very much the same function in descriptor revision as epistemic entrenchment has in AGM belief change [183, 190].

14.3 Connections with AGM

Standard revision by a sentence (we will now call it sentential revision) can easily be constructed as a special case of descriptor revision. If we have a descriptor revision $\circ$ on a belief set K, then we can construct a sentential revision $*$ as a special case, such that for all p:

$$K * p = K \circ \mathfrak{B}p$$

This operation has been axiomatically characterized, and (unsurprisingly) it does not in general satisfy the basic AGM postulates. However, it has partial meet revision as a special case. Perhaps more importantly, all transitively relational partial meet revisions can be constructed in this way from some centrolinear descriptor revision. In other words, full-blown AGM revision (satisfying all eight AGM postulates) is a special case of centrolinear descriptor revision [185, 190].

Contraction, the other major AGM operation, has a more complex relation to descriptor revision. An obvious first attempt to obtain contraction by a sentence p from descriptor revision would be to define $K - p$ as $K \circ \neg\mathfrak{B}p$. Such an operation will not in general satisfy the inclusion postulate for contraction ($K - p \subseteq K$). This may not necessarily be a disadvantage, since the inclusion postulate is contestable. When we remove a sentence p from the belief set, this is usually because we have received some new information that made us doubt it, and since this new information was accepted it should be added to the belief set. Based on that argument, we have reasons to investigate operations that remove a sentence p in a process of minimal change that may very well include the addition of some new information to the belief set. Such an operation is called a *revocation*. It can be derived from descriptor revision with the formula $K - p = K \circ \neg\mathfrak{B}p$ [182, 190].

There are several ways to construct "pure" contraction, i.e., contraction satisfying the inclusion postulate, in the framework of descriptor revision. The most straightforward of these is to extend the success condition so that it also requires inclusion to be satisfied:

$$K - p = K \circ (\{\neg\mathfrak{B}p\} \cup \{\neg\mathfrak{B}q \mid q \notin K\})$$

Several constructions of both revocation and contraction have been investigated in some detail [190, 175, 176, 178, 180, 177]. Interestingly, AGM contraction differs from AGM revision in not being subsumable under descriptor revision. It has been shown that a partial meet contraction can be reconstructed as a centrolinear revision if and only if it is also a transitively relational maxichoice contraction [184]. The latter type of operation is highly problematic. For reasons discussed in Section 3.8, maxichoice contraction is usually considered to be an utterly unrealistic limiting case.

14.4 Further Developments

The descriptor operations described above are all local, i.e., defined only for a single belief set K. However, descriptor revision can easily be extended to global (iterated) belief change [189, 190]. This means that the operation $\circ$ should not be limited to a specific belief set K but instead be applicable to all belief sets that the agent may potentially entertain at some point in time. From a formal point of view, this means that $\circ$ should take us from pairs of a belief set and a descriptor to a new belief set (instead of taking us from a descriptor to a new belief set). One way to construct global descriptor revision is to base it on distances. Then, for any belief set K and descriptor Ψ, $K \circ \Psi$ is the closest Ψ-satisfying belief set within reach from K (unless there is no such belief set, in which case $K \circ \Psi = K$). However, it can be argued that these distances should by asymmetric, or in other words that the distance from a belief set K_1 to another belief set K_2 need not be the same as the distance from K_2 to K_1. Some belief changes are made much more easily in one direction than the other.

Example 14.1. I was once in a belief state in which I did not have any opinion on whether 361 is a prime number or not. It was very easy to bring me to a belief state in which I believe it not to be a prime number. (It was sufficient to convince me that $361 = 19 \times 19$.) However, it would be far from easy to bring me back to a state with no belief in the matter [190].

In the descriptor framework, a generalization of Ramsey test conditionals is easily obtainable:

Definition 14.3. [189] A *Ramsey descriptor* is a sentence $\Psi \Rightarrow \Xi$, where Ψ and Ξ are descriptors. For a given operation $\circ$ of descriptor revision, $\Rightarrow$ has the following condition of validity:

$\Psi \Rightarrow \Xi$ holds *if and only if* $K \circ \Psi \Vdash \Xi$.

In other words, $\Psi \Rightarrow \Xi$ means that if the belief set is revised by Ψ, then the outcome will satisfy Ξ. The Gärdenfors impossiblity theorem (Section 11.3) that prevents us from inserting standard Ramsey test conditionals into belief sets cannot be generalized to the descriptor framework. Therefore, we can unproblematically extend the belief sets of descriptor revision so that they contain Ramsey descriptors, i.e., sentences of the form $\Psi \Rightarrow \Xi$.

Standard Ramsey test conditionals (that we denote with the symbol $\Box\!\!\rightarrow$) can be obtained with the following recipe:

$p \Box\!\!\rightarrow q$ if and only if $\mathfrak{B}p \Rightarrow \mathfrak{B}q$.

Ramsey descriptors can also be used to express other forms of conditionality between belief patterns [189, 188, 186, 190]. Consider for instance the following examples:

Example 14.2. "If he gives up his belief that his wife is faithful to him, then he will also lose his belief that she loves him." ($\neg\mathfrak{B}p \Rightarrow \neg\mathfrak{B}q$)

"If she gives up her belief that the first chapter of Genesis is literally true, then she will still believe that God exists." ($\neg\mathfrak{B}p \Rightarrow \mathfrak{B}q$)

"If she makes up her mind on whether this painting is a genuine Picasso or not, then she will come to believe that it is genuine" ($\mathfrak{B}p \vee \mathfrak{B}\neg p \Rightarrow \mathfrak{B}p$) [190].

These are just a few examples of the many patterns of change that can be uncovered in human epistemic behaviour and studied in a precise manner with new and more sophisticated formal tools. In spite of several decades of research on the logic of belief change, the area is still replete with unanswered questions and unexplored issues. We may still only be at the beginning of its development.

References

1. Alchourrón, C.: Philosophical foundations of deontic logic and the logic of defeasible conditionals. In: J.J. Meyer, R.J. Wieringa (eds.) Deontic Logics in Computer Science: Normative System Specification, pp. 43–84. Chicester, Wiley & Sons (1993)
2. Alchourrón, C.: Defeasible logic: Demarcation and affinities. In: Crocco, Fariña del Cerro (eds.) Conditionals: from Philosophy to Computer Science, pp. 67–102. Clarendon Press, Oxford (1995)
3. Alchourrón, C.: Detachment and defeasibility in deontic logic. Studia Logica **51**, 5–18 (1996)
4. Alchourrón, C., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. Journal of Symbolic Logic **50**, 510–530 (1985)
5. Alchourrón, C., Makinson, D.: Hierarchies of regulations and their logic. In: R. Hilpinen (ed.) New Studies in Deontic Logic: Norms, Actions, and the Foundations of Ethics, pp. 125–148 (1981)
6. Alchourrón, C., Makinson, D.: On the logic of theory change: Contraction functions and their associated revision functions. Theoria **48**, 14–37 (1982)
7. Alchourrón, C., Makinson, D.: On the logic of theory change: Safe contraction. Studia Logica **44**, 405–422 (1985)
8. Alchourrón, C., Makinson, D.: Maps between some differents kinds of contraction functions: The finite case. Studia Logica **45**, 187–198 (1986)
9. Alechina, N., Jago, M., Logan, B.: Resource-bounded belief revision and contraction. In: M. Baldoni, U. Endriss, A. Omicini, P. Torroni (eds.) Declarative Agent Languages and Technologies III, Third International Workshop, DALT 2005, Utrecht, The Netherlands, July 25, 2005, Selected and Revised Papers, *Lecture Notes in Computer Science*, vol. 3904, pp. 141–154. Springer (2006)
10. Areces, C., Becher, V.: Iterable AGM functions. In: H. Rott, M. Williams (eds.) Frontiers in Belief Revision, Applied Logic Series, pp. 261–277. Kluwer Academic Publishers (2001)
11. Areces, C., Becher, V., Fermé, E., Rodríguez, R.: Observaciones a la teoría AGM. In: Proceedings Primer Encuentro en Temas de Lógica no Standard. Vaquerías - Córdoba (1996)
12. Arló-Costa, H.: Epistemic conditionals, snakes and stars. In: Crocco, Fariña del Cerro (eds.) Conditionals: from Philosophy to Computer Science, pp. 193–239. Clarendon Press, Oxford (1995)
13. Arló-Costa, H.: Bayesian epistemology and epistemic conditionals: On the status of the export-import laws. Journal of Philosophy **98: 11**, 555–598 (2001)
14. Arló-Costa, H., Bicchieri, C.: Knowing and supposing in games of perfect information. Studia Logica **86**(3), 353–373 (2007)
15. Arló-Costa, H., Levi, I.: Two notions of epistemic validity. Synthese **109:2**, 217–262 (1996)
16. Arló-Costa, H., Parikh, R.: Conditional probability and defeasible inference. Journal of Philosophical Logic **34**, 97–119 (2005)

