

AGM 25 Years

Twenty-Five Years of Research in Belief Change

Eduardo Fermé · Sven Ove Hansson

Received: 7 October 2010 / Accepted: 22 December 2010 / Published online: 26 February 2011
© Springer Science+Business Media B.V. 2011

Abstract The 1985 paper by Carlos Alchourrón (1931–1996), Peter Gärdenfors, and David Makinson (AGM), “On the Logic of Theory Change: Partial Meet Contraction and Revision Functions” 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. In this review, the first twenty-five years of this development are summarized. The topics covered include equivalent characterizations of AGM operations, extended representations of the belief states, change operators not included in the original framework, iterated change, applications of the model, its connections with other formal frameworks, computability of AGM operations, and criticism of the model.

Keywords Belief change · Belief revision · Theory change · Partial meet contraction · Partial meet revision · AGM · Carlos Alchourrón · Peter Gärdenfors · David Makinson

1 Introduction

Many research papers have been called “seminal”, but few deserve that designation as much as the article in the *Journal of Symbolic Logic* in 1985 by Carlos Alchourrón (1931–1996), Peter Gärdenfors, and David Makinson,

E Fermé (✉)

Madeira Interactive Technologies Institute, University of Madeira, Madeira, Portugal
e-mail: ferme@uma.pt

S. O. Hansson

Department of Philosophy and the History of Technology,
Royal Institute of Technology, Stockholm, Sweden
e-mail: soh@kth.se

“On the Logic of Theory Change: Partial Meet Contraction and Revision Functions”. The immediate prehistory of this article has been told previously by Makinson [189], and now also by Gärdenfors [112]. Like all major scientific achievements, that of the AGM trio was much facilitated by previous work by others. Computer scientists such as Doyle [65] and Fagin et al. [72] had developed precise models of database updating, and philosophers such as Harper [146] and in particular Levi [170, 171] had investigated philosophical issues in belief change.

The yearly number of Web of Science citations of the AGM article has increased from about 15 per in the 1990’s to about 50 in the most recent years—a remarkable record for a logic paper. Its impact has been profound both among philosophers and in the artificial intelligence community [48].

It is the purpose of this review to exemplify the diversity and significance of the research that has been inspired by the AGM article. Section 2 covers some of the equivalent characterizations of the AGM model that have contributed to our understanding of it. In Section 3 we discuss some of the more common criticisms of AGM. Section 4 is devoted to extensions and variations of the belief state representation in AGM (belief sets), Section 5 to iterated change, and Section 6 to operations of change that have been introduced in addition to the three AGM operations (contraction, revision, and expansion). Section 7 provides an overview of applications and connections with other areas of research, and Section 8 summarizes some results on the computability of belief change operations.

2 Equivalent Characterizations

The AGM model has been expressed in several seemingly dissimilar ways.

2.1 AGM Briefly Summarized

In the basic AGM framework a person’s belief state is represented by a set K of sentences that is closed under the operation Cn of logical consequence, i.e. $K = \text{Cn}(K)$. In order to contract K by a sentence p we consider the set $K \perp p$ of all maximal subsets of K not implying p . A selection function γ selects the most plausible elements of $K \perp p$. Whenever $K \perp p$ is non-empty, $\gamma(K \perp p)$ is a non-empty subset of $K \perp p$. The *partial meet contraction* \sim_γ based on γ is defined as $K \sim_\gamma p = \bigcap \gamma(K \perp p)$. In the limiting case when p is a tautology, $K \sim_\gamma p = K$. If $\gamma(K \perp p) = K \perp p$ for all p , then \sim_γ is the *full meet contraction*, also denoted \sim . If $\gamma(K \perp p)$ is always a singleton when $K \perp p$ is non-empty, then \sim_γ is a *maxichoice contraction*. If there is a transitive relation \sqsubseteq such that $\gamma(K \perp p) = \{X \in K \perp p \mid (\forall Y \in K \perp p)(Y \sqsubseteq X)\}$ for all nonempty $K \perp p$, then \sim_γ is *transitively relational*. There are two operations that add beliefs to a belief set: $K + p = \text{Cn}(K \cup \{p\})$ (expansion) and $K *_\gamma p = \text{Cn}((K \sim_\gamma \neg p) \cup \{p\})$ (revision). In the AGM article, axiomatic characterizations were provided for all these operations. These axioms are

Table 1 The standard AGM postulates

Basic contraction postulates	Basic revision postulates
$K \div p = \text{Cn}(K \div p)$ (closure)	$K * p = \text{Cn}(K * p)$ (closure)
$K \div p \subseteq K$ (inclusion)	$K * p \subseteq K + p$ (inclusion)
If $K \not\vdash p$ then $K \div p = K$ (vacuity)	If $K \not\vdash \neg p$ then $K + p \subseteq K * p$ (vacuity)
If $\vdash p$ then $p \notin K \div p$ (success)	$p \in K * p$ (success)
If $\vdash p \leftrightarrow q$ then $K \div p = K \div q$ (extensionality)	If $\vdash p \leftrightarrow q$ then $K * p = K * q$ (extensionality)
$K \subseteq (K \div p) + p$ (recovery)	If $p \not\vdash \perp$ then $K * p \not\vdash \perp$ (consistency)
Supplementary contraction postulates	Supplementary revision postulates
$(K \div p) \cap (K \div q) \subseteq K \div (p \& q)$ (conjunctive overlap)	$K * (p \& q) \subseteq (K * p) + q$ (superexpansion)
If $p \notin K \div (p \& q)$ then $K \div (p \& q) \subseteq K \div p$ (conjunctive inclusion)	If $K * p \not\vdash \neg q$ then $(K * p) + q \subseteq K * (p \& q)$ (subexpansion)

The six basic contraction postulates exactly characterize partial meet contraction, and all eight contraction postulates exactly characterize transitively relational partial meet contraction. Similarly, the six basic revision postulates exactly characterize partial meet revision, and all eight revision postulates exactly characterize transitively relational partial meet revision. $X \vdash p$ is an alternative notation for $p \in \text{Cn}(X)$ and $\vdash p$ for $p \in \text{Cn}(\emptyset)$

listed in Table 1. Subsequently, several other equivalent characterizations have been developed.

2.2 Safe and Kernel Contraction

Alchourrón and Makinson proposed the operation of *safe contraction* that is based on a non-circular relation $<$ on the elements of K . An element a of K is safe with respect to p if and only if all inclusion-minimal p -implying subsets of K either do not contain a or contain some b such that $b < a$. The safe contraction \div based on $<$ yields as outcome the logical closure of the set of sentences in K that are safe with respect to p [6]. All safe contractions are partial meet contractions. Additional results on safe contraction have been obtained by Rott [225].

Kernel contraction, introduced in [128], is a non-relational generalization of safe contraction. Let Kp be the set of minimal p -implying subsets of K . An incision function is a function σ that selects sentences to be discarded. It satisfies the two basic properties (i) $\sigma(K \perp p) \subseteq \bigcup(K \perp p)$ and (ii) if $\emptyset \neq X \in K \perp p$, then $X \cap \sigma(K \perp p) \neq \emptyset$. The kernel contraction \approx_σ based on σ is defined by the relationship $K \approx_\sigma p = K \setminus \sigma(K \perp p)$. The partial meet contractions on a belief set coincide exactly with the kernel contractions that satisfy the additional condition of smoothness, namely that if $X \subseteq K$ and $\text{Cn}(X) \cap \sigma(K \perp p) \neq \emptyset$, then $X \cap \sigma(K \perp p) \neq \emptyset$.

2.3 Epistemic Entrenchment

Epistemic entrenchment which was introduced in [107, 113] and [109] is a binary relation \leq on the sentences in the belief set K such that in contrac-

tion, giving up beliefs with lower entrenchment is preferred to giving up those with higher entrenchment. Entrenchment is assumed to satisfy transitivity, dominance (if $p \vdash q$ then $p \leq q$), conjunctiveness (either $p \leq p \& q$ or $q \leq p \& q$), minimality (non-elements of K are the sentences with minimal entrenchment) and maximality (tautologies are the sentences with maximal entrenchment). Entrenchment-based contraction is constructed such that for all non-tautologous p : $K \div p = \{q \in K \mid q < p \vee q\}$. A contraction operator is an entrenchment-based contraction if and only if it is a transitively relational partial meet contraction. Rott has investigated similar operations that are based on entrenchment relations satisfying various weaker sets of conditions [230, 232]. Other results on entrenchment have been obtained in [78, 199] and elsewhere.

Alternatively, a contraction operator can be defined as follows: $K \div p = \{q \in K \mid q < p\}$. This is severe withdrawal (also called mild contraction or Rott's contraction). It was axiomatized independently by Pagnucco and Rott [237] and by Fermé and Rodriguez [86]. Arló-Costa and Levi have analyzed it in terms of minimal loss of informational value [12]. It has been shown to satisfy the implausible postulate of expulsiveness. (If $\not\vdash p$ and $\not\vdash q$, then either $p \notin K \div q$ or $q \notin K \div p$) [133]. However, Lindström and Rabinowicz have proposed that the truth lies somewhere between severe withdrawal and the originally proposed entrenchment-based contraction [177]. This has been called Lindström's and Rabinowicz's interpolation thesis [228].

2.4 Grove's Spheres

In Grove's possible world-modelling of belief revision, each logically closed set X is represented by the set $[X]$ of possible worlds including it, and similarly each sentence p is represented by the set $[p]$ of possible worlds including it [117]. The expansion outcome $K + p$ will then be represented by the set $[K + p] = [K] \cap [p]$, and a contraction outcome $K \div p$ by some superset $[K \div p]$ of $[K]$ that includes at least one $\neg p$ -world. The most remarkable result in this framework is the equivalence between transitively relational partial meet operations and operations based on a *sphere system* around the belief set. Let S_1, S_2, \dots be sets of possible worlds such that $[K] \subset S_1$ and $S_k \subset S_{k+1}$ for all k . The contraction \div is a sphere-based contraction on K if and only if it holds for all p that $[K \div p]$ is the set of worlds obtained by adding to $[K]$ all the $\neg p$ -worlds in the smallest sphere S_k that has any $\neg p$ -world. Sphere-based contraction coincides with transitively relational partial meet contraction (and thus with entrenchment-based contraction). It has been used for instance in analyses of iterated belief change [235].

Booth et al. have proposed a model of belief change that makes use of two relations \leq and \preceq on possible worlds, instead of only one as in the standard Grove model. Both relations represent plausibility, and \preceq is a reflexive subrelation of \leq [35].

2.5 Distance Models

Several authors have proposed models of belief change that employ representations of similarity or distance. In [125] some variants of AGM contraction operators were reconstructed in terms of a mechanism that selects those among the eligible belief sets that are closest to the original belief set. Katsuno and Mendelzon have shown how AGM revision can be reconstructed in terms of a closeness relation that is used to ensure that changes are minimal [155]. Lehmann, Magidor and Schlechta have obtained close connections between partial meet revision and models employing what they call pseudo-distances, a generalization of the common notion of a distance [168].

2.6 Specified Meet Contraction

A belief set K is finite-based if there is some finite set S of sentences such that $\text{Cn}(K) = \text{Cn}(S)$. In a cognitively realistic model, we can expect the original belief set to be finite-based, and the contraction operator should satisfy the following postulate:

Finite-based outcome: If K is finite-based, then so is $K \div p$.

Not all partial meet contractions satisfy Finite-based outcome. Those that do so have been shown to be reconstructible as *specified meet contractions*. A specified meet contraction is defined by the relationship $K \div p = K \sim f(p)$, where \sim is full meet contraction and f is a function from and to the language. Intuitively, $f(p)$ is the sentence that is really removed in order to contract by p . (Hence if we remove only q in order to contract by $q \& r$ then we can have $f(q \& r) = q$.) Various axioms for contraction, including the standard AGM postulates, have been shown to correspond exactly to some property of the sentential selector f [137, 138].

