Routledge Companion to Epistemology
ed. by S. Bernecker and D. Pritchard
London & New York, 2011

Full version
Corrected May 2016

THEORIES OF BELIEF CHANGE

ANDRÉ FUHRMANN*

1. Synchronic versus diachronic epistemic justification

Karl sees Peter on a bicycle. He believes that if Peter rides a bicycle, then he owns that bicycle. Thus, Karl believes (at time t) that

(p) Peter owns a bicycle.

Here is a traditional epistemological problem: Let K be Karl's state of belief at t. Are all of Karl's beliefs justified? In particular, are the two beliefs

(a) Peter rides a bicycle

and

(b) If Peter rides a bicycle, than he owns that bicycle

justified? For, if they are, then they jointly entail p.

This is the problem of epistemic justification concerning actual beliefs. We take a time slice of an agent's beliefs and ask whether the beliefs at the given point of time are justified. This problem of *synchronic justification* enjoys such a central status in modern epistemology that it is frequently taken to *define* the subject. In a way, this is rather unfortunate. For, there is another problem, at least as pressing as the problem of synchronic justification, which besets our states of belief.

Suppose that Karl learns from a source he recognizes as reliable that Peter does not own a bicycle. He therefore resolves to retract his belief that p. In other words, he resolves to move from K to a belief state *K-without-p*. But in K the belief that p is deductively forced by the beliefs a and b ($= a \to p$). So in order to give up p, Karl must give up at least a or b. And backtracking from a and b in the deductive network of K we may find further beliefs involved in producing p. There are thus many candidate states that all satisfy the requirement that p be no longer believed—including the initial Cartesian state, *cogito*. The challenge is to justify a choice of such a candidate state as the unique successor *K-without-p* to K. This is the

* Institut für Philosophie, Goethe-Universität Frankfurt am Main.

problem of justifying belief change or of *diachronic justification*, as we shall say.

The above example shows clearly that no unique successor belief state K' is in general determined by a given origin state K, a piece of evidence and a logical consequence relation that calculates the consequences of K and the new evidence or backtracks unwelcome consequences to subsets of K. Transitions to unique successor states essentially involve an element of *choice*.

The problem of synchronic justification consists in finding items that may plausibly replace the variable in the schema

(S) $$x \text{ R } K$$

where R indicates the relation of justification. Prominent candidates for x in (S) are fundamental beliefs, coherent integration of beliefs or reliable causal groundings.

In the diachronic case the schema under consideration is this:

(D) $$x \text{ R } (K, K')$$

where the object of justification is a pair representing a transition from an origin state K to a successor state K'.

One may ask whether the two problems are distinct. This is an intricate question. If we could somehow derive a solution to the problem of diachronic justification from an extant account of synchronic justification the apparent dominance of the synchronic problem in modern epistemology would emerge as well-deserved in retrospect. To appreciate the difficulties facing attempted reductions let us consider two simple proposals.

First, the transition (K, K') may be said to be justified just in case K' is justified at the time t' of transition while K was justified at a time t when K represented the state of belief of the agent and no change inducing evidence was available to him. This proposal falls to the obvious objection that for (K, K') to be a good transition we do not need to know whether K or K' are good belief states. In this respect the transition is a kind of inference which only aims at a guarantee that the successor state is no worse than the origin state.

The second proposal tries to take away the defect of the first. Thus one may say that the transition (K, K') is justified just in case K' enjoys a higher degrees of justification than K. This proposal may fail in both directions—assuming that a belief's degree of justification is inversely proportional to the risk of error it incurs. This assumption for the sake of the argument is not to rule out the possibility that certain notions of justification do not satisfy the condition of proportionality. Coherentist justification, for example, does not increase with avoiding risk and would thus be immune to the following counterexamples to the proposal at hand.

For one, we may have good transitions to worse belief states. Let A be any belief in K and let K/A be that subset of K that contains all beliefs as least as good (justified) as A. Then for all A, K/A is more justified than K. But most transitions to thus generated subsets of K would involve gratuitous losses of information insufficiently balanced for by protection from the risk of error.

As to the other direction, bad transitions may lead to better justified belief states. The case here is dual to the one just considered. Any expansion of a belief state results in a corpus of beliefs that is riskier and therefore justified to a lesser degree than the state before expansion. We should not now detach the conclusion that expansions are always irrational.

By contrast, the reduction of synchronic to diachronic justification seems to be a more promising enterprise. As a first attempt consider the Cartesian proposal: A belief state K (and thereby each belief in K) counts as justified if and only if it is connected by a chain of justified belief changes to an unassailable *urcorpus* of beliefs. The term 'urcorpus' is taken from Levi (see e.g. [22]). It does not matter here what one takes to be part of the *urcorpus*. (Levi, by the way, does not advocate assigning the *urcorpus* a foundational rôle.) At best, however, the Cartesian proposal affords a partial reduction of justified belief to justified belief change, since the *urcorpus* cannot itself be the result of justified belief changes. Approaches that take their inspiration in a pragmatist theory of inquiry do away with ultimate epistemic foundations. In a way, such approaches may be taken to call into question the very problem of synchronic justification: There is a real problem of legitimately fixing ones beliefs; and there is the bogus problem of legitimizing beliefs already fixed. (See Peirce [29] and Levi's "Truth, fallibility and the growth of knowledge" in [23].)

In sum, there are two *prima facie* problems of epistemic justification: the problem of justifying one's presently held beliefs and the problem of justifying a change in one's present beliefs. It is by no means obvious that the one problem can be reduced to the other. Moreover, in a first person perspective the problem of justifying belief change seems much more pressing than the problem of justifying presently held beliefs. For, on the one hand, a proposition believed is a proposition held to be true. On the other hand, attempts at justification are responses to living doubts. But it is incoherent to doubt something that at the same time is held to be true. Living doubt as to present beliefs arise when hitherto unknown evidence becomes available. Then we face the question how to accommodate the new evidence. Such accommodation needs justification. But the required justification is one that pertains to a belief change, not to the belief state prior to the new evidence.

In the light of these consideration it is all the more surprising that theories of diachronic justification have been largely neglected for a long time, with a few exceptions. Only in the 1980s—again, with a few exceptions—philosophers started to develop systematic approaches to the problem. Since then theories of belief change (or theory change or belief revision) have rapidly progressed to a presently rich field of investigation with many cross-connections to other problems, including some of particular concern to computer scientists.

2. The classical theory of belief change: AGM

The fundamental problem. Let us start by asking how much structure in a belief state we need to invest for an interesting problem of belief change to emerge.