E. Fermé and S. O. Hansson, *Belief Change*, SpringerBriefs in Intelligent Systems, https://doi.org/10.1007/978-3-319-60535-7

17. Arló-Costa, H., Pedersen, A.P.: Social norms, rational choice and belief change. In: E.J. Olsson, S. Enqvist (eds.) Belief Revision meets Philosophy of Science, *Logic, Epistemology, and the Unity of Science*, vol. 21, pp. 163–212. Springer Netherlands (2011)
18. Arrow, K.J.: Social Choice and Individual Values. Wiley (1951)
19. Aucher, G.: A combined system for update logic and belief revision. Master's thesis, ILLC, University of Amsterdam, Amsterdam, the Netherlands (2003). ILLC report MoL-2003-03
20. Aucher, G.: Interpreting an action from what we perceive and what we expect. Journal of Applied Non-Classical Logics **17**(1), 9–38 (2007)
21. Aucher, G.: Generalizing AGM to a multi-agent setting. Logic Journal of the IGPL **18**(1), 530–559 (2010)
22. Aumann, R.J.: Irrationality in game theory. In: P. Dasgupta, D. Gale, O. Hart, E. Maskin (eds.) Economic Analysis of Markets and Games (Essays in Honor of Frank Hahn), pp. 214–227. MIT Press (1992)
23. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P.: The Description Logic Handbook: Theory, Implementation, and Applications. Cambridge University Press, New York, NY, USA (2003)
24. Baltag, A., Moss, L., Solecki, S.: The logic of public announcements, common knowledge, and private suspicions. In: I. Gilboa (ed.) Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 98), pp. 43–56 (1998)
25. Baltag, A., Smets, S.: Dynamic belief revision over multi-agent plausibility models. In: W. van der Hoek, M. Wooldridge (eds.) Proceedings of LOFT 2006 (7th Conference on Logic and the Foundations of Game and Decision Theory), pp. 11–24 (2006)
26. Baral, C., Kraus, S., Minker, J.: Combining multiple knowledge bases. IEEE Transactions on Knowledge and Data Engineering **3**(2), 208–220 (1991)
27. Baral, C., Kraus, S., Minker, J., Subrahmanian, V.S.: Combining knowledge bases consisting of first-order theories. Computational intelligence **8**(1), 45–71 (1992)
28. Becher, V.: Binary functions for theory change. Ph.D. thesis, University of Buenos Aires (1999)
29. Becher, V., Fermé, E., Lazzer, S., Oller, C., Palau, G., Rodríguez, R.: Some observations on Carlos Alchourrón's theory of defeasible conditionals. In: P. McNamara, H. Prakken (eds.) Norms, Logics and Information Systems, New studies on Deontic Logic and Computer Science, pp. 219–230. IOS Press, Amsterdam/Tokyo/Washington DC (1999)
30. Belnap, N.: Rescher´s hypothetical reasoning: An amendment. In: E. Sosa (ed.) The Philosophy of Nicholas Rescher: Discussion and Replies, pp. 19–28. D. Reidel, Dordrecht (1979)
31. Benferhat, S., Kaci, S., Le Berre, D., Williams, M.A.: Weakening conflicting information for iterated revision and knowledge integration. Artificial Intelligence **153**(1-2), 339–371 (2004)
32. van Benthem, J.: Semantic parallels in natural language and computation. In: H.D. Ebbinghaus, J. Fernandez-Prida, M. Garrido, D. Lascar, M.R. Artalejo (eds.) Logic Colloquium '87, pp. 331–375. North-Holland, Amsterdam (1989)
33. van Benthem, J.: Logic and the flow of information. In: D. Prawitz, B. Skyrms, D. Westerstahl (eds.) Proceedings of the 9th International Congress of Logic, Methodology and Philosophy of Science (1991). Elsevier Science B.V. (1994)
34. van Benthem, J.: Dynamic logic of belief revision. Journal of Applied Non-Classical Logics **17**(2), 129–155 (2007)
35. Billington, D., Antoniou, G., Governatori, G., Maher, M.: Revising nonmonotonic theories: The case of defeasible logic. In: KI-99: Advances in Artificial Intelligence, pp. 101–112. Springer (1999)
36. Bochman, A.: A Logical Theory of Nonmonotonic Inference and Belief Change. Springer (2001)
37. Boella, G., da Costa Pereira, C., Pigozzi, G., Tettamanzi, A., van der Torre, L.: The role of beliefs in goal dynamics: prolegomena to a constructive theory of intentions. Logic Journal of the IGPL **4**(18), 559–578 (2010)
38. Boella, G., Pigozzi, G., van der Torre, L.: Normative framework for normative system change. In: Decker, Sichman, Sierra, Castelfranchi (eds.) Proceedings of 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009), pp. 169–176. Budapest (2009)

39. Bonanno, G.: A simple modal logic for belief revision. Synthese **147**, 193–228 (2005)
40. Bonanno, G.: Axiomatic characterization of the AGM theory of belief revision in a temporal logic. Artificial Intelligence **171**, 144–160 (2007)
41. Bonanno, G.: Temporal interaction of information and belief. Studia Logica **86**, 381–407 (2007)
42. Bonanno, G.: Rational choice and AGM belief revision. Artificial Intelligence **173**, 1194–1203 (2009)
43. Bonanno, G.: Revealed preference, iterated belief revision and dynamic games. In: G. Bonanno, J. Delgrande, H. Rott (eds.) Information processing, rational belief change and social interaction, no. 09351 in Dagstuhl Seminar Proceedings. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany, Dagstuhl, Germany (2009). URL `http://drops.dagstuhl.de/opus/volltexte/2009/2232`
44. Booth, R., Chopra, S., Meyer, T., Ghose, A.: Double preference relations for generalised belief change. Artificial Intelligence **174:16-17**, 1339–1368 (2010)
45. Booth, R., Fermé, E., Konieczny, S., Pino Pérez, R.: Credibility-limited revision operators in propositional logic. In: Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR 2012), pp. 116–125 (2012)
46. Booth, R., Fermé, E., Konieczny, S., Pino Pérez, R.: Credibility-limited improvement operators. In: ECAI 2014: 21st European Conference on Artificial Intelligence, vol. 263, pp. 123–128. IOS Press (2014)
47. Booth, R., Meyer, T.: Admissible and restrained revision. Journal of Artificial Intelligence Research **26**, 127–151 (2006)
48. Booth, R., Meyer, T.: Equilibria in social belief removal. Synthese **26**, 1–27 (2010)
49. Booth, R., Meyer, T., Varzinczak, I.J.: Next steps in propositional Horn contraction. In: Proceedings of the International Joint Conference on Artificial Intelligence, pp. 702–707 (2009)
50. Booth, R., Meyer, T., Varzinczak, I.J., Wassermann, R.: A contraction core for Horn belief change: Preliminary report. In: Proceedings of the 13th International Workshop on Non-Monotonic Reasoning (NMR). Toronto (2010)
51. Booth, R., Meyer, T., Wong, K.S.: A bad day surfing is better than a good day working: How to revise a total preorder. In: Proceedings of KR2006, Tenth International Conference on the Principles of Knowledge Representation and Reasoning, pp. 230–238. AAAI Press (2006)
52. Booth, R., Richter, E.: On revising fuzzy belief bases. Studia Logica **68**, 1–30 (2001)
53. Borgida, A.: Language features for flexible handling of exceptions in information systems. ACM Transactions on Database Systems **10**(4), 565–603 (1985)
54. Boutilier, C.: Revision sequences and nested conditionals. In: Proceedings of the 13th International Joint Conference on Artificial Intelligence (IJCAI'93), pp. 519–525 (1993)
55. Boutilier, C.: Iterated revision and minimal change of conditional beliefs. Journal of Philosophical Logic **25**, 263–305 (1996)
56. Brewka, G.: Belief revision in a framework for default reasoning. In: A. Fuhrmann, M. Morreau (eds.) The Logic of Theory Change: Workshop, Konstanz, FRG, October 13–15, 1989 Proceedings, *Lecture Notes in Artificial Intelligence*, vol. 465, pp. 206–222. Springer (1991)
57. Cantwell, J.: On the logic of small changes in hypertheories. Theoria **63**, 54–89 (1997)
58. Cantwell, J.: Some logics of iterated belief change. Studia Logica **63**(1), 49–84 (1999)
59. Carnota, R., Rodríguez, R.: AGM theory and artificial intelligence. In: E.J. Olsson, S. Enqvist (eds.) Belief Revision meets Philosophy of Science, *Logic, Epistemology, and the Unity of Science*, vol. 21, pp. 1–42. Springer Netherlands (2011)
60. Castelfranchi, C.: Reasons: Belief support and goal dynamics. Mathware & Soft Computing **3**, 233–247 (1996)
61. Castelfranchi, C.: Representation and integration of multiple knowledge sources: issues and questions. In: Cantoni, D. Gesú, Setti, Tegolo (eds.) Human & Machine Perception: Information Fusion, pp. 235–254. Plenum Press (1997)
62. Castelfranchi, C., D'Aloisi, D., Giacomelli, F.: A framework for dealing with belief-goal dynamics. In: M. Gori, G. Soda (eds.) Topics in Artificial Intelligence: 4th Congress of the Italian Association for Artificial Intelligence AI*IA '95 Florence, Italy, October 11–13, 1995 Proceedings, pp. 237–242. Springer (1995)