3 Criticism of the Model

Much of the critical discussion on the AGM model has referred either to the postulates for partial meet contraction and revision (Sections 3.1 and 3.2) or to the use of belief sets to represent belief states (Sections 3.3 and 3.4).

On philosophical issues relating to AGM see also [135] and [230].

3.1 The Recovery Postulate

By far the most criticized of the postulates is one of the basic postulates for contraction:

Recovery: $K \subseteq (K \div p) + p$.

Recovery is based on the intuition that we get all of the beliefs back after first contracting and then expanding by the same belief [106]. However, counterexamples have been constructed in which the recovery postulate seems to give rise to implausible results. Suppose for instance that your original belief set K contained the belief s that Cleopatra had a son. You receive information that makes you contract the belief that she had a child ($s \vee d$, where d denotes that she had a daughter). You will then arrive at a belief set $K \div (s \vee d)$, and clearly $s \notin K \div (s \vee d)$. Now suppose that you receive more reliable information that makes you expand your belief set by $s \vee d$. We can then expect to have $s \notin K \div (s \vee d) + (s \vee d)$, contrary to Recovery [121]. In a retort, Makinson argued that the theories considered in this and similar examples are implicitly “clothed” with additional justificational structure. In his view, Recovery remains appropriate for “naked”, logically closed theories [187]. In a somewhat similar vein, Glaister argued that this contraction is better represented by a multiple contraction than by a contraction by $s \vee d$ [115].

Belief base models do not in general satisfy Recovery. (See Section 4.1.) In Makinson’s terminology, operations that satisfy the other five basic contraction postulates but not Recovery are called withdrawals [185].

3.2 The Success Postulates

Partial meet revision satisfies the following postulate:

$$\textit{Revision success: } p \in K * p$$

Several authors have found this to be an implausible feature of belief revision, even if p is not a contradiction [55, 132]. Similarly, one of the AGM postulates for partial meet contraction:

$$\textit{Contraction success: } \text{If } \not\vdash p, \text{ then } p \notin K \div p.$$

has been contested on the ground that there may be sentences other than tautologies that an epistemic agent may refuse to withdraw [226]. Operations have been proposed that violate these two postulates. (See Section 6.2)

3.3 Are Belief Sets Too Large?

The use of a logically closed belief set to represent the belief state implies that all beliefs are treated as if they have independent status. Suppose that 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 deselect it. 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 Section 4.1) have largely been constructed in order to make this distinction.

Rott has pointed out that the logical closure of belief sets is unrealistic since it assumes that epistemic agents are ideally competent concerning logic [230]. In the same article he argued that contrary to common assumptions, AGM is not based on a principle of minimal change.

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 $(\{X \mid (\exists p)(X = K \div p)\})$ should be finite [127]. Partial meet contraction does not in general satisfy either of these two finitistic criteria. This has led to the development of finitistic models such as belief base models (Section 4.1) and specified meet contraction (Section 2.6).

3.4 Lack of Information in the Belief Set

Belief sets have also been criticized for lacking important information. Most importantly, AGM contraction or revision in its original form is a “one shot” operation. After contracting K by p with the operation \sim_γ we obtain a new belief set $K \sim_\gamma p$ but no 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 *categorical matching*, according to which the representation of a belief state after a change should have the same format as that of the belief state before the change [114]. In studies of iterated revision, various ways to extend the belief state representation to solve this problem have been investigated. (See Section 5.)

The lack of modal and conditional sentences in belief sets has often been pointed out, but attempts to include them have given rise to severe difficulties. The same applies to introspective beliefs, i.e. the agent’s beliefs about her own belief state [91]. The inclusion of sentences referring to preferences and norms has been somewhat more successful. (See Section 7.9.)

4 Extended Representations of Belief States

Many of the modifications of the AGM model that have been proposed consist in extensions of the belief state representation that make it contain more information in addition to what is contained in the belief set.

4.1 Belief Bases

A belief base is a set B of sentences such that a sentence p is believed if and only if $p \in \text{Cn}(B)$ [184, p. 357]. Operations on belief bases have been extensively investigated [56, 77, 93, 95, 119, 120, 123, 129, 201, 229, 264].

There are two major interpretations of belief bases. One of them, supported by Dalal [56], uses belief bases as mere expressive devices; hence if $\text{Cn}(B_1) = \text{Cn}(B_2)$ then B_1 and B_2 represent the same belief state and yield the same outcome under all operations of change. (“Irrelevance of syntax”, see [202].)

The other, more common approach treats inclusion in the belief base as epistemically significant. The belief base contains those sentences that have an epistemic standing of their own. Suppose that the belief set contains the sentence s , “Shakespeare wrote Hamlet”. Due to logical closure it then also contains the sentence $s \vee d$, “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 [136]. In this approach, belief bases increase the expressive power of the belief representation, since two belief bases with the same logical closure can represent different ways to hold the same beliefs. The two belief bases $\{p, q\}$ and $\{p, p \leftrightarrow q\}$ have the same logical closure but will expectedly behave differently under operations of change; revision by $\neg p$ will result in $\{\neg p, q\}$ respectively $\{\neg p, p \leftrightarrow q\}$ [122]. In this model, changes are made on the belief base, and the merely derived sentences cannot survive when the basis of their derivation is lost. (This has been called a *filtering condition* [95].)

An input-driven operation \circ on a belief base B gives rise to a *base-generated* operation \circ' on $K = \text{Cn}(B)$, such that $K \circ' p = \text{Cn}(B \circ p)$ for all p . Axiomatic characterizations have been obtained of partial meet contraction and revision on belief bases [122, 126] and of the base-generated operations that they give rise to on belief sets [127]. Kernel contraction turns out to be a more general operation than partial meet contraction when applied to belief bases [73, 128].

Belief bases make it possible to distinguish between different inconsistent belief states. This can be used to construct two types of revision operators, depending on whether the negation of the added sentence is contracted before or after its addition:

$$B * p = B \div \neg p + p \text{ (internal revision, Levi identity) [4]}$$

$$B * p = B + p \div \neg p \text{ (external revision, reversed Levi identity) [127]}$$

The recovery postulate does not hold for partial meet contraction on belief bases [121]. Johnson and Shapiro have investigated conditions under which Recovery, or closely related properties, hold in belief base contraction, and argued for the plausibility of some of these conditions [153].

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 [202]. This approach was further developed by Weyderts [266].

Epistemic entrenchment cannot be straightforwardly transferred to a belief base framework. However, Williams showed how contraction can be based on *ensconcement* relations, a related type of transitive and connective relation on belief bases [269]. On *ensconcement*, see also [83].

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 necessary to accommodate new facts. Rules are not removed but can instead be changed. Hence, suppose that the belief base contains the fact $a \& b$ and the two rules $a \rightarrow c$ and $b \rightarrow c$. After revision by the new fact $\neg c$, a new belief base is obtained that contains the facts $a \& b$ and $\neg c$ and the two rules $(a \& \neg b) \rightarrow c$ and $(b \& \neg a) \rightarrow c$ [64].

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 [27].

It is commonly assumed that the belief base approach corresponds to foundationalist epistemology, whereas the original AGM framework represents a coherentist view [110, 255]. Doyle accepted this view but argued that the epistemic concern for conservatism that underlies the coherentist view applies equally to the foundations approach [66]. In [143] it was argued that the original AGM approach is incompatible with important characteristics of coherentism. In [134] coherentism was expressed in a framework where a belief base B is assigned to the belief set K . Then a logically closed subset K' of K is coherent if and only if there is some sentence p such that $K' = \text{Cn}(B \sim_{\gamma} p)$.

4.2 Probability and Plausibility

The AGM model represents 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 [234]. It seems difficult to construct a reasonably manageable model that covers both the logic-related and the probabilistic properties of belief change [165, 183].

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 [176]. Makinson further investigates this and other connections between the two frameworks [191]. Insights from AGM can be used as an impetus for considering accounts of conditional probability in which $\text{p}(q, r)$, the probability of q given r , is not defined in the standard way. Furthermore, the notion of non-prioritized revision (see Section 6.2) can be transferred to a probabilistic context where it corresponds to “vacuous” conditionalizing when the condition is too unbelievable to be taken seriously.

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.” [30]. He 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 $p(A)$ in $[0, 1]$, representing its probability, they assign to it an element $Pl(A)$ of a partially ordered set. $Pl(A)$ is called the “plausibility” 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 operator that is essentially equivalent with partial meet revision [90, 92].

Several other authors have presented probability-based and plausibility-based belief revision models that have close connections with the AGM model [10, 13, 18, 69].

4.3 Ranking Models

In 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 [249–251]. A sentence p is assigned the value $\kappa(p) = \min\{\kappa(w) \mid p \text{ holds in } w\}$. 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 $p \rightarrow x$ -conditionalization of κ is defined by: $\kappa_{p \rightarrow x}(q) = \min\{\kappa(q \mid p), \kappa(q \mid \neg p) + x\}$

Contractions, expansions and revisions can all be represented as conditionalizations (depending 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 [147]. A generalization of Spohn’s ranking functions has been proposed by Weydert [267].

4.4 Extensions of the Language

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-truthfunctional expressions into the language has interesting and often surprisingly drastic effects.

Among the several formal interpretations of non-truthfunctional *conditionals*, such as counterfactuals, the Ramsey test is particularly well suited to belief revision. 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 $p \square \rightarrow q$ denote “if p then q ”, or more precisely: “if p were the case, then q would be the

case". The Ramsey test says that $p \Box \rightarrow q$ holds if and only if $q \in K * p$ [253]. In order to treat conditional statements like $p \Box \rightarrow q$ on par with statements about actual facts, they will have to be included in the belief set when they are assented to by the agent, thus: $p \Box \rightarrow q \in K$ if and only if $q \in K * p$.

However, inclusion in the belief set of conditionals that satisfy the Ramsey test requires radical changes in the logic of belief change. Hence, contraction cannot then satisfy the inclusion postulate ($K \div p \subseteq K$) since contraction typically provides support for conditional sentences that were not supported by the original belief state. If I give up my belief that John is mentally retarded, then I gain support for the conditional sentence "If John has lived 30 years in London, then John understands the English language" [124].

A famous impossibility theorem by Gärdenfors shows that the Ramsey test is incompatible with a set of plausible postulates for revision [108]. The crucial part of the proof consists in showing that the Ramsey test implies the following monotonicity condition: If $K \subseteq K'$ then $K * p \subseteq K' * p$. 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 is to reject the Ramsey test as a criterion for the validity of conditional sentences [222]. Another, proposed by Levi, is to accept the test as a criterion of validity but deny that such conditional sentences should be included in the belief set when they are valid [172]. 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 [9, 11]. 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 $p \Box \rightarrow q$ and a belief set K , there is some sentence $r_{p \Box \rightarrow q}^K$ such that $p \Box \rightarrow q$ holds in the belief state represented by K if and only if $r_{p \Box \rightarrow q}^K \in K$. In this way we can have the Ramsey test in the following form that is not blocked by the impossibility result [178, 179]: $r_{p \Box \rightarrow q}^K \in K$ if and only if $q \in K * p$. 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 [124, 223]. Ryan and Schobbens have related the Ramsey test to update rather than revision (cf. Section 6.1) and found the test to be compatible and indeed closely connected with update operators [238].

Kern-Isberner has proposed a framework in which conditional sentences can be elements of belief sets, and revisions can be performed with conditional sentences as inputs [158]. A partly similar approach has been developed by Weydert [268].

The inclusion of *modal sentences* in belief sets has been investigated by Fuhrmann. Let $\diamond p$ denote that p is possible, and let $\diamond p \in K$ hold if and only if $\neg p \notin K$. This seemingly reasonable definition gives rise to problems similar to those exhibited in Gärdenfors's theorem, and essentially the same types of solutions have been discussed [94].