We ascribe beliefs to agents. We usually do this by using language, though what we ascribe, using language, are not linguistic items. When using language to ascribe beliefs, we proceed on the innocuous presupposition that sentences may be used to represent beliefs. When we refer to the beliefs of an agent, we may therefore do so by representing these beliefs by a set of sentences. Of course, the language needs to be expressive enough to represent the beliefs we are interested in representing. Should that prove difficult, we may start by restricting our attention to an easily representable kind of beliefs.

Beliefs have consequences. A believer is committed to the consequences of his beliefs. There are believers who fail to live up to their commitments. But we do expect believers to acknowledge their commitments once they are clearly and patiently pointed out to them. This expectation is not of an inductive but of a normative kind: One *ought* to acknowledge the consequences of one's beliefs. Thus belief states—as they ought to be—are closed under consequence.

We are now ready to formulate the fundamental problem of belief change. Suppose an agent's belief state K contains beliefs A and B from which follows (by consequence) a further belief C. By closure under consequence, C is also in K. Suppose now that a decision has been taken to remove C from K. How should this decision be implemented? This is the given change problem. At least one of A or B has to give way. What should the resulting belief state not containing C be like?

At this point we could stop with the simple observation that K without A, or K without B, or K without both, or even the closure under consequence of $\{$cogito$\}$ are all successors to K satisfying the condition that C be no longer believed. Belief change would thus be relational rather than functional: There simply is no unique successor to K.

Since relational belief change is unhelpful (though see Lindström and Rabinowicz [27] for a dissenting view) we need to assume some more structure on belief states which delivers a ranking of possible successors such that each change problem issues in a unique successor.

To summarize, we need two assumptions for posing the problem. First, beliefs must be representable in some language for which we may define a relation of consequence. Equivalently we may represent beliefs by sets of possible worlds, with inclusion doing duty for (the converse of) consequence; see e.g. Grove [18] or Stalnaker [32]. Second, the beliefs held in a belief state are closed under consequence. Some philosophers have argued that for certain purposes it might be advisable to keep track of the base beliefs from which complete belief states are generated by closing under consequence: see Fuhrmann [14] and Hansson [19]. Even if one prefers base changes over theory changes, the problem of belief change emerges in the same way, i.e. only once we consider logical commitments. Finally, a third assumption is needed for solving the problem: Belief states are not merely closed sets of beliefs; they have further structure ranking some (clusters of) beliefs higher than others.

The sources of the AGM-theory. The classical theory of belief change has been developed in the 1980s by Carlos Alchourrón, Peter Gärdenfors and David Makinson (hence AGM for short). The "pure" theory of belief change resulted as an abstraction from two more concrete and well-known problems.

Alchourrón's principal interest lay in the concept of legal derogation.

> Suppose that A is a set of regulations, y is some proposition that is implied by A, and that for some reason a legislative body wants to eliminate y. In such a situation, the body may decide to reject y, with the intention of thereby rejecting implicitly whatever in A implies y, retaining the remainder. This we shall call derogation. ([2], p. 127.)

Alchourrón and Makinson presented their solution to the formal problem of derogation in [2].

Gärdenfors tried to elaborate the idea of belief change implicit in the following much quoted passage from Frank P. Ramsey:

> If two people are arguing 'If p will q?' and are both in doubt as to p, they are adding p hypothetically to their stock of knowledge and arguing on that basis about q. ('General propositions and causality' (1929), in [30].)

Thus the theory of belief change to be developed was thought to be constrained by the acceptance test for conditionals suggested by Ramsey. We shall come back to this constraint below. Gärdenfors [17] is the first book-length treatment of theory change. The three authors joined forces in the standard exposition [1] of the AGM-theory. The development of the AGM-theory has been much inspired and critically accompanied by the work of Isaac Levi; see in particular [23, 24, 25]. (Rott [31] summarizes the rise of the AGM-theory over 20 years, provides important missing links and draws out connections with defeasible and practical reasoning. Hansson [20] is a textbook-style exposition of the AGM-theory containing also some important extensions.)

In the AGM-theory it is assumed that the language used to represent belief states contains all boolean connectives and that the consequence operation (or relation) under which belief states are to be closed is a closure operation in the usual Tarskian sense, including classical logical consequence. As pointed out in Fuhrmann [15] and [16], for the most part of the theory no such assumptions are really necessary. In fact, the theory enjoys such a degree of abstraction that it covers not only belief change but changes of any items that may be represented in a closure systems.

The consequence operation Cn is thus expected to satisfy the following conditions (for any sets X and Y of sentences):

$$X \subseteq \mathrm{Cn}(X);$$

$$\text{if } X \subseteq Y, \text{ then } \mathrm{Cn}(X) \subseteq \mathrm{Cn}(Y);$$
$$\mathrm{Cn}(\mathrm{Cn}(X)) = \mathrm{Cn}(X);$$
$$\text{if } A \in \mathrm{Cn}(X), \text{ then } \exists X' \subseteq X : X' \text{ is finite and } A \in \mathrm{Cn}(X').$$

Sometimes we may also write relationally:

$$X \vdash A \text{ for } A \in \mathrm{Cn}(X) \quad \text{and} \quad A \equiv B \text{ for } \mathrm{Cn}(A) = \mathrm{Cn}(B).$$

Since belief states K are closed under logical consequence we have (for every sentence A)

$$A \in K \text{ iff } K \vdash A.$$

Suppose now that we receive input information that triggers a decision to incorporate A into one's present state of belief K. There are two cases to distinguish.

First, A is consistent with K. In that case incorporation is simple: We add A to K and then close again under consequence. We write $K + A$ for this operation of *expanding* K by A. Expansion is easily defined: $K + A := \mathrm{Cn}(K \cup \{A\})$, where Cn is the operation of taking the consequences of a set of beliefs.

Second, A is inconsistent with K, i.e. $\neg A$ is in K. Here simple expansion would result in an inconsistent belief state–of which classically there is only one, the trivial belief state including all beliefs. Thus a more subtle operation is called for which will be called the *revision* of K by A, $K * A$.

It is natural to think of revision as a composed operation: First we have to withdraw beliefs from K so as to make our beliefs compatible with A; then we may add A by expansion. The result of retracting from K to a weaker belief state is called a *contraction*. In the case at hand, we would need to contract K by $\neg A$; then we could expand the result, $K - \neg A$, by A without danger of inconsistency. This reduction of revisions to contractions followed by an expansion is the *Levi Identity*:

$$K * A = (K - \neg A) + A.$$

Given the Levi Identity the theory of belief changes reduces to a theory of contractions (and expansions). We note in passing that one may also take the other direction, from contractions to revisions, via the so-called Harper (or Gärdenfors) Identity: $K - A = (K * \neg A) \cap K$.