63. Castelfranchi, C., Paglieri, F.: The role of beliefs in goal dynamics: prolegomena to a constructive theory of intentions. Synthese **155**(2), 237–263 (2007)
64. Coniglio, M., Carnielli, W.: Transfers between logics and their applications. Studia Logica **72**, 367–400 (2002)
65. Cross, C., Thomason, R.: Conditionals and knowledge-base update. In: P. Gärdenfors (ed.) Belief Revision, no. 29 in Cambridge Tracts in Theoretical Computer Science, pp. 247–275. Cambridge University Press (1992)
66. Dalal, M.: Investigations into a theory of knowledge base revision: Preliminary report. In: Seventh National Conference on Artificial Intelligence, (AAAI-88), pp. 475–479. St. Paul (1988)
67. Darwiche, A., Pearl, J.: On the logic of iterated belief revision. Artificial Intelligence **89**(1–2), 1–29 (1997)
68. Del Val, A.: Non monotonic reasoning and belief revision: Syntactic, semantic, foundational, and coherence approaches. Journal of Applied Non-Classical Logics **7**, 213–240 (1997)
69. Delgrande, J.: Horn clause belief change: Contraction functions. In: G. Brewka, J. Lang (eds.) Proceedings of the Eleventh International Conference on the Principles of Knowledge Representation and Reasoning, pp. 156–165. AAAI Press, Sydney (2008)
70. Delgrande, J., Dubois, D., Lang, J.: Iterated revision as prioritized merging. In: P. Doherty, J. Mylopoulos, C. Welty (eds.) Tenth International Conference on Principles of Knowledge Representation and Reasoning, pp. 210–220. The AAAI Press/The MIT Press, Lake District, UK (2006)
71. Delgrande, J., Schaub, T.: A consistency-based approach for belief change. Artificial Intelligence **151**(1-2), 1–41 (2003)
72. Delgrande, J., Wassermann, R.: Horn clause contraction functions: Belief set and belief base approaches. In: International Conference on the Principles of Knowledge Representation and Reasoning, pp. 143–152. Toronto (2010)
73. Di Giusto, P., Governatori, G.: A new approach to base revision. In: P. Barahona, J.J. Alferes (eds.) Progress in Artificial Intelligence, pp. 327–341. Springer-Verlag (1999)
74. van Ditmarsch, H.: Prolegomena to dynamic logic for belief revision. Synthese (Knowledge, Rationality & Action) **147**, 229–275 (2005)
75. van Ditmarsch, H., van der Hoek, W., Kooi, B.: Dynamic Epistemic Logic, *Synthese Library*, vol. 337. Springer (2007)
76. van Ditmarsch, H., Kooi, B.: The secret of my success. Synthese **151**, 201–232 (2006)
77. van Ditmarsch, H., Labuschagne, W.: My beliefs about your beliefs: A case study in theory of mind and epistemic logic. Knowledge, Rationality & Action (Synthese) **155**, 191–209 (2007)
78. Doyle, J.: A truth maintenance system. Artificial Intelligence **12**, 231–272 (1979)
79. Doyle, J.: Reason maintenance and belief revision: Foundations versus coherence theories. In: P. Gärdenfors (ed.) Belief Revision, no. 29 in Cambridge Tracts in Theoretical Computer Science, pp. 29–51. Cambridge University Press (1992)
80. Doyle, J.: Mechanics and mental change. In: B.O. Küppers, U. Hahn, S. Artmann (eds.) Evolution of Semantic Systems, pp. 127–150. Springer (2013)
81. Dragoni, A.F.: A model for belief revision in a multi-agent environment. In: Y. Demazeau, E. Werner (eds.) Decentralized A.I., pp. 103–112. Elsevier Science Publisher (1992)
82. Dubois, D., Prade, H.: Epistemic entrenchment and possibilistic logic. Artificial Intelligence **50**(2), 223–239 (1991)
83. Dubois, D., Prade, H.: Belief change and possibilistic logic. In: P. Gärdenfors (ed.) Belief Revision, no. 29 in Cambridge Tracts in Theoretical Computer Science, pp. 142–182. Cambridge University Press (1992)
84. Dupin de Saint-Cyr, F., Lang, J.: Belief extrapolation (or how to reason about observations and unpredicted change). In: D. Fensel, F. Giunchiglia, D. McGuinness, M. Williams (eds.) Proceedings of the 8th International Conference on Principles of Knowledge Representation and Reasoning (KR2002), pp. 497–508. Morgan Kaufmann (2002)
85. Eiter, T., Gottlob, G.: On the complexity of propositional knowledge base revision, updates, and counterfactuals. Artificial Intelligence **57**, 227–270 (1992)

86. Fagin, R., Ullman, J., Vardi, M.: On the semantics of updates in databases: Preliminary report. In: Proceedings of Second ACM SIGACT-SIGMOD Symposium on Principles of Database Systems, pp. 352–365 (1983)
87. Falappa, M., Fermé, E., Kern-Isberner, G.: On the logic of theory change: Relations between incision and selection functions. In: G. Brewka, S. Coradeschi, A. Perini, P. Traverso (eds.) Proceedings 17th European Conference on Artificial Intelligence, ECAI06, pp. 402–406 (2006)
88. Falappa, M., Kern-Isberner, G., Simari, G.R.: Explanations, belief revision and defeasible reasoning. Artificial Intelligence **141**, 1–28 (2002)
89. Falappa, M.A., García, A.J., Kern-Isberner, G., Simari, G.R.: Stratified belief bases revision with argumentative inference. Journal of Philosophical Logic **42**(1), 161–193 (2011)
90. Falappa, M.A., Kern-Isberner, G., Simari, G.R.: Belief revision and argumentation theory. In: G. Simari, I. Rahwan (eds.) Argumentation in Artificial Intelligence, pp. 341–360. Springer US (2009)
91. Fermé, E., Krevneris, M., Reis, M.D.L.: An axiomatic characterization of ensconcement-based contraction. Journal of Logic and Computation **18**(5), 739–753 (2008)
92. Fermé, E., Reis, M.D.L.: System of spheres-based multiple contractions. Journal of Philosophical Logic **41**(1), 29–52 (2012)
93. Fermé, E., Reis, M.D.L.: Epistemic entrenchment-based multiple contractions. The Review of Symbolic Logic **6**, 460–487 (2013)
94. Fermé, E.: Actualización de bases de conocimiento usando teorías de cambio de creencia. In: 3rd Ibero-American Conference on Artificial Intelligence 1992, pp. 419–436 (1992)
95. Fermé, E.: On the logic of theory change: Contraction without recovery. Journal of Logic, Language and Information **7**, 127–137 (1998)
96. Fermé, E.: A little note about maxichoice and epistemic entrenchment. In: R. de Queiroz, W. Carnielli (eds.) Proceedings Workshop on Logic, Language, Information and Computation Wollic´99, Itatiaia, Brasil, pp. 111–114 (1999)
97. Fermé, E.: Revising the AGM postulates. Ph.D. thesis, University of Buenos Aires (1999)
98. Fermé, E.: Irrevocable belief revision and epistemic entrenchment. Logic Journal of the IGPL **8**(5), 645–652 (2000)
99. Fermé, E.: Five faces of recovery. In: H. Rott, M.A. Williams (eds.) Frontiers in Belief Revision, Applied Logic Series, pp. 247–259. Kluwer Academic Publishers (2001)
100. Fermé, E., Garapa, M., Reis, M.D.: On ensconcement and contraction. Journal of Logic and Computation p. exx008 (2017)
101. Fermé, E., Hansson, S.O.: Selective revision. Studia Logica **63:3**, 331–342 (1999)
102. Fermé, E., Hansson, S.O.: Shielded contraction. In: H. Rott, M. Williams (eds.) Frontiers in Belief Revision, Applied Logic Series, pp. 85–107. Kluwer Academic Publishers (2001)
103. Fermé, E., Mikalef, J., Taboada, J.: Credibility-limited functions for belief bases. Journal of Logic and Computation **13:1**, 99–110 (2003)
104. Fermé, E., Rodriguez, R.: A brief note about the Rott contraction. Logic Journal of the IGPL **6**(6), 835–842 (1998)
105. Fermé, E., Rodríguez, R.: Semi-contraction: Axioms and construction. Notre Dame Journal of Formal Logic **39**(3), 332–345 (1998)
106. Fermé, E., Rodríguez, R.: DFT and belief revision. Análisis Filosófico **27**(2), 373–393 (2006)
107. Fermé, E., Rott, H.: Revision by comparison. Artificial Intelligence **157**, 5–47 (2004)
108. Fermé, E., Saez, K., Sanz, P.: Multiple kernel contraction. Studia Logica **73**, 183–195 (2003)
109. Foo, N.: Observation on AGM entrenchment. Tech. rep., University of Sydney (1990). Computer Science Technical Report 389.
110. Fotinopoulos, A., Papadopoulos, V.: Semantics for Horn contraction. In: 7th PanHellenic Logic Symposium, pp. 42–47 (2009)
111. Friedman, N., Halpern, J.Y.: Modeling belief in dynamic systems. Part I: Foundations. Artificial Intelligence **95:2,**, 257–316 (1997)
112. Friedman, N., Halpern, J.Y.: Belief revision: A critique. Journal of Logic, Language, and Information **8**, 401–420 (1999)