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 dissimilar to those for conditionals are obtained in this case as well [180]. Similar results were obtained by Friedman and Halpern [91].

Dupin de Saint-Cyr and Lang introduced *temporally labelled sentences* into belief revision and proposed a belief change operator, called belief extrapolation, in which predictions are based on initial observations and a principle of minimal change [70]. Bonanno has developed logics that contain both a next-time temporal operator and a belief operator [31, 32].

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 [42].

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 [96, 99].

4.5 Changes in Norms, Preferences, Goals, and Desires

Norms Although the AGM model was partly the outcome of attempts to formalize changes in norms [5], authors who tried to apply the AGM model to norms have found it in need of rather extensive modifications to make it suitable for that purpose. Boella et al. analyzed normative change in a framework with norms represented by sets of pairs $\langle p, q \rangle$, to be read "if p , then it is obligatory that q ". In this framework, close analogues of the AGM postulates give rise to inconsistency [29]. Governatori and Rotolo proposed a model for changes in legislation that among several other aspects also includes an explicit representation of time in order to account for phenomena such as retroactivity [116]. 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. The problem of how to prioritize among conflicting items is similar to the selection of sentences for removal [142].

Preferences A model of changes in preferences can be obtained by replacing the standard AGM language by 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. Partial meet contraction can be used, but some modifications of the AGM model seem to be necessary in applications to preferences [118, 130, 166].

Goals and Desires Intention selection is a process aimed at removing contradictions, to end up with a consistent set of intentions [208]. Paglieri and Castelfranchi have proposed a model of Data-oriented Belief Revision (DBR) in which attention is paid to the mutual influences between beliefs and goals [49–52, 207]. Boella et al. have also analyzed the role of goals in belief revision [28]. Their model is similar to DBR in its selection criteria, but it puts more emphasis on avoiding wishful thinking.

5 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 probably the most studied problem in the literature on belief change.

5.1 Revising Epistemic States

In order to solve the problem of iterated change we need a belief state representation that contains more information than the belief set, so that it can guide additional changes. Furthermore, the operation of change has to yield a complete such belief state representation as its outcome, not merely a new belief set. There are several ways to represent such an extended epistemic state. The most common of these is a preorder on the set of possible worlds, or equivalently a complete sphere system (cf. Section 2.4). The belief set can be inferred from this preorder; it is simply the intersection of the worlds in the highest equivalence class (the innermost sphere). An operation of change gives rise to a new preorder (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 [57]. To present it, let Ψ be the current belief state of the agent, and \leq_{Ψ} the preorder that represents it. Similarly, let $\Psi \circ p$ be the belief state obtained after revising by p , and $\leq_{\Psi \circ p}$ the preorder that represents it. (Following tradition we will focus on revision; relatively little has been written on iterated contraction.)

Darwiche and Pearl proposed the following conditions for iteration (μ and ϕ are possible worlds):

- (DP1): For any $\mu \models p$ and $\phi \models p$, $\mu \leq_{\Psi} \phi$ iff $\mu \leq_{\Psi \circ p} \phi$.
- (DP2): For any $\mu \models \neg p$ and $\phi \models \neg p$, $\mu \leq_{\Psi} \phi$ iff $\mu \leq_{\Psi \circ p} \phi$.
- (DP3): For any $\mu \models p$ and $\phi \models \neg p$, if $\mu <_{\Psi} \phi$, then $\mu <_{\Psi \circ p} \phi$.
- (DP4): For any $\mu \models p$ and $\phi \models \neg p$, if $\mu \leq_{\Psi} \phi$, then $\mu \leq_{\Psi \circ p} \phi$.

According to (DP1), the order among the p -worlds remains unchanged after revision by p . According to (DP2) the order among the $\neg p$ -worlds remains

unchanged after revision by p . (DP3) says that if a p -world is strictly preferred to a $\neg p$ -world, then that strict preference is maintained after revision by p . (DP4) says that if a p -world is weakly preferred to a $\neg p$ -world, then that weak preference is maintained after revision by p .

These conditions have been shown to correspond to the following postulates for iterated revision: [57, Theorem 13.]

(DP1) If $q \vdash p$, then $(\Psi \circ p) \circ q = \Psi \circ q$.

(DP2) If $q \vdash \neg p$, then $(\Psi \circ p) \circ q = \Psi \circ q$.

(DP3) If $\Psi \circ q \vdash p$, then $(\Psi \circ p) \circ q \vdash p$.

(DP4) If $\Psi \circ q \not\vdash \neg p$, then $(\Psi \circ p) \circ q \not\vdash \neg p$

These four postulates have become the benchmark for iterated revision, and new proposals are almost invariably compared to them. However, Jin and Thielscher [149] and Booth and Meyer [37] have pointed out that these postulates are too permissive since they do not rule out operators by which all newly acquired information is given up as soon as an agent learns a fact that contradicts some of its current beliefs. To avoid this they proposed the following additional condition:

(Ind): For any $\mu \models p$ and $\phi \models \neg p$, if $\mu \leq_{\Psi} \phi$, then $\mu <_{\Psi \circ p} \phi$.

5.2 Major Classes of Iterable Operators

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 operators into three classes according to their ability to remember and to take the revision history into account:

Operators Without Memory In this case, each belief set is revised in a pre-determined way, independently of how it was obtained, i.e.:

If $\Psi \circ p$ and $\Upsilon \circ p$ have the same belief set, then so have $\Psi \circ p \circ r$ and $\Upsilon \circ p \circ r$.

Full meet revision is a trivial example of an iterable operator without memory. Areces and Becher have analyzed this class of operators [8].

Operators with Full Memory In this case the full history of changes is conserved, so that rollbacks of previous changes are possible. Operators with full memory have been proposed, for example, by Brewka [46], by Lehmann [169], and by Konieczny and Pérez [160]. Falappa et al. proposed another type of revision in which discarded beliefs can be reused [74].

Operators 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 identify the previous states. Most of the proposed iterable revision operators are of this type. In a recent review Rott recognized three major types of iterable revision operators. They all have partial memory [235]:

Conservative revision, originally called natural revision, has been studied by Boutilier [44, 45] and Rott [233]. 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 top of the preorder which is otherwise left unchanged. The main characteristic of this operator is:

$$(Nat): \quad \text{If } \mu \notin [\Psi \circ p] \text{ and } \phi \notin [\Psi \circ p], \text{ then } \mu \preceq_{\Psi} \phi \text{ iff } \mu \preceq_{\Psi \circ p} \phi.$$

Moderate revision, also called lexicographic revision, was originally studied by Nayak [198] and by Nayak et al. [200]. When revising by p it rearranges the preorder by putting the p -worlds at top (but conserving their relative order) and the $\neg p$ -worlds at bottom (but conserving their relative order). It has the following property.

$$(Lex): \quad \text{If } \mu \models p \text{ and } \phi \models \neg p, \text{ then } \mu <_{\Psi \circ p} \phi.$$

Radical revision is similar to moderate revision, but it differs in making the new belief irrevocable, i.e., impossible to remove. Segerberg proposed this type of revision and characterized it axiomatically [246]. It is further investigated in [80]. In radical revision by p , 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 operator is:

$$(Irr): \quad [(\Psi \circ p) \circ \neg p] = \emptyset.$$

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 [61]. Several other types of iterable operators have been proposed, see for instance [34, 41, 47, 151, 159, 204].

6 Alternative Operators of Change

In the original AGM model there are three major types of operators: contraction, revision, and expansion. Subsequently a large number of additional types of operators have been proposed.

6.1 Update

In 1992, Katsuno and Mendelzon presented a type of operator of change that they called update. Katsuno and Mendelzon [156] whereas revision operators are suited to capture changes that reflect evolving knowledge about a static

situation, update operators are intended to represent changes in beliefs that result from changes in the objects of belief [272]. There are important formal differences between update and AGM revision; in particular the AGM postulate Vacuity (If $K \not\vdash \neg p$ then $K + p \subseteq K * p$) does not hold for updates. Update and its relation with revision have been further studied by Becher [23] and others.

6.2 Non-prioritized Change

In AGM revision, new information has primacy. According to the success postulate for revision, 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. Belief revision that violates the success postulate is called non-prioritized belief revision. Revision operators can be classified according to their outcomes into the following five categories, of which the first is prioritized and the other four non-prioritized [79]. We assume that the input sentence p is consistent.

All. The input is accepted without any constraint: $p \in K \circ p$.
Satisfied by AGM revision, external revision [126], and update [156].

All or nothing. Either the new input is accepted, or the belief set is left unchanged: $p \in K \circ p$ or $K \circ p = K$.
Satisfied by screened revision [188], credibility-limited revision [141], and some applications of an improvement operator [161].

All or inconsistency. The input is accepted, but its negation may survive in the belief set: $p \in K \circ p \neq \text{Cn}(\{\perp\})$ or $K \circ p = \text{Cn}(\{\perp\})$.
Satisfied by irrevocable belief revision [242] and also by revision by comparison in the collapsed case [88].

All or less. Either the input is accepted or it induces a contraction process: $p \in K \circ p$ or $K \circ p \subseteq K$.
Satisfied by semi-revision [131] and revision by comparison in the unsuccessful case [88].

All or a part. Either the whole input or a part of it is accepted. There is then some q such that $K * p \vdash q$, $\vdash p \rightarrow q$ and $K * p = K * q$.
Satisfied by selective revision [81], Rabinowicz's and Schlechta's revision [216, 241], and Lin's revision [174].

For an overview with emphasis on 'all or nothing' constructions, see [132].

The success postulate for contraction requires that all non-tautological beliefs are retractable. This is not a fully realistic requirement, since actual agents are known to have beliefs of a non-logical nature that nothing can bring them to give up. In *shielded contraction*, the success postulate does not hold in general; some non-tautological beliefs are shielded from contraction and

cannot be given up. Shielded contraction has close connections with credibility-limited revision [82, 84, 188].

Chopra, Ghose and Meyer have extended the framework of non-prioritized belief revision by introducing non-beliefs as possible inputs that are ranked along with beliefs [53].

6.3 Changes in the Strength of Beliefs

An operation of change can raise or lower the position of a sentence in the ordering without affecting the belief set (but affecting how the belief state responds to new inputs). Hence an operator of improvement, as proposed by by Konieczny and Pérez [161] increases the plausibility of p by moving some of the p -worlds to a higher position in the preorder. This can facilitate the acceptance of p in later, additional operations, so that we can have $p \notin K \circ p$ but $p \in K \circ p \circ p$.

In quantitative theories of belief change, such as probabilistic and ranking theories, the degree of acceptance of each sentence is represented by a numerical value. 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 [154]. These difficulties are largely avoided if an operation of change is constructed to adjust the position of an input sentence in a ordering to be the same as that of a reference sentence. Since two sentences are involved, Rott called such operators two-dimensional [235]. This kind of operation was studied by Cantwell, who 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 [47]. Fermé and Rott proposed the operation of *revision by comparison*. It accepts the input sentence p to the same degree as a previously believed sentence q , unless the negation of the input sentence p is more plausible than the reference sentence q , in which case q will be removed from the outcome [88]. Revision by comparison violates the DP postulates since it collapses distinctions between 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 [236].

6.4 Resource-bounded Change and Inconsistency Management

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 [264]. Harman has put forward a highly useful list of principles that should be valid for any resource-bounded agent [145]. As was observed by Gabbay and Hunter, for a real agent it may not be necessary to restore consistency, but it may be sufficient to have rules that specify how to act when an inconsistency arises [101]. The two features of resource-boundedness that have attracted

most attention among researchers are finitude and inconsistency tolerance. Both belief bases and specified meet contraction have been constructed largely in order to avoid the infinite structures of the standard AGM model.

Doyle has investigated characteristics of real agents such as mental inertia and constitutional elasticity [67]. His *reason maintenance system* (RMS) is intended to capture these characteristics. Alechina et al. used RMS to construct a resource-bounded operator of contraction [7].