The strategy of the theory is simple. On the one hand we transpose plausible conditions on contractions (or revisions) into the language of set-theory, thereby giving an implicit characterization of the operation. On the other hand, we try to give a recipe for solving a given contraction problem, thereby producing models in which contractions can explicitly be defined. Finally, the two approaches need to be linked by a representation result: An operation satisfies the contraction postulates just in case it can be defined as a contraction in a model. If postulates on the one hand and models on the other hand are independently plausible, such a representation result makes for a *prima facie* stable theory of contractions and of belief change in general.

Of the four models furnished in the AGM-theory we shall here present two: partial meet contractions and systems of spheres. The other two are safe contractions, as in [2], and models in terms of epistemic entrenchment, for which see [17] and [31].

The AGM-Postulates. The postulates for contraction come in two groups:

Basic

(C1)	Closure	$K - A = \mathrm{Cn}(K - A)$
(C2)	Success	$A \notin K - A$, if $\nvdash A$
(C3)	Inclusion	$K - A \subseteq K$
(C4)	Vacuity	If $A \notin K$, then $K - A = K$
(C5)	Recovery	$K \subseteq (K - A) + A$
(C6)	Congruence	If $A \equiv B$, then $K - A = K - B$

Supplementary

(C7)	Conjunction 1	$K - A \cap K - B \subseteq K - (A \wedge B)$
(C8)	Conjunction 2	If $A \notin K - (A \wedge B)$, then $K - (A \wedge B) \subseteq K - A$

Likewise the postulates for the revision operation:

Basic

(R1)	Closure	$K * A = \mathrm{Cn}(K * A)$
(R2)	Success	$A \in K * A$
(R3)	Inclusion	$K * A \subseteq K + A$
(R4)	Preservation	If $\neg A \notin K$, then $K * A = K + A$
(R5)	Consistency	If $\bot \in K * A$, then $\vdash \neg A$
(R6)	Congruence	If $A \equiv B$, then $K * A = K * B$

Supplementary

(R7)	Conjunction 1	$K * (A \wedge B) \subseteq (K * A) + B$
(R8)	Conjunction 2	If $\neg B \notin K * A$, then $K * A) + B \subseteq K * (A \wedge B)$

The labels must suffice here to indicate the rationale for each of the postulates; only Recovery (C5) will be considered more closely in a moment.

The postulates harmonize: We can derive the revision postulates from the contraction postulates via the Levi-Identity $K * A = (K - \neg A) + A$. (Note that for this derivation Recovery is not needed.) Conversely we can derive the contraction postulates (including Recovery) from the revision postulates via the Harper Identity $K - A = (K * \neg A) \cap K$.

Recovery requires that in contracting we should avoid gratuitous loss of information. Without such a condition contractions are insufficiently constrained by the other postulates: they may result in too small belief states. For example, the operation defined by $K - A = \mathrm{Cn}(\emptyset)$ would satisfy all the remaining C-postulates (Makinson [28]).

The idea that contractions should minimize losses of information is essential to a theory of rational belief changes. Thus the theory requires a condition like Recovery. Yet Recovery stands out in at least three other respects.

First, as pointed out above, Recovery is redundant in deriving revisions from contractions.

Second, unlike the other postulates, Recovery is sensitive to the underlying logic: it depends essentially on closure under *classical* consequence. For every $B \in K$ we have by closure for every A, $\neg A \vee B \in K$. When we then contract by A, there is no logical compulsion to remove $\neg A \vee B$. Thus in the models considered below we invariably have $\neg A \vee B$ in every plausible contraction candidate whence also in $K - A$. But then expanding $K - A$ by A reinstates any B that was previously removed from K. (The last move proceeds by disjunctive syllogism, an inference called into question by relevant logicians among others; see e.g. [11].)

Third, under certain circumstances Recovery seems implausible. In order to remove A from K, we typically need to remove stronger beliefs that imply A from K. But when we then reinsert the weaker A, why and how should we then regain those stronger beliefs? The circumstances that call Recovery into question are easily identified: Recovery fails indeed when the objects of contractions are open belief bases instead of closed belief states. For certain purposes it may be preferable to change states of beliefs by making all incisions in an axiomatic, usually finite representation of such a state; see [14, 15, 19]. This is also the more general format for a theory of belief change. The AGM-theory, however, operates under the assumptions that the units of change are *closed* sets of belief, i.e. theories in the logical sense.

The Ramsey-Test for conditionals. One of the initial motivations for developing a theory of belief revision was the problem of characterizing the operation of "adding something hypothetically to a stock of belief" such that an acceptance or assertion condition for certain conditionals would result, as suggested by Ramsey. Here then is the *Ramsey Test* for counterfactual conditionals:

RT. $$A \sqsupset B \in K \text{ iff } B \in K * A.$$

One might take *RT* either (from right to left) as a recipe for generating a logic of conditionals, given a theory of revisions, or (from left to right) as generating a theory of revisions from a given logic of conditionals, or simply as a bridging principle requiring conditionals and revisions to walk hand in hand.

Any such aspirations have to be laid at rest by observing that *RT* requires revisions to be monotonic:

MR. $$\text{If } K \subseteq H, \text{ then } K * A \subseteq H * A.$$

Monotonicity is not a desirable feature for revisions; counterexamples are not difficult to find. Yet we do not need to remain at the level of assessing intuitions. For, from (MR) we quickly slip to triviality. It can be shown that (MR), Inclusion, Preservation and Consistency are compatible only on pain of triviality. The result was found by Gärdenfors [17] and will be proved in a moment. It is closely related to Lewis's impossibility theorem, in [26], concerning the identification of probabilities of conditionals with conditional probability and Arrow's theorem concerning social choice under certain conditions; see [21].[1]

We suppose that there are A, B, C and K such that A, B and C are consistent with K and pairwise disjoint. This is to say,

(A) $\neg A, \neg B$ and $\neg C$ are not in K;
(B) $A \vdash \neg B$, $B \vdash \neg C$, and $C \vdash \neg A$.

Let $H = K * A * B \vee C$ and let $H' = K + A \vee B * B \vee C$. We now reduce the set $\{\mathrm{R2}, \mathrm{R4}, \mathrm{R5}, \mathrm{MR}, \mathrm{A}, \mathrm{B}\}$ to absurdity.

First step: It follows from (A) that $\nvdash \neg(B \vee C)$. Using (R5) we infer from this that H is consistent. Since, by (R2), $B \vee C \in H$, we have $\neg(B \vee C) \notin H$. Hence, either (1a) $\neg B \notin H$ or (1b) $\neg C \notin H$.