113. Friedman, N., Halpern, J.Y.: Modeling belief in dynamic systems. Part II: Revision and update. Journal of AI Research **10**, 117–167 (1999)
114. Fuhrmann, A.: Relevant logic, modal logic and theory change. Ph.D. thesis, Department of Philosophy and Automated Reasoning Project, Institute of Advanced Studies, Australian National University, Canberra (1988)
115. Fuhrmann, A.: Reflective modalities and theory change. Synthese **81**(1), 115–134 (1989)
116. Fuhrmann, A.: Theory contraction through base contraction. Journal of Philosophical Logic **20**, 175–203 (1991)
117. Fuhrmann, A.: Everything in flux: Dynamic ontologies. In: S. Lindström, R. Sliwinski, K. Segerberg (eds.) Odds and Ends: Philosophical Essays Dedicated to Wlodek Rabinowicz on the Occasion of his Fiftieth Birthday, pp. 111–125. Uppsala Philosophical Studies 45 (1996)
118. Fuhrmann, A.: An Essay on Contraction. Studies in Logic, Language and Information. CSLI Publications, Stanford (1997)
119. Fuhrmann, A.: Solid belief. Theoria **63**, 90–104 (1997)
120. Fuhrmann, A.: Travels in ontological space. In: J. Nida-Rümelin, G. Meggle (eds.) Analyomen 2, Volume I: Logic, Epistemology, Philosophy of Science, pp. 68–77. De Gruyter (1997)
121. Fuhrmann, A., Hansson, S.O.: A survey of multiple contraction. Journal of Logic, Language and Information **3**, 39–74 (1994)
122. Gabbay, D., Hunter, A.: Making inconsistency respectable. In: P. Jorrand, J. Kelemen (eds.) Fundamentals of Artificial Intelligence Research (FAIR '91), *Lecture Notes in Artificial Intelligence*, vol. 535, pp. 19–32. Springer (1991)
123. Gabbay, D., Rodrigues, O., Russo, A.: Revision by translation. In: B. Bouchon-Meunier, R.R. Yager, L.A. Zadeh (eds.) Information, Uncertainty, Fusion, pp. 3–31. Kluwer Academic Publishers (1999)
124. Gabbay, D., Rodrigues, O., Russo, A.: Belief revision in non-classical logics. The Review of Symbolic Logic **1**(03), 267–304 (2008)
125. Gabbay, D.M., Pigozzi, G., Woods, J.: Controlled revision: An algorithmic approach for belief revision. Journal of Logic and Computation **13**(1), 3–22 (2003)
126. Gallier, J.R.: Autonomous belief revision and communication. In: P. Gärdenfors (ed.) Belief Revision, no. 29 in Cambridge Tracts in Theoretical Computer Science, pp. 220–246. Cambridge University Press (1992)
127. Garapa, M., Fermé, E., Reis, M.D.L.: Studies on brutal contraction and severe withdrawal. Studia Logica **105**, 1–30 (2016)
128. Gärdenfors, P.: Conditionals and changes of belief. Acta Philosophica Fennica **30**, 381–404 (1978)
129. Gärdenfors, P.: Conditionals and changes of belief. American Philosophical Quarterly **18**(3), 203–211 (1978)
130. Gärdenfors, P.: Rules for rational changes of belief. In: T. Pauli (ed.) Philosophical Essays Dedicated to Lennart Åqvist on his Fiftieth Birthday, no. 34 in Philosophical Studies, pp. 88–101 (1982)
131. Gärdenfors, P.: Epistemic importance and minimal changes of belief. Australasian Journal of Philosophy **62**, 136–157 (1984)
132. Gärdenfors, P.: Belief revisions and the Ramsey test for conditionals. Philosophical Review **95**, 81–93 (1986)
133. Gärdenfors, P.: Knowledge in Flux: Modeling the Dynamics of Epistemic States. The MIT Press, Cambridge (1988)
134. Gärdenfors, P.: The dynamics of belief systems: Foundations versus coherence theories. Revue Internationale de Philosophie **44**, 24–46 (1990)
135. Gärdenfors, P.: Belief revision and nonmonotonic logic: Two sides of the same coin? In: J. van Eijck (ed.) Logics in Artificial Intelligence European Workshop JELIA '90 Amsterdam, The Netherlands, September 10–14, 1990, *Lecture Notes in Computer Science*, vol. 478, pp. 52–54. Springer (1991)

136. Gärdenfors, P.: Belief revision: An introduction. In: P. Gärdenfors (ed.) Belief Revision, no. 29 in Cambridge Tracts in Theoretical Computer Science, pp. 1–28. Cambridge University Press (1992)
137. Gärdenfors, P.: Notes on the history of ideas behind AGM. Journal of Philosophical Logic **40**(2), 115–120 (2011)
138. Gärdenfors, P., Makinson, D.: Revisions of knowledge systems using epistemic entrenchment. In: M.Y. Vardi (ed.) Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge, pp. 83–95. Morgan Kaufmann, Los Altos (1988)
139. Gärdenfors, P., Rott, H.: Belief revision. In: D.M. Gabbay, C.J. Hogger, J.A. Robinson (eds.) Handbook of Logic in Artificial Intelligence and Logic Programming, vol. 3, Epistemic and Temporal Reasoning, pp. 35–132. Oxford University Press (1993)
140. Genin, K., Kelly, K.T.: Learning, theory choice, and belief revision (2016). Studia Logica. In press.
141. Glaister, S.M.: Recovery recovered. Journal of Philosophical Logic **29**, 171–206 (2000)
142. Governatori, G., Rotolo, A.: Changing legal systems: legal abrogations and annulments in defeasible logic. Logic Journal of the IGPL **18: (1)**, 157–194 (2010)
143. Grove, A.: Two modellings for theory change. Journal of Philosophical Logic **17**, 157–170 (1988)
144. Grüne-Yanoff, T., Hansson, S.O.: From belief revision to preference change. In: T. Grüne-Yanoff, S.O. Hansson (eds.) Preference Change: Approaches from Philosophy, Economics and Psychology, pp. 159–184. Springer (2009)
145. Halpern, J.Y.: Reasoning about knowledge: an overview. In: J.Y. Halpern (ed.) Proceedings of the First Conference on Theoretical Aspects of Reasoning About Knowledge, pp. 1–17. Morgan Kaufmann, Los Altos (1988)
146. Hansson, S.O.: New operators for theory change. Theoria **55**, 114–132 (1989)
147. Hansson, S.O.: Belief base dynamics. Ph.D. thesis, Uppsala University (1991)
148. Hansson, S.O.: Belief contraction without recovery. Studia Logica **50**, 251–260 (1991)
149. Hansson, S.O.: A dyadic representation of belief. In: P. Gärdenfors (ed.) Belief Revision, no. 29 in Cambridge Tracts in Theoretical Computer Science, pp. 89–121. Cambridge University Press (1992)
150. Hansson, S.O.: In defense of base contraction. Synthese **91**, 239–245 (1992)
151. Hansson, S.O.: In defense of the Ramsey test. The Journal of Philosophy **89**, 522–540 (1992)
152. Hansson, S.O.: Reversing the Levi identity. Journal of Philosophical Logic **22**, 637–669 (1993)
153. Hansson, S.O.: Theory contraction and base contraction unified. Journal of Symbolic Logic **58**, 602–625 (1993)
154. Hansson, S.O.: Kernel contraction. Journal of Symbolic Logic **59**, 845–859 (1994)
155. Hansson, S.O.: Taking belief bases seriously. In: D. Prawitz, D. Westerståhl (eds.) Logic and Philosophy of Science in Uppsala, pp. 13–28. Kluwer Academic Publishers, Dordrecht (1994)
156. Hansson, S.O.: Changes in preference. Theory and Decision **38**, 1–28 (1995)
157. Hansson, S.O.: Hidden structures of belief. In: A. Fuhrmann, H. Rott (eds.) Logic, Actions and Information, pp. 79–100. de Gruyter, Berlin (1996)
158. Hansson, S.O.: Closure-invariant rationality postulates. In: E. Ejerhed, S. Lindström (eds.) Logic, Action and Cognition: Essays in Philosophical Logic, pp. 113–136. Springer Netherlands, Dordrecht (1997)
159. Hansson, S.O.: Semi-revision. Journal of Applied Non-Classical Logic **7**(1-2), 151–175 (1997)
160. Hansson, S.O.: What's new isn't always best. Theoria **63**, 1–13 (1997)
161. Hansson, S.O.: Recovery and epistemic residue. Journal of Logic, Language and Information **8**(4), 421–428 (1999)
162. Hansson, S.O.: A survey of non-prioritized belief revision. Erkenntnis **50**, 413–427 (1999)
163. Hansson, S.O.: A Textbook of Belief Dynamics. Theory Change and Database Updating. Applied Logic Series. Kluwer Academic Publishers, Dordrecht (1999)