In the AGM model there is only one inconsistent belief set, namely the whole language. Belief bases fare much better in this respect; there are many different inconsistent belief bases that can represent different inconsistent states [120]. This feature of belief bases was employed in Hansson's and Wasserman's model of *local change* [144]. Given a belief base B and a sentence r , the r -compartment of B is the subset of B that is relevant for r . 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 also investigated a construction with a short-term memory in which recently computed results are temporarily stored [263]. Local change can be used for diagnosis [265], i.e. the process of finding the faulty compartment of a malfunctioning system [218].

In a similar vein, Parikh [212] proposed a principle for relevance-sensitive change according to which, if the belief set can be split into two independent parts (expressed in different sublanguages), then revision of one part does not affect the other. Peppas provided a semantics for this principle in terms of systems of spheres [213]. Kourousias and Makinson have investigated conditions under which Parikh's relevance-sensitive condition is satisfied [164, 190].

Another approach to inconsistencies is to use a paraconsistent logic (i.e. a logic where $\{p \wedge \neg p\} \not\sim q$ does not hold in general) [215, 254].

6.5 Multiple Change

In the original AGM model the input is a single sentence. In models of multiple change, the input is a set of sentences. Arguments have been given why no two of the four contractions $K \div p \div q$, $K \div q \div p$, $K \div \{p, q\}$, and $K \div (p \vee q)$ should be expected to coincide in general [100]. However, it has been argued that the order-dependence of iterated operations is in itself problematic since ideally we should treat the pieces of information that we receive on an equal footing, independently of the order in which they arrive [61, 140].

Fuhrmann and Hansson identified two types of multiple contraction. In *package contraction* all members of an input set A are removed, hence $(K \div A) \cap A = \emptyset$. In *choice contraction* it is only required that at least one member of the input set be removed, i.e. $A \not\subseteq K \div A$ [100]. Zhang has added a third option, *set contraction*. Its purpose is not to remove the input but to make the outcome compatible with the input. Hence, the outcome $K \ominus A$ satisfies the property $(K \ominus A) \cup A \vDash \perp$ [273].

Most of the major AGM-related contraction operators have been generalized to multiple contraction: multiple partial meet contraction [100, 119, 173],

multiple kernel contraction [89], multiple specified meet contraction [140], and a multiple version of Grove's sphere system [85, 217]. Spohn has proposed a ranking-theoretic account of multiple package contraction [252]. Fuhrmann has investigated the subtraction $p - q$ that asserts p with the exception of what q says, and the merge operation $p \circ q$ that extracts the maximized consistent content from p and q jointly [97, 98]. Finally, Zhang has investigated operations that are both iterated and multiple [274].

6.6 Indeterministic Change

The AGM model and most other models of belief change are deterministic in the sense that the outcome is always well-determined. There is no scope for chance in determining the outcome of the change. In indeterministic belief change, an operation can have more than one admissible outcome. Indeterministic belief change has been studied by Gallier [105] and by Lindström and Rabinowicz [177]. The latter authors gave up the assumption that epistemic entrenchment satisfies connectedness. This resulted in Grove's sphere systems with "fallbacks" that are not linearly ordered but still all include the original belief set.

6.7 Some Other Operators of Change

Consolidation An inconsistent belief base can be consolidated, i.e. made consistent by removing some of its elements. The consolidation of B is denoted $B!$. A plausible way to perform consolidation is to contract by falsum (contradiction), i.e. $B! = B \div \perp$ [131]. In Olsson's *coherence consolidation*, enough of an incoherent belief base is removed to make it coherent [205].

Replacement replaces one sentence by another. Such an operator has two variables, such that in $K|_q^p$, p has been replaced by q . 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, since contraction, revision, and expansion can all be defined in terms of it. ($K \div p = K|_T^p$, $K * p = K|_p^\perp$, and $K + p = |_p^\top$.) Partial meet replacement has been axiomatically characterized, and it also has a semantic account in terms of possible worlds [139].

Reconsideration introduced by Johnson and Shapiro [151, 152] is a non-prioritized operation on belief bases. It represents changes that are performed in hindsight in order to eliminate negative effects caused by previously performed changes, such as the reintroduction of previously removed beliefs. This operation can be produced by an algorithm that examines a subset of the retracted basic beliefs using dependency relationships.

Multiagent Belief Change AGM operations can be extended to cover changes performed jointly by more than one agent, who are assumed to combine

their beliefs in order to obtain a common, consistent set of beliefs. Konieczny and Perez have proposed an operator called merging that generalizes several previously proposed methods for conflict-solving combinations of information from two agents [162, 163]. For more than two agents, early work by Revesz [219] and Lin [174], based on distances between models, has been followed by several other approaches [19, 68, 182, 193, 194].

7 Applications and Connections

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

7.1 Non-Monotonic and Defeasible Logic

In spite of the differences between belief revision and non-monotonic logic [111] it is possible to translate concepts, models, and results between these areas [192]. Non-monotonic reasoning can be expressed by an inference operator \sim such that $A \sim p$ denotes that A is a good enough reason to believe that p , or that p is a plausible consequence of A . However, contrary to Cn , C does not satisfy monotony (If $A \subseteq B$ then $\text{Cn}(A) \subseteq \text{Cn}(B)$). Given a belief set K representing the background beliefs we can translate formulas between the two frameworks as follows:

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

Belief revision and non-monotonic logic have become increasingly interconnected. Satoh provided a model of belief change (“minimal revision”) that satisfies the AGM postulates except Vacuity and Conjunctive overlap. It can be interpreted as a non-monotonic reasoning operator [240]. Lindström [175] and Rott [227] 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 [186, 224]. Billington et al. further clarified the relationships between AGM revision and non-monotonic inference [26].

In the last years of his life Alchourrón published a series of articles on the logic of defeasible conditionals [1–3]. 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 its presuppositions. The connection between this approach and AGM was initially somewhat unclear [24], but in [87] axioms were given that relate it to a generalized version of AGM revision for an implicit underlying belief set K . This 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 [211].

7.2 Description Logic

Description logics have been successful in detecting incoherences in databases, but provide little support for resolving these incoherences. 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. Using proposals by Benferhat et al. [25] and Meyer et al. [196] developed strategies for solving incoherences in a framework based on description logic. Ribeiro and Wassermann proposed a belief revision approach to finding and repairing inconsistencies in ontologies represented in description logics [220].

7.3 Horn Clause Contraction Functions

A Horn clause is a disjunction of one positive literal and at least one negated literal, such as $p_1 \vee \neg p_2 \vee \neg p_3 \vee \neg p_4$. With few exceptions [77, 221], studies of Horn clause belief change started only recently. Delgrande [60] investigated the contraction of theories expressed in propositional Horn logic and proposed two approaches that he called entailment-based and inconsistency-based [60]. For each of these he proposed a special kind of partial meet (and maxichoice) contraction. Delgrande and Wassermann [63] found that the application of contraction operators to remainder sets (as in AGM) has undesirable properties in Horn clause logic. Instead they developed an account of maxichoice Horn contraction that operates on weak remainder sets (defined semantically instead of syntactically).

Booth et al. provided a generalization of Delgrande's partial meet constructions and proposed that the latter are only a subset of the appropriate constructions of contraction in Horn logic [39] (cf. [100]). Booth et al. [40] also showed that their construction corresponds exactly to standard kernel contraction, as applied to Horn clauses. For further work on contraction in Horn logic, see [167, 276].

7.4 Game Theory

Belief change and game theory are related in several ways. Booth and Meyer [38] investigated equilibria in belief merging. The key idea is that a social belief removal function can be 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 [36]. Zhang studied bargaining from another viewpoint, proposing a logical axiomatization of bargaining solutions that is based on postulates from AGM and game theory [275].

The classical analysis of centipedes in the game theoretical literature (by Aumann, Binmore and others) shows that adequate models should be able to represent belief-contravening hypotheses and therefore also the notion

of supposition. Samet [239] proposed that hypothetical knowledge can be represented by the operator: $K_i^H(E) = \bigcup \{P \in \Pi_i: T_i(P, H) \subseteq E\}$, where Π_i is a partition cell and T_i is a transformation function mapping partition cells and hypotheses to partition cells. This transformation function has the flavour of belief change but it is constrained by properties not used in the literature on belief change. Arló-Costa and Bicchieri proposed an operator of this sort that satisfies some of the AGM properties. They 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 choose the backward induction solution in centipede games. This result was obtained without assuming that rationality is commonly known (as in [20]) or commonly hypothesized by the players (as in [239]).

7.5 Argumentation

Falappa et al. [76] have argued that belief revision and argumentation theory are complementary approaches. By combining the two, the variety and complexity of reasoning processes is better accounted for than if only one of them is used. In [75], they combine ATMS (Assumption-based truth-maintenance systems) [58] with belief revision and 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 A of reasons to believe it. A partial acceptance revision operator is constructed such that A is initially accepted, which leads to the creation of $B \cup A$ as a (possibly inconsistent) intermediate belief base from which inconsistencies are removed, giving rise to a consistent revised belief base $B \circ A$. The operator \circ is an operator of external revision in the sense explained in Section 4.1.

Paglieri and Castelfranchi proposed Data-oriented Belief Revision (DBR) as an alternative to AGM [50, 208]. This model combines belief revision with argumentation, following Toulmin's account of argumentation. The application of DBR to argumentation is primarily intended to highlight structural communalities between arguments and belief-supporting networks [209, 210]. The model contains two basic informational categories, data and beliefs. Contrary to beliefs, data are allowed to be contradictory. When a belief is abandoned, this does not entail removal of the corresponding data from the agent's memory, i.e. disbelieving is not forgetting.

7.6 Modal and Dynamic Logics

We can distinguish between three ways to integrate belief revision with a modal logic.¹ First, we can add epistemic or doxastic modal operators to make it explicit in the logical language that the belief set consists of *beliefs*. Secondly,

¹This section relies heavily on personal communications from Hans van Ditmarsch.

we can formalize the execution of expansion, revision, and contraction with dynamic modal operators, similar to those for program execution. Thirdly, we can add both epistemic and dynamic modal operators, to integrate belief and belief change in one language; this has been done in dynamic doxastic logic by Segerberg and collaborators and largely separately from that, later 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. $\Diamond \Box B$ 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 [129].

Dynamic Modalities for Belief Revision In various publications, van Benthem, de Rijke, and Fuhrmann [59, 95, 256, 257] 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 formulas with both Hintikka-style doxastic operators [148] and dynamic modal operators for belief change [243–245]. It was defined by Krister Segerberg as a logical framework for reasoning about doxastic change. The basic DDL represents an agent that has 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 this model in order to include such formulas [181]. The extended model allows not only for introspective agents but also for iterated change, which has been studied by John Cantwell [47].

Dynamic Epistemic Logic (DEL) also studies changes in information and investigates actions with epistemic impact on agents [21, 214, 262]. Like DDL it has epistemic operators for belief or for 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 $p \wedge \neg Bp$ [148, 197]. 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 [260].

Belief revision in the typical AGM sense is more problematic. If the agent believes in p (i.e., Bp is true), then a public announcement of $\neg p$ will make her or his beliefs inconsistent. To model belief revision in DEL, we need Kripke models where knowledge, belief, and degrees of belief (or conditional belief) can all be encoded. To achieve this we can add *plausibility relations* to the Kripke models, and identify belief in ϕ with truth of ϕ in the most plausible of the epistemically accessible states.

For example, suppose that two states s and t are both 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 . 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 is when the state s is eliminated from consideration (a ‘hard’ update, as for the execution of public announcements), after which the t is the most plausible (namely the only remaining) state. Again, the agent believes that p is false. These issues were addressed by Aucher [17], van Benthem [258], van Ditmarsch and Labuschagne [259, 261], and Baltag and Smets [22] (with many follow-up papers). Conditional reasoning and reasoning with different degrees of belief can also be modelled in such settings.