Since $A \vdash A \vee B$, we have (2) $K + A \vee B \subseteq K + A$. By (A) we assume that $\neg A \notin K$. It thus follows by (R4) that $K + A = K * A$, whence by (2), (3) $K + A \vee B \subseteq K * A$. Applying (MR) to (3) we obtain $K + A \vee B * B \vee C \subseteq K * A * B \vee C$, i.e. (4) $H' \subseteq H$. Together with (1) this gives: either (5a) $\neg B \notin H'$ or (5b) $\neg C \notin H'$.

Second step: Let us now consider $K + A \vee B$. Assume—for *reductio*!—that (6) $\neg(B \vee C) \in K + A \vee B$. Then $A \vee B \to \neg(B \vee C)$, resp. its equivalent $\neg B \wedge \neg(A \wedge C)$ is in K. According to (B), the second conjunct is a tautology. Hence, $\neg B \in K$, contrary to our assumption (A). It follows that (6) is false, i.e. (7) $\neg(B \vee C) \notin K + A \vee B$.

By (R4) it follows from (7) that (8) $K + A \vee B + B \vee C \subseteq K + A \vee B * B \vee C\,(= H')$. By the definition of expansion we have $K + A \vee B + B \vee C = K + (A \vee B) \wedge (B \vee C)$. By assumption (B) we have $(A \vee B) \wedge (B \vee C) \equiv B$. Putting this together, we get (9) $K + A \vee B + B \vee C = K + B$. It follows from (8) and (9) that (10) $K + B \subseteq H'$. Since by (R2) $B \in K + B$, assumption (B) gives us $\neg C \in K + B$, whence, by (10), $\neg C \in H'$. It now follows from (5) that (5a) $\neg B \notin H'$.

The *third step* starts from considering $K + A \vee C$ and then basically repeats the second. We verify that (7') $\neg(B \vee C) \notin K + A \vee C$ and conclude via $(A \vee C) \wedge (B \vee C) \equiv C$, that (10') $K + C \subseteq H'$ and, using (B) again, that $\neg B \in H'$, contary to (5a). This concludes the *reductio* for $\{\mathrm{R2}, \mathrm{R4}, \mathrm{R5}, \mathrm{MR}, \mathrm{A}, \mathrm{B}\}$.[2]

Various ways of constraining RT in such a way that its spirit is maintained have been explored. The general lesson seems to be this: Caution needs to be exercised when trying to systematically connect factual beliefs ("objectual beliefs") with beliefs reflecting the behaviour or content of belief

[1] The following proof, due to Gärdenfors, replaces an incorrect one, printed in the original contribution to the *Handbook*.
[2] End of correction.

states ("meta-beliefs"). This is confirmed by a similar impossibility result concerning the *Levi Test* for serious possibility: (Levi does not propose such a test. He notes that such a test for epistemic possibility may be suggested and counsels caution.)

LT. $$\Diamond A \in K \text{ iff } \neg A \notin K.$$

The modality in *LT* is what is called a reflective modality in [12]; all such tests of a reflective kind, including *LT* can be shown to be inconsistent.

Partial meet contraction. As pointed out above, belief change is fundamentally a choice problem. The partial meet models of AGM model this element of choice by brute force by simply assuming the existence of an appropriate choice function.

Suppose we decide to remove A from K. We might proceed as follows. First we collect all maximal subsets of K that do not entail A. (Let $K \perp A$ denote the family of such *remainders*.) Then we select the "best" of them. (Let s_K be a *choice function* that is defined on the family of subsets of K; if $K \perp A$ happens to be empty, then we let $s(K \perp A) = \{K\}$.) In the sequel we omit the subscript K to the choice function s. But it is important to remember that the choice functions used represent the theoretical preferences of a given belief set K. These preferences may change after a contraction has been performed. It is for this reason that the AGM-theory does not cover in any interesting way iterated belief changes. For an attempt at tackling iterated changes see [8]. Finally we take what is common to all such best subsets. Thus

$$K - A := \bigcap s(K \perp A).$$

It is not difficult to verify that this definition of a contraction operation satisfies all the basic C-postulates.

Conversely, every contraction operation satisfying the basic C-postulates can be defined as a partial meet contraction. We prove this in a moment. Together the two observations yield the sought representation result: Essentially the same operation may be characterized implicitly by the postulates (C1-6) or explicitly defined as a partial meet contraction. Thus postulates and models support each other.

To prove that the postulates pick out partial meet contractions we assume that $K - (\ )$ satisfies (C1-6) and define a canonical selection function for K, σ, as follows:

$$(\sigma) \qquad \sigma(K \perp A) := \begin{cases} \{X \in K \perp A : K - A \subseteq X\}, & \text{if } K \perp A \neq \emptyset; \\ \{K\}, & \text{otherwise.} \end{cases}$$

We need to show that σ (for K) is a selection function s_K of the above introduced kind. This is done by verifying five claims:

(1) σ is well-defined, i.e. $K \perp A = K \perp B$ entails $\sigma(K \perp A) = \sigma(K \perp B)$.
(2) $\sigma(K \perp A) = \{K\}$, if $K \perp A = \emptyset$.
(3) $\sigma(K \perp A) \subseteq K \perp A$, if $K \perp A \neq \emptyset$.
(4) $\sigma(K \perp A) \neq \emptyset$.
(5) $K - A = \bigcap \sigma(K \perp A)$.

(1-4) are more or less immediate reflections of the C-postulates. It will thus suffice to give the argument for (5). One direction, $K - A \subseteq \bigcap \sigma(K \perp A)$, is immediate from (σ) and Inclusion (C3). It remains to show

$$(*) \qquad \text{If } B \notin K - A, \text{ then } B \notin \bigcap \sigma(K \perp A)$$

Case $A \notin K$. Then by Vacuity (C4) $K - A = K$ and so, according to (σ), $K = \bigcap \sigma(K \perp A)$, whence (*) holds.