164. Hansson, S.O.: Coherentist contraction. Journal of Philosophical Logic **29**, 315–330 (2000)
165. Hansson, S.O.: Ten philosophical problems in belief revision. Journal of Logic and Computation **13**, 37–49 (2003)
166. Hansson, S.O.: In praise of full meet contraction. Análisis Filosófico **27**(1), 134–146 (2006)
167. Hansson, S.O.: Logic of belief revision. In: E.N. Zalta (ed.) The Stanford Encyclopedia of Philosophy. The Metaphysics Research Lab. Center for the Study of Language and Information. Stanford University (2006). URL = http://plato.stanford.edu/entries/logic-belief-revision/
168. Hansson, S.O.: Contraction based on sentential selection. Journal of Logic and Computation **17**, 479–498 (2007)
169. Hansson, S.O.: Specified meet contraction. Erkenntnis **69**, 31–54 (2008)
170. Hansson, S.O.: Replacement: A Sheffer stroke for belief revision. Journal of Philosophical Logic **38**, 127–149 (2009)
171. Hansson, S.O.: Multiple and iterated contraction reduced to single-step single-sentence contraction. Synthese **173**, 153–177 (2010)
172. Hansson, S.O.: Global and iterated contraction and revision: An exploration of uniform and semi-uniform approaches. Journal of Philosophical Logic **41 (1)**, 143–172 (2011)
173. Hansson, S.O.: Eradication. Journal of Applied Logic **10**(1), 75–84 (2012)
174. Hansson, S.O.: Finite contractions on infinite belief sets. Studia Logica **100**(5), 907–920 (2012)
175. Hansson, S.O.: Blockage contraction. Journal of Philosophical Logic **42**(2), 415–442 (2013)
176. Hansson, S.O.: Bootstrap contraction. Studia Logica **101**(5), 1013–1029 (2013)
177. Hansson, S.O.: Cognitive realism in belief revision. In: E. Fermé, D.M. Gabbay, G.R. Simari (eds.) Trends in Belief Revision and Argumentation Dynamics, Logic and Cognitive Systems, pp. 57–74. College Publications (2013)
178. Hansson, S.O.: Maximal and perimaximal contraction. Synthese **190**(16), 3325–3348 (2013)
179. Hansson, S.O.: Outcome level analysis of belief contraction. Review of Symbolic Logic **6**(2), 183–204 (2013)
180. Hansson, S.O.: Repertoire contraction. Journal of Logic, Language and Information **22**(1), 1–21 (2013)
181. Hansson, S.O.: Contraction, revision, expansion: Representing belief change operations. In: R. Trypuz (ed.) Krister Segerberg on Logic of Actions, *Outstanding Contributions to Logic*, vol. 1, pp. 135–151. Springer Netherlands (2014)
182. Hansson, S.O.: Descriptor revision. Studia Logica **102**(5), 955–980 (2014)
183. Hansson, S.O.: Relations of epistemic proximity for belief change. Artificial Intelligence **217**, 76–91 (2014)
184. Hansson, S.O.: AGM contraction is not reconstructible as a descriptor operation. Journal of Logic and Computation (2015). Online first
185. Hansson, S.O.: A monoselective presentation of AGM revision. Studia Logica **103**, 1019–1033 (2015)
186. Hansson, S.O.: Alternatives to the Ramsey test. In: C. Beierle, G. Brewka, M. Thimm (eds.) Computational Models of Rationality. Essays Dedicated to Gabriele Kern-Isberner on the Occasion of her 60th Birthday, pp. 84–97. College Publications (2016)
187. Hansson, S.O.: Blockage revision. Journal of Logic, Language, and Information **25**, 37–50 (2016)
188. Hansson, S.O.: The co-occurrence test for non-monotonic inference. Artificial Intelligence **234**, 190–195 (2016)
189. Hansson, S.O.: Iterated descriptor revision and the logic of Ramsey test conditionals. Journal of Philosophical Logic **45**, 429–450 (2016)
190. Hansson, S.O.: Descriptor Revision. Belief Change Through Direct Choice (2017). Springer
191. Hansson, S.O., Fermé, E., Cantwell, J., Falappa, M.: Credibility-limited revision. Journal of Symbolic Logic **66**(4), 1581–1596 (2001)
192. Hansson, S.O., Makinson, D.: Applying normative rules with restraint. In: M.L. Chiara, K. Doets, D. Mundici, J. van Benthem (eds.) Logic and Scientific Method, pp. 313–332. Springer Netherlands, Dordrecht (1997)

193. Hansson, S.O., Olsson, E.: Levi contraction and AGM contraction: A comparison. Notre Dame Journal of Formal Logic **36**, 103–119 (1995)
194. Hansson, S.O., Olsson, E.: Providing foundations for coherentism. Erkenntnis **51**, 243–265 (1999)
195. Hansson, S.O., Wassermann, R.: Local change. Studia Logica **70 (1)**, 49–76 (2002)
196. Harman, G.: Change in View-Principles of Reasoning. MIT Press, Cambridge York (1986)
197. Harper, W.: Rational conceptual change. In: PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association, Volume Two: Symposia and Invited Papers, pp. 462–494. The University of Chicago Press (1977)
198. Hild, M., Spohn, W.: The measurement of ranks and the laws of iterated contraction. Artificial Intelligence **172**, 1195–1218 (2008)
199. Hintikka, J.: Knowledge and Belief: An Introduction to the Logic of the Two Notions. Cornell University Press, Ithaca, NY (1962)
200. Horty, J.: Norm change in the common law. In: S.O. Hansson (ed.) David Makinson on Classical Methods for Non-Classical Problems, Outstanding Contributions to Logic, pp. 335–355. Springer Netherlands (2014)
201. Jin, Y., Thielscher, M.: Iterated belief revision, revised. Artificial Intelligence **171**, 1–18 (2007)
202. Jin, Y., Thielscher, M.: Reinforcement belief revision. Journal of Logic and Computation **18**, 783–813 (2008)
203. Johnson, F.L.: Dependency-directed reconsideration: An anytime algorithm for hindsight knowledge-base optimization. Ph.D. thesis, State University of New York at Buffalo (2006)
204. Johnson, F.L., Shapiro, S.C.: Dependency-directed reconsideration: Belief base optimization for truth maintenance systems. In: Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI-05), pp. 313–320. AAAI Press (2005)
205. Katsuno, H., Mendelzon, A.: Propositional knowledge base revision and minimal change. Journal of Artificial Intelligence **52**, 263–294 (1991)
206. Katsuno, H., Mendelzon, A.: On the difference between updating a knowledge base and revising it. In: P. Gärdenfors (ed.) Belief Revision, no. 29 in Cambridge Tracts in Theoretical Computer Science, pp. 183–203. Cambridge University Press (1992)
207. Keller, A.M., Winslett, M.: On the use of an extended relational model to handle changing incomplete information. IEEE Transactions on Software Engineering **11**(7), 620–633 (1985)
208. Kelly, K.: Iterated belief revision, reliability, and inductive amnesia. Erkenntnis **50**, 11–58 (1998)
209. Kern-Isberner, G.: A thorough axiomatization of a principle of conditional preservation in belief revision. Annals of Mathematics and Artificial Intelligence **40(1-2)**, 127–164 (2004)
210. Kern-Isberner, G.: Linking iterated belief change operations to nonmonotonic reasoning. In: G. Brewka, J. Lang (eds.) Proceedings of the 11th International Conference on Knowledge Representation and Reasoning, KR'2008, pp. 166–176. AAAI Press, Menlo Park, CA (2008)
211. Konieczny, S., Medina Grespan, M., Pino Pérez, R.: Taxonomy of improvement operators and the problem of minimal change. In: Proceedings of the 12th International Conference on Principles of Knowledge Representation and Reasoning (KR 2010), pp. 161–170 (2010)
212. Konieczny, S., Pino Pérez, R.: A framework for iterated revision. Journal of Applied Non-Classical Logics **10**(3-4), 339–367 (2000)
213. Konieczny, S., Pino Perez, R.: Improvement operators. In: Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR'08), pp. 177–186 (2008)
214. Konieczny, S., Pino Pérez, R.: Logic based merging. Journal of Philosophical Logic **40**(2), 239–270 (2011)
215. Konieczny, S., Pino Pérez, R.: Merging information under constraints: a logical framework. Journal of Logic and Computation **12**(5), 773–808 (2002)
216. Kourousias, G., Makinson, D.: Parallel interpolation, splitting, and relevance in belief change. Journal of Symbolic Logic **72:3**, 994–1002 (2007)
217. Kyburg, H.E.: Probability and the Logic of Rational Belief. Wesleyan University Press, Middletown (1961)