7.7 Belief Change by Translation Between Logics

The AGM paradigm can be extended to non-classical logics by translating the source logic into classical logic, performing the change and then translating back. This method was proposed independently by two groups. Gabbay et al. [102, 103] respectively Coniglio and Carnielli [54]. The latter defined logics as two-sorted first-order structures, and argued that this broad definition encompasses a wide class of logics with theoretical interest as well as interest 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. They define a model of belief change,

called Wide Belief Revision Systems, to define belief revision for non-standard logics.

7.8 Truth

In eliminative induction a number of possible hypotheses concerning some state of affairs are presumed, and rivals are progressively eliminated by new evidence. The process is an idealization, since in practice no closed set of initial theories is usually available. Kelly found similarities between the process of eliminative induction and theory change, and proposed the use of belief revision to arrive at informative, true, empirical belief on the basis of increasing information. He analyzed several algorithms of iterated revision (see Section 5), assuming that learning is based on the outcomes of sequential experiments [157].

7.9 Use of Choice Functions and Related Preference Orderings

The classical theory of rational choice was developed by mathematical economists [15, 16, 247]. It occupies a central role in the philosophy of the social sciences.

Choice rationality is concerned with how to choose rationally among a set of alternatives. Formal requirements are imposed on choices from different, overlapping alternative sets. The standard presumption is that selection functions are rationalizable in terms of some underlying binary preference relation. Formally, this is represented in structures containing a set of alternatives Ω and a function f that takes us from any subset E of Ω to a subset of E . The main objective of rational choice theory is to investigate the conditions under which the function f can be rationalized by a total pre-order R on Ω in the sense that, for every $E \subseteq \Omega$, $f(E)$ coincides with the best elements of E according to R . However, as discussed by Sen [248], social norms and menu dependence can make rationalizability in terms of an underlying preference relation impossible.

The structures studied in rational choice theory have a close connection with the selection functions employed in belief revision. In his book [231], 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.

Olsson [206] conceded 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 revision and choice. Instead he has reconnected the AGM theory with its roots in counterfactual reasoning and rational choice.

Further studies in this area have been reported by Arló-Costa and Pedersen [14] and by Bonanno [33].

8 Computability and Implementation

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 1980's several algorithms for the implementation of belief change operations were proposed. Most of them were constructed to recognize 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, [65], Borgida [43], Winslett [272], Dalal [56] and Satoh [240]. Katsuno and Mendelzon [155] 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) [195], and the second to Gabbay et al. [104]. Recent works on implementation also include proposals by Williams [270], Williams and Sims [271], and Delgrande and Schaub [62].

A core aspect in implementation is the space and time required for computation. One of the first studies of the cost of belief change algorithms was performed by Eiter and Gottlob [71]. 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 they are located on the second level of the polynomial hierarchy” [203].

However, specific algorithms can reduce this computational cost. The determination of what properties have to be resigned in order to gain efficiency is a major challenge for future studies. An interesting step was taken by Jin and Thielscher who proposed a model called Reinforcement Belief Revision that 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 [150].

References

1. Alchourrón, C. (1993). 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). Chichester: Wiley & Sons.
2. Alchourrón, C. (1995). Defeasible logic: Demarcation and affinities. In C. Press (Ed.), *Conditionals: From philosophy to computer science* (pp. 67–102). Oxford: Crocco, Fariña del Cerro.
3. Alchourrón, C. (1996). Detachment and defeasibility in deontic logic. *Studia Logica*, 51, 5–18.
4. Alchourrón, C., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50, 510–530.

5. Alchourrón, C., & Makinson, D. (1981). 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).
6. Alchourrón, C., & Makinson, D. (1985). On the logic of theory change: Safe contraction. *Studia Logica*, 44, 405–422.
7. Alechina, N., Jago, M., & Logan, B. (2006). 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, 25 July 2005, Selected and revised papers. Lecture notes in computer science* (Vol. 3904, pp. 141–154).
8. Areces, C., & Becher, V. (2001). Iterable AGM functions. In H. Rott, & M.-A. Williams (Eds.), *Frontiers in belief revision. Applied logic series* (pp. 261–277). Kluwer Academic Publishers.
9. Arló-Costa, H. (1995). Epistemic conditionals, snakes and stars. In A. H. Series editor: D. M. Gabbay, L. Fari nas del Cerro, G. Crocco (Ed.), *Conditionals: From philosophy to computer science. Studies in logic and computation* (Vol. 5, pp. 193–239). Oxford: Oxford University Press.
10. Arló-Costa, H. (2001). Bayesian epistemology and epistemic conditionals: On the status of the export-import laws. *Journal of Philosophy*, XCVIII(11), 555–598.
11. Arló-Costa, H., & Levi, I. (1996). Two notions of epistemic validity. *Synthese*, 109(2), 217–262.
12. Arló-Costa, H., Levi, I. (2006). Contraction: On the decision-theoretical origins of minimal change and entrenchment. *Synthese*, 152(1), 129–154.
13. Arló-Costa, H., & Parikh, R. (2005). Conditional probability and defeasible inference. *Journal of Philosophical Logic*, 34, 97–119.
14. Arló-Costa, H., & Pedersen, A. P. (2011). Social norms, rational choice and belief change. In E. Olsson (Ed.), *Belief revision meets philosophy of science* (pp 163–212). Springer.
15. Arrow, K. J. (1951). *Social choice and individual values*. New York: Wiley.
16. Arrow, K. J. (1970). *Collective choice and social welfare*. Holden-Day.
17. Aucher, G. (2003). *A combined system for update logic and belief revision*. Master's thesis, ILLC, University of Amsterdam, Amsterdam, The Netherlands. ILLC report MoL-2003-03.
18. Aucher, G. (2007). Interpreting an action from what we perceive and what we expect. *Journal of Applied Non-Classical Logics*, 17(1), 9–38.
19. Aucher, G. (2010). Generalizing AGM to a multi-agent setting. *Logic Journal of the IGPL*, 18(1), 530–559.
20. Aumann, R. J. (1992). 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). Cambridge: MIT Press.
21. Baltag, A., Moss, L., & Solecki, S. (1998). 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).
22. Baltag, A., & Smets, S. (2006). 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).
23. Becher, V. (1999). *Binary functions for theory change*. Ph.D. thesis, University of Buenos Aires.
24. Becher, V., Fermé, E., Lazzar, S., Oller, C., Palau, G., & Rodríguez, R. (1999). Some observations on Carlos Alchourrón's theory of defeasible conditionals. In M. N. P. Prakke (Ed.), *Norms, logics and information systems. New studies on deontic logic and computer science* (pp. 219–230). Amsterdam/Tokyo/Washington DC: Ios Press.
25. Benferhat, S., Kaci, S., Le Berre, D., & Williams, M.-A. (2004). Weakening conflicting information for iterated revision and knowledge integration. *Artificial Intelligence*, 153(1–2), 339–371.
26. Billington, D., Antoniou, G., Governatori, G., & Maher, M. J. (1999). Revising non-monotonic theories: The case of defeasible logic. In *KI* (pp. 101–112).
27. Bochman, A. (2001). *A logical theory of nonmonotonic inference and belief change*. Springer.

28. Boella, G., da Costa Pereira, C., Pigozzi, G., Tettamanzi, A., & van der Torre, L. (2010). The role of beliefs in goal dynamics: Prolegomena to a constructive theory of intentions. *Logic Journal of IGPL*, 4(18), 559–578.
29. Boella, G., Pigozzi, G., & van der Torre, L. (2009). Normative framework for normative system change. In S. Decker, J. S. Sichman, & C. Castelfranchi (Eds.), *Proc. of 8th int. conf. on autonomous agents and multiagent systems (AAMAS 2009)* (pp. 169–176). Budapest.
30. Bonanno, G. (2005). A simple modal logic for belief revision. *Synthese*, 147, 193–228.
31. Bonanno, G. (2007). Axiomatic characterization of the AGM theory of belief revision in a temporal logic. *Artificial Intelligence*, 171, 144–160.
32. Bonanno, G. (2007). Temporal interaction of information and belief. *Studia Logica*, 86, 381–407.
33. Bonanno, G. (2009). Rational choice and AGM belief revision. *Artificial Intelligence*, 173, 1194–1203.
34. Bonanno, G. (2009). Revealed preference, iterated belief revision and dynamic games. In G. Bonanno, J. Delgrande, & H. Rott (Eds.), *Information processing, rational belief change and social interaction, Dagstuhl, Germany*.
35. Booth, R., Chopra, S., Meyer, T., & Ghose, A. (2004). A unifying semantics for belief change. In L. de Mántaras R., & S. L. (Eds.) *Proceedings of the 16th European conference on artificial intelligence (ECAI 2004)* (pp. 793–797).
36. Booth, R., Chopra, S., Meyer, T., & Ghose, A. (2010). Double preference relations for generalised belief change. *Artificial Intelligence*, 174(16–17), 1339–1368.
37. Booth, R., & Meyer, T. (2006). Admissible and restrained revision. *Journal of Artificial Intelligence Research*, 26, 127–151.
38. Booth, R., & Meyer, T. (2010). Equilibria in social belief removal. *Synthese*, 26, 1–27.
39. Booth, R., Meyer, T., & Varzinczak, I. J. (2009). Next steps in propositional horn contraction. In *Proceedings of the international joint conference on artificial intelligence* (pp. 702–707).
40. Booth, R., Meyer, T., Varzinczak, I. J. & Wassermann, R. (2010). A contraction core for horn belief change: Preliminary report. In *Proceedings of the 13th international workshop on non-monotonic reasoning (NMR), Toronto*.
41. Booth, R., Meyer, T., Wong, K.-S. (2006). 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).
42. Booth, R., & Richter, E. (2001). On revising fuzzy belief bases. *Studia Logica*, 68, 1–30.
43. Borgida, A. (1985). Language features for flexible handling of exceptions in information systems. *ACM Transactions on Database Systems*, 10(4), 565–603.
44. Boutilier, C. (1993). Revision sequences and nested conditionals. In *Proc. 13th Int. joint conf. on artificial intelligence (IJCAI'93)* (pp. 519–525).
45. Boutilier, C. (1996). Iterated revision and minimal change of conditional beliefs. *Journal of Philosophical Logic*, 25, 263–305.
46. Brewka, G. (1991). Belief revision in a framework for default reasoning. In A. Fuhrmann, & M. Morreau (Eds.), *The logic of theory change* (pp. 206–222). Berlin.
47. Cantwell, J. (1997). On the logic of small changes in hypertheories. *Theoria*, 63, 54–89.
48. Carnota, R., & Rodríguez, R. (2011). The AGM impact in artificial intelligence. In E. Olsson (Ed.), *Belief revision meets philosophy of science* (pp. 1–42). Springer.
49. Castelfranchi, C. (1996). Reasons: Belief support and goal dynamics. *Mathware & Soft Computing*, 3, 233–247.
50. Castelfranchi, C. (1997). Representation and integration of multiple knowledge sources: Issue and questions. In V. Cantoni, V. Di Gesù, A. Setti, & D Tegolo (Eds.), *Human & machine perception: Information fusion* (pp. 235–254). Plenum Press.
51. Castelfranchi, C., D'Aloisi, D., & Giacomelli, F. (1995). A framework for dealing with belief-goal dynamics. In *AI*IA* (pp. 237–242).
52. Castelfranchi, C., & Paglieri, F. (2007). The role of beliefs in goal dynamics: Prolegomena to a constructive theory of intentions. *Synthese*, 155(2), 237–263.
53. Chopra, S., Ghose, A., & Meyer, T. (2003). Non-prioritized ranked belief change. *Journal of Philosophical Logic*, 32, 417–443(27).
54. Coniglio, M., & Carnielli, W. (2002). Transfers between logics and their applications. *Studia Logica*, 72, 367–400.