Case $A \in K$. (*) is trivially true if $B \notin K$. So let us suppose further that $B \in K$. We need to convince ourselves that there will be a set X such that

$$\text{(a) } K - A \subseteq X \quad \text{and} \quad \text{(b) } B \notin X \text{and} \quad \text{(c) } X \in K \perp A.$$

We proceed on the basis of a Lemma proved in [1] (Lemma 2.4):

$$\frac{X \in K \perp C \quad D \in K \quad X \nvdash D}{X \in K \perp D}$$

Then the sought conclusion follows thus:

$$Ad \text{ a)} \quad \dfrac{\dfrac{\dfrac{\dfrac{\dfrac{}{B \in K}\,\text{Hyp.}}{K - A, A \vdash B}\,\text{(C5)} \qquad \dfrac{}{K - A \nvdash B}\,\text{Hyp.}}{K - A, \neg A \nvdash B}}{K - A \nvdash A \vee B}}{\exists X \in K \perp A \vee B : K - A \subseteq X} \qquad Ad \text{ b)} \quad \dfrac{\dfrac{\dfrac{}{X \in K \perp A \vee B}\,\text{a)}}{X \nvdash A \vee B}}{B \notin X}$$

$$Ad \text{ c)} \quad \dfrac{\dfrac{}{A \in K}\,\text{Hyp.} \qquad \dfrac{}{X \in K \perp A \vee B}\,\text{a)} \qquad \dfrac{\dfrac{\dfrac{}{X \in K \perp A \vee B}\,\text{a)}}{X \nvdash A \vee B}}{X \nvdash A}}{X \in K \perp A}\,\text{Lemma}$$

This concludes the proof that contractions satisfying the basic AGM-postulates represent exactly the partial meet contractions. The result may be extended to the supplementary postulates by generating the selection function from a preference relation; see [1], §4. The representation result carries over to revisions via the Levi Identity.

Systems of spheres. The Ramsey Test, even though not tenable in its naive form, suggests a close analogy between counterfactual conditionals and the revision operation. The analogy naturally invites the thought that extant semantics for counterfactuals may be adapted to yield models of belief revision.

The Ramsey Test may be seen as an epistemic version of what one may call the *Stalnaker Test*:

$$A \sqsupset B \text{ in } w \text{ iff } B \text{ in } (w * A).$$

The $*$-operation on the right-hand-side takes the world w and sentence A to the least deviation from ("revision" of) w so as to make A true. Perhaps there is no unique such least deviation, whence we may prefer the *Lewis Test*:

$$A \sqsupset B \text{ in } w \text{ iff } B \text{ in all worlds in } (w * A).$$

The task at hand is now to determine $w * A$. A system of nested spheres centered around w is the key to modelling $w * A$ in David Lewis' semantics for counterfactuals. Put briefly, in the spheres semantics for counterfactuals worlds are ordered around w in spheres of similarity (see the left illustration below). If sphere S_2 is more distant from w than S_1, i.e. if $S_1 \subseteq S_2$, then the worlds in S_1 are more similar to w than the worlds in S_2. To assess $A \sqsupset B$ at w, we look at the closest-to-w sphere S that contains A-worlds; $w * A$ collects the A-worlds in S. To evaluate the conditional we check whether all worlds in $w * A$ satisfy B. If so, the conditional is true at w, otherwise it is false at w. This is the principal case. If there are no A-permitting spheres around w, i.e. if the antecedent is not entertainable when viewed from w, then the conditional is stipulated to be vacuously true at w.

The Ramsey Test, however, does not test for the truth of a conditional in a world w but for acceptance in a belief state K. A belief states can be modelled by the set of worlds not ruled out by what is believed. Thus, in the Ramsey Test,

$$A \sqsupset B \text{ in } K \text{ iff } B \text{ in } K * A,$$

the belief states K and $K * A$ may be interpreted as standing for sets of worlds. We may simply adapt the spheres models for counterfactuals by centering spheres on a *set* of worlds (representing K) and determining the new belief state $K * A$ as before. The sphere represent now fallback positions ordered as to comparative plausibility from the viewpoint of K. The result is shown on the right-hand-side of the following illustration.

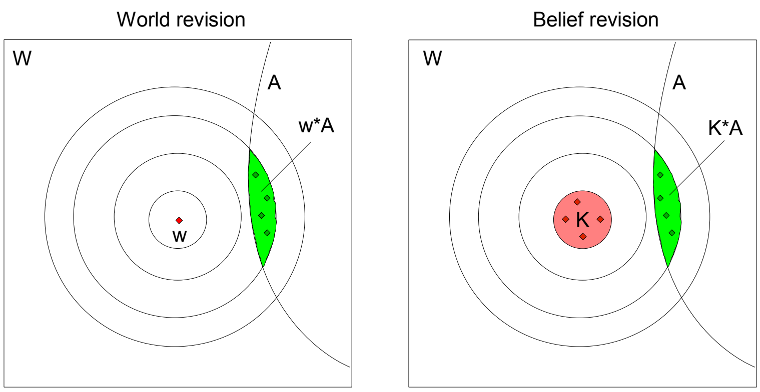

This is the proposal. It remains to verify that it works. That is to say, we need another representation result.

Let W be a universe of possible worlds, let K be a belief state and let $[K] \subseteq W$ be the representation of K in W. A system of spheres for K is a family of subsets of W such that $[K]$ is the smallest and W the largest sphere, all spheres are comparable by set-inclusion, and for any nonempty subset X of W, there is always a smallest sphere, S_X, cutting X nonemptily. Thus a spheres system for K may be pictured as an onion or Russian doll, with $[K]$ in the centre and W as the outermost shell; formal definitions may be found in [18]. We define the set of worlds representing $K * A$ thus:

$$(*) \qquad [K * A] := \begin{cases} S_{[A]} \cap [A], & \text{if } A \text{ is consistent;} \\ \mathrm{Cn}(\bot) & \text{otherwise.} \end{cases}$$

Then we can prove the following result (Grove [18]): Revisions, as defined in (*), satisfy the revision postulates (R1-8), and any revision operation satisfying (R1-8) can be defined as in (*). An analogous representation result holds for contractions *via* the Harper Identity.

Given the correspondence between revision in systems of spheres and revisions satisfying the AGM-postulates, there follows a correspondence between spheres-revisions and partial meet revisions. This should not be surprising. The latter correspondence can be made transparent without detouring via the postulates. It suffice to note a bijection between $[\neg A]$ and $K \perp A$. (Here we only look at the principal case that $A \in K$ and A is consistent.) From spheres to remainders we have for every $w \in [\neg A]$ a corresponding remainder $|[K] \cup \{w\}|$ in $K \perp A$, where $|X| = \{A : \forall y \in X : y \models A\}$, for each subset X of W. Conversely, from remainders to spheres we have for every remainder $X \in K \perp A$ a unique $w \in [\neg A]$ such that $w \in [X]$. Thus a system of spheres for K selects among the $\neg A$-worlds just in case a selection function for K takes its pick in $K \perp A$.

3. Multimodal theories of belief change

A sentence like $B \in K * A$ we read thus:

(1) After revising his beliefs so as to include A, K believes that B.

