Notre Dame Journal of Formal Logic
Volume 36, Number 1, Winter 1995

Editor's Introduction

ANDRÉ FUHRMANN

> The process [by which any individual settles into new opinions] is always the same. The individual has a stock of old opinions already, but he meets a new experience that puts them to a strain.... The result is an inward trouble to which his mind till then had been a stranger, and from which he seeks to escape by modifying his previous mass of opinions. He saves as much of it as he can, for in this matter of belief we are all extreme conservatives. So he tries to change first this opinion, and then that (for they resist change very variously), until at last some new idea comes up which he can graft upon the ancient stock with a minimum of disturbance of the latter, some idea that mediates between the stock and the new experience and runs them into one most felicitously and expediently.
>
> The new idea is then adopted as the true one. It preserves the older stock of truths with a minimum of modification, stretching them just enough to make them admit the novelty, but conceiving that in ways as familiar as the case leaves possible.
>
> William James, *Lectures on Pragmatism* (1907)

1 Introduction The logic of theory change studies the formal structure of the process described informally by James in the passage above. For formal structure there is, and the structure is both simple and vexing at the same time.

The logic of theory change (alias belief change or belief revision) is of relatively recent origin. With hindsight some of the more abstract features of belief change may be seen already described in the work of the American pragmatist philosophers Peirce, James, and Dewey. The logic of theory change in its present day sense, however, originated with the investigations of Levi [20],[22], on the one hand, and of Alchourrón, Gärdenfors, and Makinson [3] (AGM) on the other hand. Since then the subject has progressed rapidly both in its pure and in its applied parts. Originally a branch of philosophical logic (with a view towards applications in epistemology, philosophy of science, or philosophical semantics) it has quickly transgressed disciplinary boundaries and is now most successful as a contribution to the logical foundations of artificial intelligence.

The domain of investigation is quickly described. There are beliefs, and there is a relation that compares their resistance to change. New beliefs may either fit in smoothly with old beliefs, or they may "put them to a strain." In the latter case the old beliefs together with the new beliefs form an *incoherent* whole. So here are three elements of a formal analysis:

1. a domain of *beliefs*,
2. a relation of *comparative resistance* between beliefs, and

– a property of stocks of beliefs, called *incoherence*.

(The last is not yet labeled because we shall replace it in a moment with a better understood notion.)

Given the current stock of beliefs B and a potential new belief p such that B and p cannot be believed together on pain of incoherence, the task is to somehow use the resistance relation to determine the maximum subcollection of B that we may allow to survive the pressure to accommodate p.

A common abstraction of "collections" of beliefs is to view them as *sets*. "Living" sets of belief have no doubt a much richer structure. For example, we acquire beliefs in a certain temporal order and from many different sources. Many beliefs are also tagged with an emotional coloring. As far as these aspects matter for how we should change our beliefs, they can be taken account of by the notion of comparative resistance to change. If "fresh" (or "long-held," or "warm") beliefs are to have priority over "stale" (or over "untried," or over "cold") ones, let them be more resistant; similarly with beliefs from more or less reliable sources. So sets-cum-resistance seems to provide a representation of belief states that is adequate for the present purpose.

Next we turn to coherence. This notion is intimately linked to that of consistency. However, not all forms of incoherence are outright cases of inconsistency. For example, the belief that a particular coin is fair is consistent with the belief that the next thousand tosses of that coin will land head up. But, so one wishes to say, the two beliefs are obviously at odds with each other.

The oddity can be explained by pointing out that inductive inference from the belief that the coin is fair would lead to a conclusion which is inconsistent with the further belief that the next thousand tosses will land head up. We must be prepared to draw such inductive inferences, for otherwise—if, say, we know that the fabulous Ricky Jay will toss the coin—the incoherence evaporates. Thus the kind of incoherence exhibited in this case can, in a sense, be reduced to the well-understood notion of (logical) consistency. At least we may assume that belief change initiated by the threat of inconsistency (before or after induction) is a an important enough phenomenon to merit investigation.

Consistency, in turn, is a property interdefinable with that of logical consequence: two beliefs can be held consistently just in case the one does not follow from the negation of the other. Hence, the third ingredient in a general theory of belief change is

3. a relation of logical *consequence*.

Most parts of the logics of theory change can be developed by appealing only to elementary (i.e., closure) properties of the consequence relation. The logic of theory change rarely taps the structural resources of the language in which the theories under consideration are formulated.

Now that we have introduced some of the abstractions underlying the formal investigation of belief change, let us turn to the kinds of belief changes that may occur. Given a decision to accommodate new information, two possibilities can arise: either the new information is consistent with what is already believed or it is not.

In the first case no rearrangement of old beliefs is necessary. The new information can be adopted into our stock of beliefs without any disturbance to the latter. This kind of change is simply additive (or "monotonic") and, as we shall see shortly, it can be easily characterized.

In the second case incisions into old beliefs are called for. Before the new piece of information can be accepted as belief, the old stock of belief has to be altered so as to make room for consistently receiving the new belief. In that case then, belief change is a two-step process: first the old stock of beliefs must be contracted so as to open one's mind towards the new candidate belief; then the new belief may be added without threat of inconsistency.