218. Lang, J., van der Torre, L.: From belief change to preference change. In: M. Ghallab, C.D. Spyropoulos, N. Fakotakis, N.M. Avouris (eds.) ECAI 2008: 18th European Conference on Artificial Intelligence, Patras, Greece, July 21-25, 2008, Proceedings, *Frontiers in Artificial Intelligence and Applications*, vol. 178, pp. 351–355. IOS Press (2008)
219. Lehmann, D.J.: Belief revision, revised. In: Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI'95), pp. 1534–1540 (1995)
220. Levi, I.: Subjunctives, dispositions, and chances. Synthese **34**, 423–455 (1977)
221. Levi, I.: Iteration of conditionals and the Ramsey test. Synthese **76**, 49–81 (1988)
222. Levi, I.: The fixation of belief and its undoing: changing beliefs through inquiry. Cambridge University Press, Cambridge (1991)
223. Levi, I.: Contraction and informational value (1997). (manuscript. Available at http://www.columbia.edu/ levi/contraction.pdf)
224. Levi, I.: Mild Contraction: Evaluating Loss of Information Due to Loss of Belief. Oxford University Press, Cambridge (2005)
225. Li, J.: A note on partial meet package contraction. Journal of Logic, Language and Information **7**, 139–142 (1998)
226. Liberatore, P., Schaerf, M.: Arbitration: A commutative operator for belief revision. In: Proceedings of the Second World Conference on the Fundamentals of Artificial Intelligence (WOCFAI'95 (1995)
227. Liberatore, P., Schaerf, M.: Arbitration (or how to merge knowledge bases). Knowledge and Data Engineering, IEEE Transactions on Knowledge and Data Engineering **10**(1), 76–90 (1998)
228. Lin, J.: Integration of weighted knowledge bases. Artificial Intelligence **83 (2)**, 363–378 (1996)
229. Lindström, S.: A semantic approach to nonmonotonic reasoning: Inference operations and choice. Uppsala Prints and Preprints in Philosophy 6, Department of Philosophy, Uppsala University (1991)
230. Lindström, S., Rabinowicz, W.: On probabilistic representation of non-probabilistic belief revision. Journal of Philosophical Logic **19**, 69–101 (1989)
231. Lindström, S., Rabinowicz, W.: Epistemic entrenchment with incomparabilities and relational belief revision. In: A. Fuhrmann, M. Morreau (eds.) The Logic of Theory Change: Workshop, Konstanz, FRG, October 13–15, 1989, *Lecture Notes in Artificial Intelligence*, vol. 465, pp. 93–126. Springer (1991)
232. Lindström, S., Rabinowicz, W.: Belief revision, epistemic conditionals and the Ramsey test. Synthese **91**, 195–237 (1992)
233. Lindström, S., Rabinowicz, W.: Conditionals and the Ramsey test. In: D. Gabbay, P. Smets (eds.) Handbook of Defeasible Reasoning and Uncertainty Management Systems, vol. 3 (Belief Change), pp. 147–188. Kluwer (1998)
234. Lindström, S., Rabinowicz, W.: Belief change for introspective agents. In: B. Hansson, S. Halldén, Nils-Eric-Sahlin (eds.) Spinning Ideas. Electronic Essays Dedicated to Peter Gärdenfors on His Fiftieth Birthday. http://www.lucs.lu.se/spinning/ (1999)
235. Lindström, S., Rabinowicz, W.: DDL unlimited. Dynamic doxastic logic for introspective agents. Erkenntnis **51**, 353–385 (1999)
236. Liu, W.: A framework for multi-agent belief revision. Ph.D. thesis, University of Newcastle (2002)
237. Makinson, D.: The paradox of the preface. Analysis **25**, 205–207 (1965)
238. Makinson, D.: How to give it up: A survey of some recent work on formal aspects of the logic of theory change. Synthese **62**, 347–363 (1985)
239. Makinson, D.: On the status of the postulate of recovery in the logic of theory change. Journal of Philosophical Logic **16**, 383–394 (1987)
240. Makinson, D.: The Gärdenfors impossibility theorem in nonmonotonic contexts. Studia Logica **49**, 1–6 (1990)
241. Makinson, D.: On the force of some apparent counterexamples to recovery. In: E.G. Valdés, W. Krawietz, G. von Wright, R. Zimmerling (eds.) Normative Systems in Legal and Moral Theory: Festschrift for Carlos Alchourrón and Eugenio Bulygin, pp. 475–481. Duncker & Humblot, Berlin (1997)

242. Makinson, D.: Screened revision. Theoria **63**, 14–23 (1997)
243. Makinson, D.: Propositional relevance through letter-sharing. Journal of Applied Logic **7**(4), 377–387 (2009)
244. Makinson, D.: Conditional probability in the light of qualitative belief change. Journal of Philosophical Logic **40**(2), 121–153 (2011)
245. Makinson, D., Gärdenfors, P.: Relation between the logic of theory change and nonmonotonic logic. In: A. Fuhrmann, M. Morreau (eds.) The Logic of Theory Change: Workshop, Konstanz, FRG, October 13–15, 1989, *Lecture Notes in Artificial Intelligence*, vol. 465, pp. 185–205. Springer (1991)
246. Malheiro, B., Jennings, N.R., Oliveira, E.: Belief revision in multi-agent systems. In: Proceedings of the 11th European Conference on Artificial Intelligence (ECAI 94), pp. 294–298 (1994)
247. Malheiro, B., Oliveira, E.: Solving conflicting beliefs with a distributed belief revision approach. In: M. Monard, J. Sichman (eds.) Advances in Artificial Intelligence, *Lecture Notes in Computer Science*, vol. 1952, pp. 146–155. Springer (2000)
248. Mares, E.D.: A paraconsistent theory of belief revision. Erkenntnis **56**(2), 229–246 (2002)
249. Martins, J., Shapiro, S.: A model for belief revision. Artificial Intelligence **35**, 25–79 (1988)
250. Meyer, T., Heidema, J., Labuschagne, W., Leenen, L.: Systematic withdrawal. Journal of Philosophical Logic **31:5**, 415–443 (2002)
251. Meyer, T., Lee, K., Booth, R.: Knowledge integration for description logics. In: Proceedings of the 7th International Symposium on Logical Formalizations of Commonsense Reasoning, pp. 645–650. AAAI Press (2005)
252. Moore, G.E.: A reply to my critics. In: P. Schilpp (ed.) The Philosophy of G.E. Moore, pp. 535–677. Northwestern University, Evanston IL (1942). The Library of Living Philosophers (volume 4)
253. Nayak, A.: Iterated belief change based on epistemic entrenchment. Erkenntnis **41**, 353–390 (1994)
254. Nayak, A., Pagnucco, M., Peppas, P.: Dynamic belief revision operators. Artificial Intelligence **146:2**, 193–228 (2003)
255. Nebel, B.: A knowledge level analysis of belief revision. In: Proceedings of the 1st International Conference of Principles of Knowledge Representation and Reasoning, pp. 301–311. Morgan Kaufmann (1989)
256. Nebel, B.: Syntax-based approaches of belief revision. In: P. Gärdenfors (ed.) Belief Revision, no. 29 in Cambridge Tracts in Theoretical Computer Science, pp. 52–88. Cambridge University Press (1992)
257. Nebel, B.: How hard is it to revise a belief base? In: D. Dubois, H. Prade (eds.) Handbook of Defeasible Reasoning and Uncertainty Management Systems, Volume 3: Belief Change, pp. 77–145. Kluwer Academic Publishers, Dordrecht (1998)
258. Newell, A.: The knowledge level. Artificial Intelligence **18**, 87–127 (1982)
259. Niederée, R.: Multiple contraction: A further case against Gärdenfors' principle of recovery. In: A. Fuhrmann, M. Morreau (eds.) The Logic of Theory Change: Workshop, Konstanz, FRG, October 13–15, 1989, *Lecture Notes in Artificial Intelligence*, vol. 465, pp. 322–334. Springer (1991)
260. Nittka, A.: A method for reasoning about other agents' beliefs from observations. Ph.D. thesis, University of Leipzig (2008)
261. Olsson, E.: Coherence. Ph.D. thesis, Department of Philosophy. Uppsala University (1997)
262. Olsson, E.: Belief revision, rational choice and the unity of reason. Studia Logica **73**, 219–240 (2003)
263. Olsson, E., Enqvist, S. (eds.): Belief Revision Meets Philosophy of Science, *Logic, Epistemology, and the Unity of Science*, vol. 21. Springer Netherlands (2011)
264. Paglieri, F.: See what you want, believe what you like: Relevance and likeability in belief dynamics. In: L. Cañamero (ed.) Papers from the AISB'05 Symposium, Agents That Want and Like: Motivational and Emotional Roots of Cognition and Action, pp. 90–97. Hatfield, AISB (2005)