55. Cross, C., & Thomason, R. (1992). Conditionals and knowledge-base update. In P. Gärdenfors (Ed.), *Belief revision. Cambridge tracts in theoretical computer science* (Vol. 29, pp. 247–275). Cambridge: Cambridge University Press.
56. Dalal, M. (1988). Investigations into a theory of knowledge base revision: Preliminary report. In *Seventh national conference on artificial intelligence, (AAAI-88), St. Paul* (pp. 475–479).
57. Darwiche, A., & Pearl, J. (1996). On the logic of iterated belief revision. *Artificial Intelligence*, 89, 1–29.
58. de Kleer, J. (1986). An assumption-based TMS. *Artificial Intelligence*, 28(2), 127–162.
59. de Rijke, M. (1994). Meeting some neighbours. In J. van Eijck, & A. Visser (Eds.), *Logic and information flow* (pp. 170–195). Cambridge, MA.
60. Delgrande, J. (2008). 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). Sydney.
61. Delgrande, J., Dubois, D., & Lang, J. (2006). 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). Lake District, UK.
62. Delgrande, J., & Schaub, T. (2003). A consistency-based approach for belief change. *Artificial Intelligence*, 151(1–2), 1–41.
63. Delgrande, J., & Wassermann, R. (2010). Horn clause contraction functions: Belief set and belief base approaches. In *International conference on the principles of knowledge representation and reasoning, Toronto*.
64. Di Giusto, P., & Governatori, G. (1999). A new approach to base revision. In P. Barahona, & Alferes, J. J. (Eds.), *Progress in artificial intelligence* (pp. 327–341). Berlin: Springer.
65. Doyle, J. (1979). A truth maintenance system. *Artificial Intelligence*, 12, 231–272.
66. Doyle, J. (1992). Reason maintenance and belief revision: Foundations versus coherence theories. In P. Gärdenfors (Ed.), *Belief revision. Cambridge tracts in theoretical computer science* (Vol. 29, pp. 29–51). Cambridge: Cambridge University Press.
67. Doyle, J. (2011). Mechanics and mental change. In S. Artmann (Ed.), *Evolution of semantic systems*. Springer (to appear).
68. Dragoni, A. F. (1992). A model for belief revision in a multi-agent environment. In Y. Demazeau, & E. Werner (Eds.), *Decentralized A.I.* (pp. 103–112).
69. Dubois, D., & Prade, H. (1991). Epistemic entrenchment and possibilistic logic. *Artificial Intelligence*, 50(2), 223–239.
70. Dupin de Saint-Cyr, F., & Lang, J. (2002). Belief extrapolation (or how to reason about observations and unpredicted change). In D. M. D. Fensel, F. Giunchiglia, & M. Williams (Eds.), *Proceedings of the 8th international conference on principles of knowledge representation and reasoning (KR2002)* (pp. 497–508).
71. Eiter, T., & Gottlob, G. (1992). On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence*, 57, 227–270.
72. Fagin, R., Ullman, J., & Vardi, M. (1983). 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).
73. Falappa, M., Fermé, E., & Kern-Isberner, G. (2006). 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).
74. Falappa, M., García, A. J., Kern-Isberner, G., & Simari, G. R. (2010). Stratified belief bases revision with argumentative inference (manuscript).
75. Falappa, M., Kern-Isberner, G., & Simari, G. R. (2002). Belief revision, explanations and defeasible reasoning. *Artificial Intelligence*, 141, 1–28.
76. Falappa, M. A., Kern-Isberner, G., & Simari, G. R. (2009). Belief revision and argumentation theory. In G. Simari, & I. Rahwan (Eds.), *Argumentation in artificial intelligence* (pp. 341–360). US: Springer.
77. Fermé, E. (1992). Actualización de bases de conocimiento usando teorías de cambio de creencia. *Iberamia*, 92, 419–436.
78. Fermé, E. (1999). 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).

79. Fermé, E. (1999). *Revising the AGM postulates*. Ph.D. thesis, University of Buenos Aires.
80. Fermé, E. (2000). Irrevocable belief revision and epistemic entrenchment. *Logic Journal of the IGPL*, 8(5), 645–652.
81. Fermé, E., & Hansson, S. O. (1999). Selective revision. *Studia Logica*, 63(3), 331–342.
82. Fermé, E., & Hansson, S. O. (2001). Shielded contraction. In H. Rott, & M.-A. Williams (Eds.), *Frontiers in belief revision. Applied logic series* (pp. 85–107). Kluwer Academic Publishers.
83. Fermé, E., Krevneris, M., & Reis, M. (2008). An axiomatic characterization of ensconcement-based contraction. *Journal of Logic and Computation*, 18(5), 739–753.
84. Fermé, E., Mikalef, J., & Taboada, J. (2003). Credibility-limited functions for belief bases. *Journal of Logic and Computation*, 13(1), 99–110.
85. Fermé, E., & Reis, M. D. L. (2010). System of spheres-based multiple contractions. *Journal of Philosophical Logic* (to appear).
86. Fermé, E., & Rodriguez, R. (1998). A brief note about the Rott contraction. *Logic Journal of the IGPL*, 6(6), 835–842.
87. Fermé, E., & Rodríguez, R. (2006). DFT and belief revision. *Análisis Filosófico*, XXVII(2), 373–393.
88. Fermé, E., & Rott, H. (2004). Revision by comparison. *Artificial Intelligence*, 157, 5–47.
89. Fermé, E., Saez, K., & Sanz, P. (2003). Multiple kernel contraction. *Studia Logica*, 73, 183–195.
90. Friedman, N., & Halpern, J. Y. (1997). Modeling belief in dynamic systems. Part I: Foundations. *Artificial Intelligence*, 95(2), 257–316.
91. Friedman, N., & Halpern, J. Y. (1999). Belief revision: A critique. *Journal of Logic, Language, and Information*, 8, 401–420.
92. Friedman, N., & Halpern, J. Y. (1999). Modeling belief in dynamic systems. Part II: Revision and update. *Journal of AI Research*, 10, 117–167.
93. Fuhrmann, A. (1988). *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.
94. Fuhrmann, A. (1989). Reflective modalities and theory change. *Synthese*, 81(1), 115–134.
95. Fuhrmann, A. (1991). Theory contraction through base contraction. *Journal of Philosophical Logic*, 20, 175–203.
96. Fuhrmann, A. (1996). Everything in flux: Dynamic ontologies. In S. Lindström, R. Sliwinski (Eds.), *Odds and ends: Philosophical essays dedicated to Wlodek Rabinowicz on the occasion of his fiftieth birthday. Uppsala Philosophical studies* (Vol. 45, pp. 111–125).
97. Fuhrmann, A. (1997). *An Essay on contraction. Studies in logic, language and information*. Stanford: CSLI Publications.
98. Fuhrmann, A. (1997). Solid belief. *Theoria*, 63, 90–104.
99. Fuhrmann, A. (1997). Travels in ontological space. In G. Meggle, & A. Mundt (Eds.), *Analyomen 2. Proceedings of the 2nd conference perspectives in analytical philosophy* (pp. 68–77).
100. Fuhrmann, A., & Hansson, S. O. (1994). A survey of multiple contraction. *Journal of Logic, Language and Information*, 3, 39–74.
101. Gabbay, D., & Hunter, A. (1991). Making inconsistency respectable. In P. Jorrand, & J. Kelemen (Eds.), *Fundamental of artificial intelligence research (FAIR '91). Lecture notes in artificial intelligence* (Vol. 535, pp. 19–32).
102. Gabbay, D., Rodrigues, O., & Russo, A. (1999). Revision by translation. In R. R. Y. B. Bouchon-Meunier, & L. A. Zadeh (Eds.), *Information, uncertainty, fusion* (pp. 3–31).
103. Gabbay, D., Rodrigues, O., & Russo, A. (2008). Belief revision in non classical logics. *The Review of Symbolic Logic*, 1(03), 267–304.
104. Gabbay, D. M., Pigozzi, G., & Woods, J. (2003). Controlled revision—an algorithmic approach for belief revision. *Journal of Logic and Computation*, 13(1), 3–22.
105. Gallier, J. R. (1992). Autonomous belief revision and communication. In P. Gärdenfors (Ed.), *Belief revision. Cambridge tracts in theoretical computer science* (Vol. 29, pp. 220–246). Cambridge: Cambridge University Press.
106. Gärdenfors, P. (1982). Rules for rational changes of belief. In T. Pauli (Ed.), *Philosophical essays dedicated to Lennart Åqvist on his fiftieth birthday* (pp. 88–101).

107. Gärdenfors, P. (1984). Epistemic importance and minimal changes of belief. *Australasian Journal of Philosophy*, 62, 136–157.
108. Gärdenfors, P. (1986). Belief revisions and the Ramsey test for conditionals. *Philosophical Review*, 95, 81–93.
109. Gärdenfors, P. (1988). *Knowledge in flux: Modeling the dynamics of epistemic states*. Cambridge: MIT Press.
110. Gärdenfors, P. (1990). The dynamics of belief systems: Foundations versus coherence theories. *Reveu Internationale de Philosophie*, 44, 24–46.
111. Gärdenfors, P. (1991). Belief revision and nonmonotonic logic: Two sides of the same coin?. In J. Mvan Eijck (Ed.), *Logics in AI. European workshop JELIA '90 Amsterdam, The Netherlands, 10–14 September, 1990 Proceedings. Lecture notes in computer science* (Vol. 478, pp. 52–54).
112. Gärdenfors, P. (2011). Notes on the history of ideas behind AGM (this issue).
113. Gärdenfors, P., & Makinson, D. (1988). 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). Los Altos.
114. Gärdenfors, P., & Rott, H. (1993). Belief revision. In D. M. Gabbay, C. J. Hogger, & J. A. Robinson (Eds.), *Handbook of logic in artificial intelligence and logic programming. Epistemic and temporal reasoning* (Vol. 3, pp. 35–132). Oxford University Press.
115. Glaister, S. M. (2000). Recovery recovered. *Journal of Philosophical Logic*, 29, 171–206.
116. Governatori, G., & Rotolo, A. (2010). Changing legal systems: Legal abrogations and annulments in defeasible logic. *Logic Journal of IGPL*, 18(1), 157–194.
117. Grove, A. (1988). Two modellings for theory change. *Journal of Philosophical Logic*, 17, 157–170.
118. Grüne-Yanoff, T., & Hansson, S. O. (2009). 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.
119. Hansson, S. O. (1989). New operators for theory change. *Theoria*, 55, 114–132.
120. Hansson, S. O. (1991). *Belief base dynamics*. Ph.D. thesis, Uppsala University.
121. Hansson, S. O. (1991). Belief contraction without recovery. *Studia Logica*, 50, 251–260.
122. Hansson, S. O. (1992). A dyadic representation of belief. In P. Gärdenfors (Ed.), *Belief revision. Cambridge tracts in theoretical computer science* (Vol. 29, pp. 89–121). Cambridge: Cambridge University Press.
123. Hansson, S. O. (1992). In defense of base contraction. *Synthese*, 91, 239–245.
124. Hansson, S. O. (1992). In defense of the Ramsey test. *The Journal of Philosophy*, 89, 522–540.
125. Hansson, S. O. (1992). Similarity semantics and minimal changes of belief. *Erkenntnis*, 37, 401–429.
126. Hansson, S. O. (1993). Reversing the Levi identity. *Journal of Philosophical Logic*, 22, 637–669.
127. Hansson, S. O. (1993). Theory contraction and base contraction unified. *Journal of Symbolic Logic*, 58, 602–625.
128. Hansson, S. O. (1994). Kernel contraction. *Journal of Symbolic Logic*, 59, 845–859.
129. Hansson, S. O. (1994). Taking belief bases seriously. In D. Prawitz, & D. Westerståhl (Eds.), *Logic and philosophy of science in Uppsala* (pp. 13–28). Dordrecht: Kluwer Academic Publishers.
130. Hansson, S. O. (1995). Changes in preference. *Theory and Decision*, 38, 1–28.
131. Hansson, S. O. (1997). Semi-revision. *Journal of Applied Non-Classical Logic*, 7(1–2), 151–175.
132. Hansson, S. O. (1999). A survey of non-prioritized belief revision. *Erkenntnis*, 50, 413–427.
133. Hansson, S. O. (1999). *A textbook of belief dynamics. Theory change and database updating. Applied logic series*. Dordrecht: Kluwer Academic Publishers.
134. Hansson, S. O. (2000). Coherentist contraction. *Journal of Philosophical Logic*, 29, 315–330.
135. Hansson, S. O. (2003). Ten philosophical problems in belief revision. *Journal of Logic and Computation*, 13, 37–49.