It is not difficult to see that subtracting a belief from a body of opinions is not usually a straightforward matter, for a stock of beliefs is not simply an odd bag of isolated opinions. Beliefs are usually part of a more or less tight web; the knots are enmeshed with each other by various relations of support. With many beliefs it is the case that they are supported by a set of other beliefs without being supported by any particular member of that set. When such an essentially multiply supported belief is to be retracted, at least one of the supporting beliefs needs to be removed too; otherwise the belief to be retracted will be reinstated. This situation immediately raises the question which one of the supporting beliefs to remove. Contraction thus involves an element of choice.

Given that the operation of addition, or *expansion* as we shall say, poses no challenge for a theory of belief change, the theory focuses naturally on the notion of a contraction. (In a moment we shall see that the logic of theory change can also be developed with a different operation—revision—occupying center stage.)

2 A sketch of the classical (AGM) theory The "classical" version of the theory of theory change (now usually called "the AGM theory") was first formulated in the survey paper of AGM [3]. In this paper the authors took a decision to focus on theories in the classical sense, i.e., sets of sentences (in some propositional language) closed under classical deducibility. They also presented a "canonical" set of postulates (now called "the AGM postulates") that should govern reasonable change operations. The principal construction offered was that of a partial meet contraction, in later work supplemented by modelings in terms of an epistemic entrenchment ordering on sentences. Both kinds of modelings as well as the postulates will be described below. As it turns out, the various modelings are equivalent and match with the postulates. These equivalences lend an impressive stability to the AGM theory.

First, however, we need to delineate the scope of the theory and describe in more detail the kinds of changes that may occur. Theories change in response to certain triggers of change. A trigger can be the sudden availability of new and relevant information which may strengthen, complete, or undermine the theory. "Irrational" factors, such as historical, social, or psychological pressure, may also lead to changing a theory. To prevent a common misunderstanding: an investigation of the triggers of theory change is not part of those theories collected under the heading 'logic of theory change'. Any view about the triggers of change is compatible with any version of the logic of theory change.

Also, certain "large scale" changes fall outside the scope of the present theory. Thus, meaning shifts in the terms of a theory or so-called paradigm replacements (in the sense of Kuhn and his followers) are not among the topics to be dealt with. Instead, the AGM theory is primarily about *very small changes* (as Segerberg has put it), that is, sentence-by-sentence adjustments. (However, the theory can be generalised to also cover changes by *sets* of sentences; see Fuhrmann and Hansson [9].) A change operation, in the sense of AGM, takes a theory A and a sentence α to deliver a successor to A. Three directions of change exhaust the range of possibilities:

- *expansions*: the result of expanding a theory A is a larger theory;
- *contractions*: the result of contracting A is a smaller theory;
- *revisions*: the result of revising A is a theory that is neither contained in nor contains the original theory.

Let us consider an example theory, $A = \mathrm{Cn}(p,\, p \rightarrow q)$, and ask how A may be expanded, contracted or revised.

Expansion. Suppose we obtain new information r which we should like to incorporate into our theory A. If r is consistent with A, we may just add r to A and close the result under logical consequence: $\mathrm{Cn}(A \cup \{r\})$. Thus the expansion of a theory may simply be defined in terms of the consequence operation underlying the notion of a theory.

Contraction. Suppose that we learn that q is not the case. We now need to remove q from A (resulting in a theory $A - q$). But q can be removed from A in many ways. We can either remove p or remove $p \rightarrow q$ or remove even both. Obviously the information we have about A does not suffice to nominate nonarbitrarily one of the many subtheories of A not containing q as *the* contraction of A by q.

Revision. A similar situation arises when we try to revise A. Suppose we obtain new information that $\neg q$ is the case. We should like to adjust A so as *consistently* to include $\neg q$ (resulting in a theory $A * \neg q$). As in the case of contraction, this requires an incision into the theory so as effectively to remove q from A. Again, we are faced with a choice which cannot be resolved without taking into account further, as yet unspecified, properties of A.

Whereas expansions may be defined solely in terms of logical consequence, further parameters are needed in order to resolve the choice situation faced in our simple example above. There is one constraint on resolving the choice situation which derives from the fact that information is precious and should not be discarded without necessity. This is the *Maxim of Minimal Mutilation*: keep incisions into theories as small as possible! The principal difficulty for a formal characterization of contractions and revisions is somehow to respect the maxim. There is, however, substantial disagreement about how much respect the maxim deserves. Among the AGM postulates the maxim is reflected in the condition that one may fully recover from a contraction by some sentence by adding that sentence to the contracted theory. This condition of recovery has been criticized on various grounds. Levi [23] has proposed a construction of contractions that does not in general satisfy recovery. Hansson and

Olsson (in the present issue) provide an axiomatic characterization of Levi contractions and compare them with those favored by AGM.

It will have occurred to the reader that contractions and revisions may be interdefinable. If a sentence α is to be consistently added to a theory A, then one should first make the theory consistent with α and then expand by α. That is to say, to revise A by α should amount to first contracting A by $\neg\alpha$ and then expanding the result by α. This idea is given concise expression in the *Levi Identity*.