265. Paglieri, F.: Belief dynamics: From formal models to cognitive architectures, and back again. Ph.D. thesis, Università degli Studi di Siena (2006)
266. Paglieri, F., Castelfranchi, C.: Argumentation and data-oriented belief revision: On the two-sided nature of epistemic change. In: CMNA IV: 4th workshop on Computational Models of Natural Argument, pp. 5–12. Valencia: ECAI 2004 (2004)
267. Paglieri, F., Castelfranchi, C.: The Toulmin test: Framing argumentation within belief revision theories. In: D. Hitchcock, B. Verheij (eds.) Arguing on the Toulmin Model, pp. 359–377. Springer (2006)
268. Pagnucco, M.: The role of abductive reasoning within the process of belief revision. Ph.D. thesis, Department of Computer Science, University of Sydney (1996)
269. Pagnucco, M., Nayak, A., Foo, N.: Abductive expansion: Abductive inference and the process of belief change. In: C. Zhang, J. Debenham, D. Lukose (eds.) Proceedings of the Seventh Australian Joint Conference on Artificial Intelligence (AI94), pp. 267–274. Armidale, Australia (1994)
270. Parikh, R.: Beliefs, belief revision, and splitting languages. In: Logic, Language, and Computation, *CSLI Lecture Notes*, vol. 96-2, pp. 266–268. Springer-Verlag (1999)
271. Peppas, P., Chopra, S., Foo, N.Y.: Distance semantics for relevance-sensitive belief revision. In: Principles of Knowledge Representation and Reasoning: Proceedings of the Ninth International Conference (KR2004), Whistler, Canada, June 2–5, 2004, pp. 319–328 (2004)
272. Plaza, J.: Logics of public communications. In: M. Emrich, M. Pfeifer, M. Hadzikadic, Z. Ras (eds.) Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems: Poster Session Program, pp. 201–216. Oak Ridge National Laboratory (1989)
273. Priest, G.: Paraconsistent belief revision. Theoria **67:3**, 214–228 (2001)
274. Rabinowicz, W.: Global belief revision based on similarities between worlds. In: S.O. Hansson, W. Rabinowicz (eds.) Logic for a change, no. 9 in Uppsala Prints and Preprints in Philosophy, pp. 80–105. Department of Philosophy, Uppsala University (1995)
275. Reis, M.D.L.: On theory multiple contraction. Ph.D. thesis, Universidade da Madeira (2011)
276. Reis, M.D.L., Fermé, E.: Possible worlds semantics for partial meet multiple contraction. Journal of Philosophical Logic **41**(1), 7–28 (2012)
277. Reis, M.D.L., Fermé, E., Peppas, P.: Construction of system of spheres-based transitively relational partial meet multiple contractions: An impossibility result. Artificial Intelligence **233**, 122–141 (2016)
278. Reis, M.D.L., Fermé, E., Peppas, P.: Two axiomatic characterizations for the system of spheres-based (and the epistemic entrenchment-based) multiple contractions. Annals of Mathematics and Artificial Intelligence **78**(3-4), 181–203 (2016)
279. Reiter, R.: A theory of diagnosis from first principles. Artificial Intelligence **32**, 57–95 (1987)
280. Restall, G., Slaney, J.K.: Realistic belief revision. In: Proceedings of the Second World Conference on Foundations of Artificial Intelligence. WOCFAI, vol. 95, pp. 367–378 (1995)
281. Revesz, P.Z.: On the semantics of theory change: Arbitration between old and new information. In: Proceedings of the Twelfth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Databases, pp. 71–92 (1993)
282. Ribeiro, M.M., Wassermann, R.: Base revision for ontology debugging. Journal of Logic and Computation **19**(5), 721–743 (2009)
283. de Rijke, M.: Meeting some neighbours. In: J. van Eijck, A. Visser (eds.) Logic and Information Flow, pp. 170–195. MIT Press, Cambridge MA (1994)
284. Rodrigues, O., Benevides, M.: Belief revision in pseudo-definite sets. In: Proceedings of the 11th Brazilian Symposium on Artificial Intelligence (SBIA '94) (1994)
285. Rott, H.: Ifs, though and because. Erkenntnis **25**, 345–37 (1986)
286. Rott, H.: Conditionals and theory change: Revision, expansions, and additions. Synthese **81**, 91–113 (1989)
287. Rott, H.: A nonmonotonic conditional logic for belief revision. Part 1: Semantics and logic of simple conditionals. In: A. Fuhrmann, M. Morreau (eds.) The Logic of Theory Change: Workshop, Konstanz, FRG, October 13–15, 1989, *Lecture Notes in Artificial Intelligence*, vol. 465, pp. 135–181. Springer (1991)