Perhaps the logical form of (1) may be analysed in a number of ways. But one particularly natural analysis would be this: There is a sentence B that stands in the skopus of an *epistemic operator* (“K believes that”) which in turn stands in the skopus of an *action operator* (“after revising his beliefs”) *parameterized* by a sentence (“so as to include A”). This suggests conducting the whole theory of belief change in a multimodal language containing a belief operator B and action operators $[a_0], [a_1], [a_2], \ldots$ with each action of the form $*\langle \mathit{sentence} \rangle$ (for revision) or $-\langle \mathit{sentence} \rangle$ (for contraction).

The idea of recasting the AGM-theory in a modal language has been aired repeatedly. Early attempts are [5] and [13]. There are presently two approaches to belief revision in a modal setting, Public Announcement and Dynamic Epistemic Logic (PAL and DEL), and Doxastic Dynamic Logic (DDL). In the remainder of this survey we give a brief exposition of DDL, as presented in [21], and then add a few remarks on PAL and DEL.

What are the advantages of presenting a theory of belief change in a modal language? First, we may help ourselves to the rich reservoir of semantical techniques available for modal languages. This is a cheap point and may indeed amount to nothing unless the techniques employed deliver clear benefits. But in one respect at least there is every reason to expect such benefits; they will be indicated in connection with the third point below.

Second, the modal language sketched above is more expressive than the set-theoretic AGM-language. We may express metabeliefs such as that the subject believes that after revision by A she will believe B: $\mathsf{B}[*A]\mathsf{B}B$. A direct translation back to AGM would render that formula as "$(B \in K*A) \in K$"—which is ill-formed. Alternatively the language of AGM would have to be stratified—a notoriously unwieldy and artificial device.

Third, the modal language is in some respects less expressive and thereby less misleading than the AGM-language. The notation $K*A$ suggests that K and A are arguments *en par* to a revision operation $*$. That is not the case. Perusing the postulates shows that nothing of interest is said about varying the belief state-parameter. Indeed, the only candidate postulate we have considered, Monotonicity for revisions, turned out to lead to desastre. The models showed why this is so. The selection functions respectively spheres determining the new, changed belief state are essentially tied to the belief state to be changed; they unfold the theoretical preferences built into that belief state. The AGM framework has no resources to study interactions or transitions between different theoretical preferences.

However, just like the AGM-language, the modal language allows to chain changes, as in

$$[-A][-B]\mathsf{B}C.$$

This is known as iterated belief change. Iterated belief change is difficult precisely because we need to know how the transition from the theoretical preferences of K to those of $K - B$ and in turn to those of $(K - B) - A$ is constrained so as to settle the question whether the latter belief state ought to include C. This is where modal semantics can be beneficial: It may deliver or suggest constraints for iterated modalities and make their evaluation more transparent.

Basic doxastic dynamics. We shall simplify the exposition of how belief change may be treated in a multimodal language by confining ourselves to a language in which belief states can only be changed by and can only contain "factual" beliefs. This simplification regrettably takes away the potential of the modal framework for a substantial theory of iterated belief change. We shall indicate prospects and problems for the more general theory in the final section.

The language is based on a set of *atoms* from which we generate the *factual formulae* by Boolean combinations. Every factual formula is a *formula*. The modal formulae are built by employing a belief and a knowledge operator, B and K, and a revision operator $[*\ ]$. If A is factual and B a formula, then $\mathsf{B}A$, $\mathsf{K}A$ and $[*A]B$ are formulae. Contractions, $[-A]B$, are defined in the spirit of the Harper Identity as $B \wedge [*\neg A]B$. Expansions may be introduced in analogy to so-called tests in dynamic logic, i.e. $[+A]B := \mathsf{B}A \rightarrow B$.

The knowledge operator K appears in **DDL** for mainly technical reasons: it will be used to signal unrevisable beliefs, those belonging to the urcorpus, to use Levi's term again. In particular, the urcorpus contains all logical truths. Note that this notion of knowledge as unrevisable belief is slightly at odds with the more usual notion that requires knowledge to entail truth.

What is the logic of these operators? Usually this depends on the kind of models one aims to describe. But in this case we have yet another clue. If we suppose that the AGM-postulates provide a reliable and complete description of belief revision, we shall be in the vicinity of the desired logic by combining the basic properties of the unary modal operators B, K and $[*A]$ (for each factual A) with translations of the AGM revision postulates into our modal language. Adjustments need then possibly be made in light of the aimed at models.

Before presenting the system of basic **DDL** (for "Doxastic Dynamic Logic") a cautionary note is in order. We know that the Ramsey Test trivializes the AGM-theory. The Ramsey Test requires a systematic connection between revisions and conditionals. The fact that we do not have a (counterfactual) conditional connective in our language for which the Ramsey Test is plausible is no sufficient reason to relax. We must rule out the possibility that some formula C in our language may step in for the conditional in the Ramsey Test, i.e. that the dangerous schema

D. $$\mathsf{B}C(A,B) \leftrightarrow [*A]\mathsf{B}B$$

be derivable for some formula C with constituents A and B. In basic **DDL** this is ruled out by the syntactic formation rules in combination with the Congruence Rule below.

The basic system of **DDL** extends classical propositional logic to a normal modal logic for each of the modalities B, K and $[*A]$ (where A is factual), i.e. it contains apart from all tautologies the following axiom schemes and schematic rules ($\Box \in \{\mathsf{B}, \mathsf{K}, [*A] : A \text{ factual }\}$, $\langle *A \rangle := \neg[*A]\neg$):

Regularity	$\Box(A \wedge B) \leftrightarrow (\Box A \wedge \Box B)$
Necessitation	$\Box\top$
Congruence Rule	$\dfrac{A \leftrightarrow B}{\Box A \leftrightarrow \Box B}$

To these we add the following translations of the AGM-postulates for revision:

Success	$[*A]\mathsf{B}A$
Inclusion	$[*\top]\mathsf{B}A \rightarrow \mathsf{B}A$
Preservation	$\neg\mathsf{B}\bot \rightarrow (\mathsf{B}A \rightarrow [*\top]\mathsf{B}A$
Consistency	$[*A]\mathsf{B}\bot \rightarrow \mathsf{K}\neg A$
Congruence	$\mathsf{K}(A \leftrightarrow B) \rightarrow [*A]\mathsf{B}C \leftrightarrow [*B]\mathsf{B}C$
Supp. 1	$[*(A \wedge B)]\mathsf{B}C \rightarrow [*A]\mathsf{B}(B \rightarrow C)$
Supp. 2	$\neg[*A]\mathsf{B}\neg B \rightarrow ([*A]\mathsf{B}(B \rightarrow C) \rightarrow [*(A \wedge B)\mathsf{B}C])$

No translation of the closure postulate is required as it is already built into the first group of axiom schemes. Finally we need to make sure that revision issues in a definite (belief) state and we give a minimal characterization of the required notion of knowledge:

Function	$\langle *A\rangle B \leftrightarrow [*A]B$
KB	$\mathsf{K}A \rightarrow \mathsf{B}A$
K*K	$\mathsf{K}A \leftrightarrow [*B]\mathsf{K}A$

Spheres semantics. The interpretation of (normal) multimodal languages generalizes that familiar from simple (normal) modal logic in a straightforward way. Each modality $\Box_i$ is matched by a binary relation R_i (on the set W of evaluation points) such that we have a clause of this form:

$$w \models \Box_i A \text{ iff } \forall v : R_i wv \Rightarrow v \models A.$$

In the case at hand, however, some modalities come with a structure which we should like to make available for logical investigation. Thus in our semantics, we should like the relations concerned to reflect that structure, as in

$$w \models [*B]A \text{ iff } \forall v : R_{[B]} wv \Rightarrow v \models A,$$

where $[B]$ is the proposition denoted by B. A clause of this form results on adapting again the spheres semantics.