136. Hansson, S. O. (2006). 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. <http://plato.stanford.edu/entries/logic-belief-revision/>.
137. Hansson, S. O. (2007). Contraction based on sentential selection. *Journal of Logic and Computation*, 17, 479–498.
138. Hansson, S. O. (2008). Specified meet contraction. *Erkenntnis*, 69, 31–54.
139. Hansson, S. O. (2009). Replacement—a sheffer stroke for belief revision. *Journal of Philosophical Logic*, 38, 127–149.
140. Hansson, S. O. (2010). Multiple and iterated contraction reduced to single-step single-sentence contraction. *Synthese*, 173, 153–177.
141. Hansson, S. O., Fermé, E., Cantwell, J., & Falappa, M. (2001). Credibility-limited revision. *Journal of Symbolic Logic*, 66(4), 1581–1596.
142. Hansson, S. O., & Makinson, D. (1997). Applying normative rules with restraint. In M. L. Dalla Chiara et al. (Eds.), *Logic and scientific method* (pp. 313–332). Kluwer Academic Publishers.
143. Hansson, S. O., & Olsson E. (1999). Providing foundations for coherentism. *Erkenntnis*, 51, 243–265.
144. Hansson, S. O., & Wassermann, R. (2002). Local change. *Studia Logica*, 70(1), 49–76.
145. Harman, G. (1986). *Change in view—principles of reasoning*. Cambridge York: MIT Press.
146. Harper, W. (1977). Rational conceptual change. In T. U. of Chicago Press (Ed.), *PSA: Proceedings of the biennial meeting of the philosophy of science association. Symposia and invited papers* (Vol. 2, pp. 462–494).
147. Hild, M., & Spohn, W. (2008). The measurement of ranks and the laws of iterated contraction. *Artificial Intelligence*, 172, 1195–1218.
148. Hintikka, J. (1962). *Knowledge and belief: An introduction to the logic of the two notions*. Ithaca, NY: Cornell University Press.
149. Jin, Y., & Thielscher, M. (2007). Iterated belief revision, revised. *Artificial Intelligence*, 171, 1–18.
150. Jin, Y., & Thielscher, M. (2008). Reinforcement belief revision. *Journal of Logic and Computation*, 18, 783–813.
151. Johnson, F. L. (2006). *Dependency-directed reconsideration: An anytime algorithm for hindsight knowledge-base optimization*. Ph.D. thesis, State University of New York at Buffalo.
152. Johnson, F. L., & Shapiro, S. C. (2005). 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).
153. Johnson, F. L., & Shapiro, S. C. (2005). Improving recovery for belief bases. In L. Morgenstern, & M. Pagnucco (Eds.), *IJCAI-05 workshop on nonmonotonic reasoning, action, and change (NRAC'05)* (pp. 65–70). Edinburgh.
154. Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press.
155. Katsuno, H., & Mendelzon, A. (1991). Propositional knowledge base revision and minimal change. *Journal of Artificial Intelligence*, 52, 263–294.
156. Katsuno, H., & Mendelzon, A. (1992). On the difference between updating a knowledge base and revising it. In P. Gärdenfors (Ed.), *Belief revision. Cambridge tracts in theoretical computer science* (Vol. 29, pp. 183–203). Cambridge: Cambridge University Press.
157. Kelly, K. (1998). Iterated belief revision, reliability, and inductive amnesia. *Erkenntnis*, 50, 11–58.
158. Kern-Isberner, G. (2004). A thorough axiomatization of a principle of conditional preservation in belief revision. *Annals of Mathematics and Artificial Intelligence*, 40(1–2), 127–164.
159. Kern-Isberner, G. (2008). Linking iterated belief change operations to nonmonotonic reasoning. In G. Brewka, & J. Lang (Eds.), *Proceedings 11th international conference on knowledge representation and reasoning, KR'2008* (pp. 166–176). Menlo Park, CA.
160. Konieczny, S., & Pérez, R. P. (2000). A framework for iterated revision. *Journal of Applied Non-Classical Logics*, 10(3–4), 339–367.
161. Konieczny, S., & Perez, R. P. (2008). Improvement operators. In *Eleventh international conference on principles of knowledge representation and reasoning (KR'08)* (pp. 177–186).

162. Konieczny, S., & Perez, R. P. (2011). Logic based merging (this issue).
163. Konieczny, S., & Pérez, R. P. (2002). Merging information under constraints: A logical framework. *Journal of Logic and Computation*, 12(5), 773–808.
164. Kourousias, G., & Makinson, D. (2007). Parallel interpolation, splitting, and relevance in belief change. *Journal of Symbolic Logic*, 72(3), 994–1002.
165. Kyburg, H. E. (1961). *Probability and the logic of rational belief*. Middletown: Wesleyan University Press.
166. Lang, J., & van der Torre, L. (2008). 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, 21–25 July, 2008, Proceedings. Frontiers in artificial intelligence and applications* (Vol. 178, pp. 351–355).
167. Langlois, M., Szörényi, R. H. S. B., & Turán, G. (2008). Horn complements: Towards horn-to-horn belief revision. In *Twenty-third AAAI conference on artificial intelligence (AAAI-08)* (pp. 466–471). AAAI Press.
168. Lehmann, D., Magidor, M., & Schlechta, K. (2001). Distance semantics for belief revision. *Journal of Symbolic Logic*, 66(1), 295–317.
169. Lehmann, D. J. (1995). Belief revision, revised. In *Proceedings of the fourteenth international joint conference on artificial intelligence (IJCAI'95)* (pp. 1534–1540).
170. Levi, I. (1977). Subjunctives, dispositions, and chances. *Synthese*, 34, 423–455.
171. Levi, I. (1980). *The enterprise of knowledge*. Cambridge: MIT Press.
172. Levi, I. (1988). Iteration of conditionals and the Ramsey test. *Synthese*, 76, 49–81.
173. Li, J. (1998). A note on partial meet package contraction. *Journal of Logic, Language and Information*, 7, 139–142.
174. Lin, J. (1996). Integration of weighted knowledge bases. *Artificial Intelligence*, 83(2), 363–378.
175. Lindström, S. (1991). *A semantic approach to nonmonotonic reasoning: Inference operations and choice*. Uppsala Prints and Preprints in Philosophy 6, Dep. of Philosophy, Uppsala University.
176. Lindström, S., & Rabinowicz, W. (1989). On probabilistic representation of non-probabilistic belief revision. *Journal of Philosophical Logic*, 19, 69–101.
177. Lindström, S., & Rabinowicz, W. (1991). Epistemic entrenchment with incomparabilities and relational belief revision. In A. Fuhrmann, & M. Morreau (Eds.), *The logic of theory change* (pp. 93–126). Berlin.
178. Lindström, S., & Rabinowicz, W. (1992). Belief revision, epistemic conditionals and the Ramsey test. *Synthese*, 91, 195–237.
179. Lindström, S., & Rabinowicz, W. (1998). Conditionals and the Ramsey test. In D. Gabbay, & P. Smets (Eds.), *Handbook of defeasible reasoning and uncertainty management systems (Belief Change)* (Vol. 3, pp. 147–188). Kluwer.
180. Lindström, S., & Rabinowicz, W. (1999). Belief change for introspective agents. In S. H. Bengt Hansson, & Nils-Eric Sahlin (Eds.), *Spinning ideas. Electronic essays dedicated to Peter Gärdenfors on his fiftieth birthday*. <http://www.lucs.lu.se/spinning/>.
181. Lindström, S., & Rabinowicz, W. (1999). DDL unlimited. Dynamic doxastic logic for introspective agents. *Erkenntnis*, 51, 353–385.
182. Liu, W. (2002). *A framework for multi-agent belief revision*. Ph.D. thesis, University of Newcastle.
183. Makinson, D. (1965). The paradox of the preface. *Analysis*, 25, 205–207.
184. Makinson, D. (1985). How to give it up: A survey of some recent work on formal aspects of the logic of theory change. *Synthese*, 62, 347–363.
185. Makinson, D. (1987). On the status of the postulate of recovery in the logic of theory change. *Journal of Philosophical Logic*, 16, 383–394.
186. Makinson, D. (1990). The Gärdenfors impossibility theorem in nonmonotonic contexts. *Studia Logica*, 49, 1–6.
187. Makinson, D. (1997). On the force of some apparent counterexamples to recovery. In E. Garzón Valdés et al. (Eds.), *Normative systems in legal and moral theory: Festschrift for Carlos Alchourrón and Eugenio Bulygin* (pp. 475–481). Berlin.
188. Makinson, D. (1997). Screened revision. *Theoria*, 63, 14–23.
189. Makinson, D. (2003). Ways of doing logic: What was new about AGM 1985. *Journal of Logic and Computation*, 13, 5–15.

190. Makinson, D. (2009). Propositional relevance through letter-sharing. *Journal of Applied Logic*, 7(4), 377–387.
191. Makinson, D. (2010). Conditional probability in the light of qualitative belief change (this issue).
192. Makinson, D., & Gärdenfors, P. (1991). Relation between the logic of theory change and nonmonotonic logic. In Fuhrmann, & Morreau (Eds.), *The logic of theory change* (pp. 185–205). Berlin.
193. Malheiro, B., Jennings, N. R., & Oliveira, E. (1994). Belief revision in multi-agent systems. In *ECAI* (pp. 294–298).
194. Malheiro, B., & Oliveira, E. (2000). Solving conflicting beliefs with a distributed belief revision approach. In Monard, M., & Sichman, J. (Eds.), *Advances in artificial intelligence. Lecture notes in computer science* (Vol. 1952, pp. 146–155). Berlin/Heidelberg: Springer.
195. Martins, J. & Shapiro, S. (1988). A model for belief revision. *Artificial Intelligence*, 35, 25–79.
196. Meyer, T., Lee, K., & Booth, R. (2005). Knowledge integration for description logics. In *Proceedings of the 7th international symposium on logical formalizations of commonsense reasoning* (pp. 645–650).
197. Moore, G. (1942). A reply to my critics. In P. Schilpp (Ed.), *The philosophy of G.E. Moore. The library of living philosophers* (Vol. 4, pp. 535–677). Evanston IL: Northwestern University.
198. Nayak, A. (1994). Iterated belief change based on epistemic entrenchment. *Erkenntnis*, 41, 353–390.
199. Nayak, A., Nelson, P., & Polansky, H. (1996). Belief change as change in epistemic entrenchment. *Synthese*, 109, 143–174.
200. Nayak, A., Pagnucco, M., Peppas, P. (2003). Dynamic belief revision operators. *Artificial Intelligence*, 146(2), 193–228.
201. Nebel, B. (1989). A knowledge level analysis of belief revision. In *Proceedings of the 1st international conference of principles of knowledge representation and reasoning* (pp. 301–311).
202. Nebel, B. (1992). Syntax-based approaches of belief revision. In P. Gärdenfors (Ed.) *Belief revision* (pp. 52–88).
203. Nebel, B. (1998). How hard is it to revise a belief base?. In D. Dubois, & Prade, H. (Eds.), *Handbook of defeasible reasoning and uncertainty management systems. Belief change* (Vol. 3), pp. 77–145. Dordrecht: Kluwer Academic Publishers.
204. Nittka, A. (2008). *A method for reasoning about other agents' beliefs from observations*. Ph.D. thesis, University of Leipzig.
205. Olsson, E. (1997). *Coherence*. Ph.D. thesis, Department of Philosophy. Uppsala University.
206. Olsson, E. (2003). Belief revision, rational choice and the unity of reason. *Studia Logica*, 73, 219–240.
207. Paglieri, F. (2005). See what you want, believe what you like: Relevance and likeability in belief dynamics. In *Proceedings AISB'05 symposium? Agents that want and like: Motivational and emotional roots of cognition and action* (pp. 90–97).
208. Paglieri, F. (2006). *Belief dynamics: From formal models to cognitive architectures, and back again*. Ph.D. thesis, Università degli Studi di Siena.
209. Paglieri, F., & Castelfranchi, C. (2004). 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).
210. Paglieri, F., & Castelfranchi, C. (2006). The Toulmin test: Framing argumentation within belief revision theories. In B. V. D. Hitchcock (Ed.), *Arguing on the Toulmin model* (pp. 359–377). Berlin: Springer.
211. Pagnucco, M. (1996). *The role of abductive reasoning within the process of belief revision*. Ph.D. thesis, Department of Computer Science, University of Sydney.
212. Parikh, R. (1999). Beliefs, belief revision, and splitting languages. In *Logic, language, and computation. CSLI lecture notes* (Vol. 96-2, pp. 266–268). Springer.
213. Peppas, P. (2004). The limit assumption and multiple revision. *Journal of Logic and Computation*, 14(3), 355–371.