288. Rott, H.: Two methods of constructing contractions and revisions of knowledge systems. Journal of Philosophical Logic **20**, 149–173 (1991)
289. Rott, H.: On the logic of theory change: More maps between different kinds of contraction functions. In: P. Gärdenfors (ed.) Belief Revision, no. 29 in Cambridge Tracts in Theoretical Computer Science, pp. 122–141. Cambridge University Press (1992)
290. Rott, H.: Preferential belief change using generalized epistemic entrenchment. Journal of Logic, Language and Information **1**, 45–78 (1992)
291. Rott, H.: Belief contraction in the context of the general theory of rational choice. Journal of Symbolic Logic **58**, 1426–1450 (1993)
292. Rott, H.: "Just because". Taking belief bases very seriously. In: S.O. Hansson, W. Rabinowicz (eds.) Logic for a Change, no. 9 in Uppsala Prints and Preprints in Philosophy, pp. 106–124. Department of Philosophy, Uppsala University (1995)
293. Rott, H.: "Just because". Taking belief bases seriously. In: S. Buss, P. Hajek, P. Pudlak (eds.) Logic Colloquium '98: Proceedings of the Annual European Summer Meeting of the Association for Symbolic Logic, *Lecture Notes in Logic*, vol. 13. Association for Symbolic Logic, Prague (2000)
294. Rott, H.: Two dogmas of belief revision. Journal of Philosophy **97(9)**, 503–522 (2000)
295. Rott, H.: Change, Choice and Inference: A Study of Belief Revision and Nonmonotonic Reasoning. Oxford Logic Guides. Clarendon Press, Oxford (2001)
296. Rott, H.: Basic entrenchment. Studia Logica **73**, 257–280 (2003)
297. Rott, H.: Coherence and conservatism in the dynamics of belief. Part II: Iterated belief change without dispositional coherence. Journal of Logic and Computation **13**, 111–145 (2003)
298. Rott, H.: Degrees all the way down: Beliefs, non-beliefs and disbeliefs. In: F. Huber, C. Schmidt-Petri (eds.) Degrees of Belief, pp. 301–339. Springer Netherlands, Dordrecht (2009)
299. Rott, H.: Shifting priorities: Simple representations for twenty-seven iterated theory change operators. In: D. Makinson, J. Malinowski, H. Wansing (eds.) Towards Mathematical Philosophy, no. 28 in Trends in Logic, pp. 269–296. Springer Science (2009)
300. Rott, H.: Bounded revision: Two-dimensional belief change between conservative and moderate revision. Journal of Philosophical Logic **41**(1), 173–200 (2012)
301. Rott, H., Hansson, S.O.: Safe contraction revisited. In: S.O. Hansson (ed.) David Makinson on Classical Methods for Non-Classical Problems, *Outstanding Contributions to Logic*, vol. 3, pp. 35–70. Springer Netherlands (2014)
302. Rott, H., Pagnucco, M.: Severe withdrawal (and recovery). Journal of Philosophical Logic **28**, 501–547 (1999)
303. Ryan, M., Schobbens, P.Y.: Counterfactuals and updates as inverse modalities. Journal of Logic, Language and Information **6**, 123–146 (1997)
304. Samet, D.: Hypothetical knowledge and games with perfect information. Games and Economic Behavior **17**, 230–251 (1996)
305. Satoh, K.: Nonmonotonic reasoning by minimal belief revision. In: Proceedings of the International Conference on Fifth Generation Computer Systems, pp. 455–462. Springer-Verlag, Tokyo (1988)
306. Schlechta, K.: Non-prioritized belief revision based on distances between models. Theoria **63**, 34–53 (1997)
307. Segerberg, K.: Belief revision from the point of view of doxastic logic. Bulletin of the IGPL **3**, 535–553 (1995)
308. Segerberg, K.: Some questions about hypertheories. In: S.O. Hansson, W. Rabinowicz (eds.) Logic for a Change, no. 9 in Uppsala Prints and Preprints in Philosophy, pp. 136–153. Department of Philosophy, Uppsala University (1995)
309. Segerberg, K.: Irrevocable belief revision in dynamic doxastic logic. Notre Dame Journal of Formal Logic **39**(3), 287–306 (1998)
310. Segerberg, K.: Two traditions in the logic of belief: Bringing them together. In: H.J. Ohlbach, U. Reyle (eds.) Logic, Language and Reasoning: Essays in Honour of Dov Gabbay, pp. 135–147. Springer Netherlands, Dordrecht (1999)

311. Sen, A.: Quasi-transitivity, rational choice and collective decisions. Review of Economic Studies **36**(3), 381–393 (1969)
312. Sen, A.: Collective Choice and Social Welfare. Holden-Day (1970)
313. Sen, A.: Choice functions and revealed preference. Review of Economic Studies **38**, 307–317 (1971)
314. Spohn, W.: Eine Theorie der Kausalität (1983). Unpublished Habilitationsschrift, available at http://www.uni-konstanz.de/FuF/Philo/Philosophie/Spohn/spohn_files/Habilitation.pdf
315. Spohn, W.: Ordinal conditional functions: A dynamic theory of epistemic states. In: W. Harper, B. Skyrms (eds.) Causation in Decision, Belief Change and Statistics, vol. 2, pp. 105–134. D. Reidel, Dordrecht (1988)
316. Spohn, W.: A survey of ranking theory. In: F. Huber, C. Schmidt-Petri (eds.) Degrees of Belief, pp. 185–228. Springer Netherlands, Dordrecht (2009)
317. Spohn, W.: Multiple contraction revisited. In: M. Suárez, M. Dorato, M. Rédei (eds.) EPSA Epistemology and Methodology of Science, Vol. 1. Launch of the European Philosophy of Science Association, pp. 279–288. Springer, Dordrecht (2010)
318. Stalnaker, R.: A theory of conditionals. In: N. Rescher (ed.) Studies in Logical Theory, no. 2 in American Philosophical Quarterly Monograph Series. Blackwell, Oxford (1968). Also printed in W. Harper, R. C. Stalnaker and G. Pearce (eds.), *Ifs*, Reidel, Dordrecht, 1981
319. Stalnaker, R.: Iterated belief revision. Erkenntnis **70**(2), 189–209 (2009)
320. Tanaka, K.: The AGM theory and inconsistent belief change. Logique et Analyse **48**(189–192), 113–150 (2005)
321. Tarski, A.: Logic, Semantics, Metamathematics. Papers from 1923 to 1938. Translated by J. H. Woodger. Clarendon Press, Oxford (1956)
322. Testa, R., Coniglio, M., Ribeiro, M.: Paraconsistent belief revision based on a formal consistency operator. CLE e-prints 15(8), University of Campinas (2015)
323. Toulmin, S.E.: The Uses of Argument. Cambridge University Press (1958/2003)
324. Tversky, A., Kahneman, D.: Judgment under uncertainty: Heuristics and biases. Science **185**(4157), 1124–1131 (1974)
325. von Wright, G.H.: An Essay in Modal Logic. Amsterdam: North-Holland (1951)
326. Wassermann, R.: Resource bounded belief revision. Erkenntnis **50**, 429–446 (1999)
327. Wassermann, R.: Resource bounded belief revision. Ph.D. thesis, University of Amsterdam (2000)
328. Wassermann, R.: Local diagnosis. Journal of Applied Non-Classical Logics **11**, 107–129 (2001)
329. Weydert, E.: Relevance and revision: About generalizing syntax-based belief revision. In: D. Pearce, G. Wagner (eds.) Logics in AI, European Workshop, JELIA '92, *Lecture Notes in Computer Science*, vol. 633, pp. 126–138. Springer-Verlag, Germany (1992)
330. Weydert, E.: General belief measures. In: R.L. de Mántaras, D. Poole (eds.) UAI '94: Proceedings of the Tenth Annual Conference on Uncertainty in Artificial Intelligence, pp. 575–582. Morgan Kaufmann (1994)
331. Weydert, E.: Projective default epistemology. In: G. Kern-Isberner, W. Rödder, F. Kulmann (eds.) Conditionals, Information, and Inference. International Workshop, WCII 2002, Hagen, Germany, May 13–15, 2002, Revised Selected Papers, *Lecture Notes in Computer Science*, vol. 3301, pp. 65–85. Springer (2005)
332. Williams, M.A.: Two operators for theory bases. In: Proceedings of the Australian Joint Artificial Intelligence Conference, pp. 259–265. World Scientific (1992)
333. Williams, M.A.: On the logic of theory base change. In: C. MacNish, D. Pearce, L.M. Pereira (eds.) Logics in Artificial Intelligence: European Workshop JELIA '94 York, UK, September 5–8, 1994 Proceedings, no. 835 in Lecture Notes Series in Computer Science, pp. 86–105. Springer (1994)
334. Williams, M.A.: Transmutations of knowledge systems. In: J. Doyle, E. Sandewall, P. Torasso (eds.) Proceedings of the fourth International Conference on Principles of Knowledge Representation and Reasoning, pp. 619–629. Morgan Kaufmann, Bonn, Germany (1994)

335. Williams, M.A.: Iterated theory base change: A computational model. In: Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95), vol. 2, pp. 1541–1547. Montreal, Canada (1995)
336. Williams, M.A.: Anytime belief revision. In: Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97), vol. 1, pp. 74–79. Morgan Kaufmann (1997)
337. Williams, M.A., Sims, A.: Saten: An object-oriented web-based revision and extraction engine. In: International Workshop on Nonmonotonic Reasoning (NMR'2000) (2000). Online Computer Science Abstract, http://arxiv.org/abs/cs.AI/0003059/
338. Winslett, M.: Reasoning about action using a possible models approach. In: Proceedings of the Seventh American Association for Artificial Intelligence Conference, pp. 89–93 (1988)
339. Zhang, D.: Properties of iterated multiple belief revision. In: V. Lifschitz, I. Niemelä (eds.) LPNMR 2004, *Lecture Notes in Artificial Intelligence*, vol. 2923, pp. 314–325. Springer (2004)
340. Zhang, D.: A logic-based axiomatic model of bargaining. Artificial Intelligence **174**, 1307–1322 (2010)
341. Zhang, L., Hansson, S.O.: How to make up one's mind. Logic Journal of the IGPL **23**, 705–717 (2015)
342. Zhuang, Z., Pagnucco, M.: Horn contraction via epistemic entrenchment. In: T. Janhunen, I. Niemelä (eds.) Logics in Artificial Intelligence, *Lecture Notes in Computer Science*, vol. 6341, pp. 339–351. Springer (2010)