A *revision structure* on a (nonempty) *universe* W consists of a *family of systems of spheres* in W and for each proposition $P \subseteq W$ a *binary relation* R_P between systems of spheres. The systems of spheres are, as before, families of subsets of W, closed under intersection (so that the system is centered on the smallest sphere, representing ones beliefs in that sphere), completely ordered under set-inclusion (so that the system is nested), and such that for each nonempty proposition P there is a smallest sphere cutting P (so that the system is sufficiently discrete). We let $\mathfrak{S} \bullet P = \{S \in \mathfrak{S} : S \cap P \neq \emptyset\}$ denote the collection of those spheres S in a system $\mathfrak{S}$ that cut P nonemptily.

The structure is subject to four conditions. It is customary to first state these conditions before proceeding to the interpretation of the modal language. Let us here change the official order of exposition; we shall then see how the conditions help the interpretation to deliver the correct results.

Some formulae, the modal ones, receive their truth-values depending on what system of spheres we are considering. For this reason the satisfaction relation $\models$ will link systems of spheres and worlds on the one hand with formulæ on the other. For atoms p and Boolean formulae, of course, only the world-coordinate contains relevant information. Let a valuation V distribute atoms over possible worlds. Then

$$\mathfrak{S}, w \models p \text{ iff } w \in V(p),$$

$$\text{and so on for } \neg, \wedge, \text{ etc.}$$

$$\mathfrak{S}, w \models \mathsf{B}A \text{ iff } \bigcap \mathfrak{S} \subseteq [A],$$

$$\mathfrak{S}, w \models \mathsf{K}A \text{ iff } \bigcup \mathfrak{S} \subseteq [A],$$

$$\mathfrak{S}, w \models [*A]B \text{ iff } \forall \mathfrak{S}' : R_{[A]} \mathfrak{S}\mathfrak{S}' \Rightarrow \mathfrak{S}', w \models B.$$

According to the clause for B, the current beliefs of an agent as evaluated in a system of spheres are represented by the innermost sphere. The clause for K displays the notion of knowledge used here: unrevisable beliefs, i.e. beliefs that stay, no matter to what sphere we shall retreat.

Now back to the four conditions. First,

$$\text{if } R_P\mathfrak{S}\mathfrak{S}',\ \text{then } \bigcap \mathfrak{S}' = \begin{cases} P \cap \min(\mathfrak{S} \bullet P), & \text{if } (\mathfrak{S} \bullet P) \neq \emptyset \\ \{\emptyset\} & \text{otherwise.} \end{cases}$$

This *defines* the revision of one's current belief state (in $\mathfrak{S}$) by the sentence representing P: it is the shaded area in the right diagram in the earlier illustration.

The second condition guarantees that systems of spheres are large enough to carry out any revision:

$$W \setminus \bigcup \mathfrak{S} = W \setminus \bigcup \mathfrak{S}' \ (\forall \mathfrak{S}, \mathfrak{S}' \in \Sigma)$$

It replaces for technical reasons our earlier and simpler requirement that each system is bounded above by W.

The final two conditions make sure that revision operations are well-defined. The first secures that they always have a value,

$$\forall \mathfrak{S} \exists \mathfrak{S}' : R\mathfrak{S}\mathfrak{S}',$$

the second that they always have at most one value:

$$\text{if } R\mathfrak{S}\mathfrak{S}' \text{ and } R\mathfrak{S}\mathfrak{S}'',\ \text{then } \mathfrak{S}' = \mathfrak{S}''.$$

Segerberg has shown that the theorems of basic **DDL** are exactly those formulæ that are true relative to all pairs $(\mathfrak{S}, w)$ in any interpretation in a revision structure.

4. Current and future changes

The transposition of the AGM-theory into a modal framework attempts to address from a different angle certain questions and problems left open in the classical version of that theory. Four such questions arise immediately.

The first three concern higher-order beliefs. If we waive the syntactic restrictions imposed in basic **DDL**, we will be confronted with three types of higher-order beliefs.

First, there will be beliefs about changes, as in $\mathsf{B}[*A]B$: Karl believes that after revising by A, B is the case, or, counterfactually: were he to believe that A, then he would believe that B, where the place of B will typically be taken by another belief formula, $\mathsf{B}C$. Now, if it is the case that $[*A]\mathsf{B}B$, then Karl can know that fact by pure introspection: the result is determined by logic plus his present theoretical preferences. What Karl can know in this way, he should know. Thus it would seem to be quite in order to impose the following condition on how the belief operator should interact with revisions:

$$\mathsf{B}([*A]\mathsf{B}B) \leftrightarrow [*A]\mathsf{B}B.$$

But this is just an instance of the dangerous schema D! Thus adding beliefs about changes to our language amounts to a non-conservative and indeed potentially hazardous extension of the theory.

Second, there will be changes by beliefs, as in $[*\mathsf{B}A]B$: If Karl were to come to believe that he believes A, then B would be the case. At first sight one would expect that the Success condition for revisions carries over to the extended language and that we thus have in particular $[*\mathsf{B}A]\mathsf{B}A$. But now consider a Moore-sentence like

m. $$p \wedge \neg\mathsf{B}p.$$

By Success we expect $[*\mathsf{B}m]\mathsf{B}m$ from which we get by Regularity both $[*\mathsf{B}m]\mathsf{B}p$ and $[*\mathsf{B}m]\mathsf{B}\neg\mathsf{B}p$. Even in very weak logics of belief, the latter gives $[*\mathsf{B}m]\neg\mathsf{B}p$, thus contradicting the former. (It suffices to add to the basic axioms above the condition that belief be consistent, $\mathsf{B}\neg A \rightarrow \neg\mathsf{B}A$, and the introspective schema $\mathsf{B}A \rightarrow \mathsf{B}\mathsf{B}A$. Jointly these conditions describe a **KD4**-type modal logic.) So one cannot (successfully and consistently) revise one's beliefs so as to include a Moore-sentence. This is as it should be. But it also shows that there are classes of formulae that must be systematically exempted from conditions that otherwise make good sense. For an elaboration of this point as applied to sentences of the form "A and I do not know that A"—which trigger the Fitch-paradox—see [7].