214. Plaza, J. (1989). 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).
215. Priest, G. (2001). Paraconsistent belief revision. *Theoria*, 67(3), 214–228.
216. Rabinowicz, W. (1995). Global belief revision based on similarities between worlds. In S. O. Hansson, & W. Rabinowicz (Eds.), *Logic for a change* (Vol. 9, pp. 80–105). Uppsala Prints and Preprints in Philosophy. Dep. of Philosophy, Uppsala University.
217. Reis, M. D. L., & Fermé, E. (2010). Possible worlds semantics for partial meet multiple contraction. *Journal of Philosophical Logic* (to appear).
218. Reiter, R. (1987). A theory of diagnosis from first principles. *Artificial Intelligence*, 32, 57–95.
219. Revesz, P. Z. (1993). On the semantics of theory change: Arbitration between old and new information. In *Proceedings of the 12th ACM SIGACT-SIGMOD-SIGART symposium on principles of databases* (pp. 71–92).
220. Ribeiro, M. M., & Wassermann, R. (2009). Base revision for ontology debugging. *Journal of Logic and Computation*, 19(5), 721–743.
221. Rodrigues, O., & Benevides, M. (1994). Belief revision in pseudo-definite sets. In *Proceedings of the 11th Brazilian symposium on artificial intelligence (SBIA '94)*.
222. Rott, H. (1986). Ifs, though and because. *Erkenntnis*, 25, 345–37.
223. Rott, H. (1989). Conditionals and theory change: Revision, expansions, and additions. *Synthese*, 81, 91–113.
224. Rott, H. (1991). A nonmonotonic conditional logic for belief revision. Part 1: Semantics and logic of simple conditionals. In F. A., & M. Morreau (Eds.), *The logic of theory change, workshop, lecture notes in artificial intelligence* (Vol. 465). Konstanz, FRG.
225. Rott, H. (1992). On the logic of theory change: More maps between different kinds of contraction functions. In P. Gärdenfors (Ed.), *Belief revision. Cambridge tracts in theoretical computer science* (Vol. 29, pp. 122–141). Cambridge: Cambridge University Press.
226. Rott, H. (1992). Preferential belief change using generalized epistemic entrenchment. *Journal of Logic, Language and Information*, 1, 45–78.
227. Rott, H. (1993). Belief contraction in the context of the general theory of rational choice. *Journal of Symbolic Logic*, 58, 1426–1450.
228. Rott, H. (1995). Just because. Taking belief bases very seriously.. In S. O. Hansson, & W. Rabinowicz (Eds.), *Logic for a change* (Vol. 9, pp. 106–124). Uppsala Prints and Preprints in Philosophy. Dep. of Philosophy, Uppsala University.
229. Rott, H. (2000). 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). Prague.
230. Rott, H. (2000). Two dogmas of belief revision. *Journal of Philosophy*, 97(9), 503–522.
231. Rott, H. (2001). *Change, choice and inference: A study of belief revision and nonmonotonic reasoning. Oxford logic guides*. Oxford: Clarendon Press.
232. Rott, H. (2003). Basic entrenchment. *Studia Logica*, 73, 257–280.
233. Rott, H. (2003). Coherence and conservatism in the dynamics of belief. Part II: Iterated belief change without dispositional coherence. *Journal of Logic and Computation*, 13, 111–145.
234. Rott, H. (2009). Degrees all the way down: Beliefs, non-beliefs and disbeliefs. In Springer (Ed.), *Degrees of belief* (pp. 301–339). Dordrecht: Franz Huber and Christoph Schmidt-Petri.
235. Rott, H. (2009). Shifting priorities: Simple representations for twenty-seven iterated theory change operators. In H. W. D. Makinson, & J. Malinowski (Eds.), *Towards mathematical philosophy. Trends in logic* (Vol. 28, pp. 269–296). Springer Science.
236. Rott, H. (2010). *Bounded revision: Two dimensional belief change between conservative and moderate revision* (manuscript).
237. Rott, H., & Pagnucco, M. (1999). Severe withdrawal (and recovery). *Journal of Philosophical Logic*, 28, 501–547.
238. Ryan, M., & Schobbens, P.-Y. (1997). Counterfactuals and updates as inverse modalities. *Journal of Logic, Language and Information*, 6, 123–146.
239. Samet, D. (1996). Hypothetical knowledge and games with perfect information. *Games and Economic Behavior*, 17, 230–251.

240. Satoh, K. (1988). Nonmonotonic reasoning by minimal belief revision. In *FGCS* (pp. 455–462).
241. Schlechta, K. (1997). Non-prioritized belief revision based on distances between models. *Theoria*, 63, 34–53.
242. Segerberg, K. (1989). A note on the impossibility theorem of Gärdenfors. *Noûs*, 23, 351–354.
243. Segerberg, K. (1995). Belief revision from the point of view of doxastic logic. *Bulletin of the IGPL*, 3, 535–553.
244. Segerberg, K. (1995). Some questions about hypertheories. In S. O. Hansson, & W. Rabinowicz (Eds.), *Logic for a change. Uppsala prints and preprints in philosophy* (Vol. 9, pp. 136–153). Dep. of Philosophy, Uppsala University.
245. Segerberg, K. (1996). *Two traditions in the logic of belief: Bringing them together*. Uppsala Prints and Preprints in Philosophy 11, Dep. of Philosophy, Uppsala University.
246. Segerberg, K. (1997). *Irrevocable belief revision in dynamic doxastic logic*. Uppsala Prints and Preprints in Philosophy 6, Dep. of Philosophy, Uppsala University.
247. Sen, A. (1971). Choice functions and revealed preference. *Review of Economic Studies*, 38, 307–317.
248. Sen, A. (1997). Maximization and the act of choice. *Econometrica*, 65(4), 121–127.
249. Spohn, W. (1983). *Eine theorie der kausalität*. Unpublished Habilitationsschrift. Available at http://www.uni-konstanz.de/FuF/Philo/Philosophie/Spohn/spohn_files/Habilitation.pdf.
250. Spohn, W. (1988). 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) Dordrecht: D. Reidel.
251. Spohn, W. (2009). A survey of ranking theory. In C. S.-P. F. Huber (Ed.), *Degrees of belief. An anthology* (pp. 185–228). Dordrecht: Springer.
252. Spohn, W. (2010). Multiple contraction revisited. In M. R. Mauricio Suárez, & M. Dorato (Eds.), *EPSA epistemology and methodology of science. Launch of the European philosophy of science association* (Vol. 1, pp. 279–288). Dordrecht: Springer.
253. Stalnaker, R. (1968). A theory of conditionals. In N. Rescher (Ed.), *Studies in logical theory. American philosophical quarterly monograph series* (Vol. 2). Blackwell, Oxford. Also appears in *Ifs*, (ed., by W. Harper, R. C. Stalnaker and G. Pearce), Reidel, Dordrecht, 1981.
254. Tamminga, A. (2001). *Belief dynamics: (Epistemo)logical investigations*. Ph.D. thesis, Institute for Logic, Language and Computation, Universiteit van Amsterdam.
255. Val, A. D. (1997). Non monotonic reasoning and belief revision: Syntactic, semantic, foundational, and coherence approaches. *Journal of Applied Non-Classical Logics*, 7, 213–240.
256. van Benthem, J. (1989). Semantic parallels in natural language and computation. In *Logic colloquium '87*. Amsterdam.
257. van Benthem, J. (1994). Logic and the flow of information. In *Proceedings of the 9th international congress of logic, methodology and philosophy of science (1991)*. Also available as: Report LP-91-10, ILLC, University of Amsterdam.
258. van Benthem, J. (2004). Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 14, 1–26.
259. van Ditmarsch, H. (2005). Prolegomena to dynamic logic for belief revision. *Synthese (Knowledge, Rationality & Action)*, 147, 229–275.
260. van Ditmarsch, H., & Kooi, B. (2006). The secret of my success. *Synthese*, 151, 201–232.
261. van Ditmarsch, H., & Labuschagne, W. (2007). My beliefs about your beliefs—a case study in theory of mind and epistemic logic. *Knowledge, Rationality & Action (Synthese)*, 155, 191–209.
262. van Ditmarsch, H., van der Hoek, W., & Kooi, B. (2007). *Dynamic epistemic logic. Synthese library* (Vol. 337). Springer.
263. Wassermann, R. (1999). Resource bounded belief revision. *Erkenntnis*, 50, 429–446.
264. Wassermann, R. (2000). *Resource bounded belief revision*. Ph.D. thesis, University of Amsterdam.
265. Wassermann, R. (2001). Local diagnosis. *Journal of Applied Non-Classical Logics*, 11, 107–129.
266. Weydert, E. (1992). 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). Germany.

267. Weydert, E. (1994). 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).
268. Weydert, E. (2005). Projective default epistemology. In G. Kern-Isberner, W. Rödder, & F. Kulmann (Eds.), *Conditionals, information, and inference, international workshop, WCII 2002, Hagen, Germany, 13–15 May 2002, revised selected papers. Lecture notes in computer science* (Vol. 3301, pp. 65–85).
269. Williams, M.-A. (1992). Two operators for theory bases. In *Proc. Australian joint artificial intelligence conference* (pp. 259–265).
270. Williams, M.-A. (1997). Applications of belief revision. In *Proceedings of the fourteenth international joint conference on artificial intelligence* (pp. 74–79).
271. Williams, M.-A., & Sims, A. (2000). SATEN: An object-oriented web-based revision and extraction engine. In *International workshop on nonmonotonic reasoning (NMR'2000). Online computer science abstract*. <http://arxiv.org/abs/cs.AI/0003059/>.
272. Winslett, M. (1988). Reasoning about action using a possible models approach. In *AAAI* (pp. 89–93).
273. Zhang, D. (1996). Belief revision by sets of sentences. *Journal of Computer Science and Technology*, 11(2), 108–125.
274. Zhang, D. (2004). Properties of iterated multiple belief revision. In V. Lifschitz, & I. Niemelä (Eds.), *LPNMR 2004* (pp. 314–325).
275. Zhang, D. (2010). A logic-based axiomatic model of bargaining. *Artificial Intelligence*, 174, 1307–1322.
276. Zhuang, Z., & Pagnucco, M. (2010). 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 Berlin / Heidelberg