Third, there will be changes of changed belief states, as in $[*A]([*B]C)$: Were Karl to revise by A, then a further revision by B would make it the case that C. The models of basic **DDL** permit evaluating such formulae; but they have nothing interesting to say on iterated belief changes, just like the original AGM-theory. Iterated belief changes involve possible changes in theoretical preferences. Such changes find no direct representation in the modal languages considered so far. Thus one natural way to proceed is to introduce preference changes into the object language, as in [6].

Fourth, the methodological restriction to a single-agent perspective must eventually give way to considering belief change in a multi-agent environment. Most of our belief changes are made in response to or under the influence of the belief states of others. We have beliefs about what others believe and they are aware of this, and so on.

Beliefs about what other agents believe are the subject matter of Public Announcement Logic (PAL, see [3]) and the closely related paradigm of Dynamic Epistemic Logic (DEL, see [9, 10]). They study the rational commitments of agents who receive information by public announcements. It would be misleading, however, to think of these theories as logics of belief revision. In their basic forms they study the rather limited scenario of knowledge expansions in a multi-agent environment under the influence of an oracle. As reduction results show, PAL (and DEL) without a "looping" common knowledge modality does not go beyond a multi-epistemic logic.

But PAL is open to incorporate genuine belief revision operators; see e.g. [4]. It is to be expected that PAL, DEL, DDL and the logic of preference change will merge into a unified and general framework for studying belief change; an outline of such an integrated family of theories may be found in [6].

References

[1] Carlos Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50:510–530, 1985.

[2] Carlos Alchourrón and David Makinson. Hierarchies of regulations and their logic. In R. Hilpinen, editor, *New Studies in Deontic Logic*, pages 123–148. Reidel, Dordrecht, 1981.

[3] Alexandru Baltag, Lawrence S. Moss, and Slawomir Solecki. The logic of public announcements, common knowledge, and private suspicions. In *TARK '98: Proceedings of the 7th conference on Theoretical aspects of rationality and knowledge*, pages 43–56, San Francisco, CA, USA, 1998. Morgan Kaufmann Publishers Inc.

[4] Alexandru Baltag and Sonja Smets. The Logic of Conditional Doxastic Actions: A theory of dynamic multi-agent belief revision. In Sergei Artemov and Rohit Parikh, editors, *Proceedings of the Workshop on Rationality and Knowledge, ESSLLI 2006*. 2006.

[5] Johan F. A. K. van Benthem. *Language in Action: Categories, Lambdas and Dynamic Logic*. Studies in Logic and the Foundations of Mathematics. North-Holland, Amsterdam, 1991.

[6] Johan van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 14:1–26, 2004.

[7] Johan van Benthem. What one may come to know. *Analysis*, 64:95–105, 2004.

[8] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.

[9] Hans van Ditmarsch, W. van der Hoek, and B. P. Kooi. Playing Cards with Hintikka: An introduction to dynamic epistemic logic. *The Australasian Journal of Logic*, 3:108–134, 2005.

[10] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Springer, Dordrecht, 2008.

[11] John M. Dunn. Relevance logic and entailment. Handbook of Philosophical Logic. In D. M. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume 3, pages 117–224. Reidel, Dordrecht, 1986.

[12] André Fuhrmann. Reflective modalities and theory change. *Synthese*, 81:115–134, 1989.

[13] André Fuhrmann. On the modal logic of theory change. In André Fuhrmann and Michael Morreau, editors, *The Logic of Theory Change*, Lecture Notes in Artificial Intelligence 465, pages 259–281. Springer, Heidelberg-Berlin-New York, 1991.

[14] André Fuhrmann. Theory contraction through base contraction. *Journal of Philosophical Logic*, 20:256–281, 1991.

[15] André Fuhrmann. *An Essay on Contraction.* Studies in Logic, Language and Information. CSLI Publications, Stanford, 1997.

[16] André Fuhrmann. Everything in flux: Dynamic ontologies. In R. Sliwinski S. Lindström and J. Österberg, editors, *Odds and Ends. Philosophical Essays Dedicated to Wlodek Rabinowicz*, pages 111–125. Uppsala, 1997.

[17] Peter Gärdenfors. *Knowledge in Flux.* MIT Press, Cambridge, Mass., 1988.

[18] Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.

[19] Sven Ove Hansson. In defense of base contraction. *Synthese*, 91:239–245, 1992.

[20] Sven Ove Hansson. *A Textbook of Belief Dynamics.* Kluwer, Dordrecht, 1999.

[21] Hannes Leitgeb and Krister Segerberg. Dynamic doxastic logic: why, how, and where to? *Synthese*, 155:167–190, 2007.

[22] Isaac Levi. *The Enterprise of Knowledge.* MIT Press, Cambridge, Mass., 1980.

[23] Isaac Levi. *Decisions and Revisions.* Cambridge Univ. Press, Cambridge, 1984.

[24] Isaac Levi. *The Fixation of Belief and Its Undoing.* Cambridge Univ. Press, Cambridge, 1991.

[25] Isaac Levi. *Mild Contractions.* Cambridge Univ. Press, Cambridge, 2004.

[26] David Lewis. Probabilities of conditionals and conditional probabilities. *The Philosophical Review*, 85:297–315, 1976.

[27] Sten Lindström and Wlodek Rabinowicz. Epistemic entrenchment with incomparabilities and relational belief revision. In André Fuhrmann and Michael Morreau, editors, *The Logic of Theory Change*, Lecture Notes in Artificial Intelligence 465, pages 93–126. Springer, Heidelberg-Berlin-New York, 1991.

[28] David Makinson. On the status of the postulate of recovery in the logic of theory change. *Journal of Philosophical Logic*, 16:383–394, 1987.

[29] Charles S. Peirce. The fixation of belief. *Popular Science Monthly*, 12:1–15, 1877.

[30] Frank P. Ramsey. *The Foundations of Mathematics*. Routledge and Kegan Paul, Oxford, 1931.

[31] Hans Rott. *Change, Choice and Inference*. Clarendon Press, Oxford, 2001.

[32] Robert Stalnaker. Iterated belief revision. *Erkenntnis*, 70:189–209, 2009.