(LI): $\quad A * \alpha = (A - \neg\alpha) + \alpha$

(Here and in the sequel we use the simpler notation, $A + \alpha$ for $\mathrm{Cn}(A \cup \{\alpha\})$, and assume that brackets associate to the left.) The Levi identity was first explicitly proposed by Levi [21], but it is implicit in many of his earlier writings.

Note that reversing the Levi identity will not in general result in satisfactory contractions. For in the principal case where A is inconsistent with the sentence α to be revised by, the classical inference *ex falso quodlibet* will explode $A + \alpha$ to the trivial theory, consisting of all sentences of the language. It is not clear how contracting the set of all formulas by $\neg\alpha$ can lead us anywhere close to where we intuitively expect $A * \alpha$ to be. Hansson [17] reverses the Levi identity in a rather special sense. He proposes first to add α to A without closing under consequence and then to retract $\neg\alpha$ from $A \cup \{\alpha\}$.

The definition of contractions from revisions is less transparent but, on reflection, equally plausible. Suppose α is to be retracted from A. The revision of A by $\neg\alpha$ will (usually) be a consistent set containing $\neg\alpha$. Hence, α will not be in $A * \neg\alpha$. But $A * \alpha$ may be bigger than the target set $A - \alpha$: for one, it will contain $\neg\alpha$. To obtain a subset of A that does not contain α we intersect $A * \neg\alpha$ with the original theory A. Thus results the *Gärdenfors Identity*.

(GI): $\quad A - \alpha = A * \neg\alpha \cap A$

For now the reader should rest content with the plausibility of the above reductions. Below we shall show that the reductions are not only plausible but indeed provable given the AGM conditions for revisions and contractions. This fact greatly simplifies the task for a theory of theory change. One may focus on either contraction or revision and then transfer results accordingly. We shall now describe two ways of constructing contractions of a given theory. The one of these approaches assumes the availability of a selection function (possibly generated from a preference relation) on the family of subtheories of the given theory. The other approach assumes that some sentences are more "entrenched" in a theory than others, an assumption which is captured by postulating an (entrenchment) ordering of all sentences with respect to some theory.

2.1 Partial meet contractions Suppose we are to contract a theory A by some sentence α. As a first approximation towards contracting without incurring loss of information beyond necessity, we may restrict attention to the *maximal* subsets of A that do not entail α. Call such subsets of A *remainders* and let $A \perp \alpha$ be the set of remainders (of A after removing α).

In all nondegenerate cases there are many such remainders; in fact, there are too many remainders to let their intersection (so-called full meet contraction) be a viable candidate for the contraction of A by α. On the other hand, picking an arbitrary remainder brings in an element of gambling where rational choice is asked for. Besides, remainders are in a way "too large" to qualify as candidates for the contracted theory, as in so-called maxichoice contraction; see AGM [3] for some negative results. As we have noted before, it appears that the consequence operation on its own—or, rather, in conjunction with familiar set-operations—does not suffice to determine that successor to a given theory that deserves the title "contraction" (by some given sentence).

At this point we simply pad our logical apparatus with a brute but natural assumption: that among a collection of alternative remainders we can somehow pick those that are, in some sense, the most preferred ones in that collection. Note that it is not assumed that the choice can always be narrowed down to uniqueness: there may be more than one most valuable remainder. Given that we have revealed our preferences by choosing a set of remainders we define $A - \alpha$ to be the set of sentences that are common to all preferred remainders, i.e.,

$$A - \alpha = \bigcap s(A \perp \alpha),$$

where s is a mapping (a selection function) from a nonempty class of theories (the remainders) into a nonempty subset of that class (the preferred remainders). We need to take care of the special case when $A \perp \alpha$ is empty. This happens just in case α cannot be removed because it is a logical truth. In that case we simply put $s(A \perp \alpha) = \{A\}$ whence $A - \alpha = A$.

This is the most general version of the partial meet recipe for constructing contractions. It is natural, however, to assume that the selection among the remainders of a theory is based on some (preference) relation. When partial meets are thus generated from a relation (by intersecting all maximal elements), we speak of a relational partial meet contraction.

2.2 Epistemic entrenchment Although all sentences in a theory must count as fully accepted, some are more accepted than others. For example, some sentences are more central to the concerns of the theory in question than others; or some sentences may be better supported than others; or for some sentences their possibility of falsehood is more remote than for certain other. There are many sources for ranking the sentences in a theory. Without going into a detailed analysis of such sources it suffices to note that some such sources lead to an ordering of sentences that will constrain the way in which a theory may change.

The approach to modeling contraction in terms of an ordering of "epistemic entrenchment" is based on the assumption that each theory may be equipped with a way of ranking its sentences (or all sentences of the underlying language) such that when it comes to choosing between candidates for removal, the least entrenched ones ought to be given up. Gärdenfors and Makinson propose to equip each theory A with a relation $\leq_A$ between single sentences such that the following conditions are satisfied:

$\alpha \leq \beta \;\&\; \beta \leq \gamma \Longrightarrow \alpha \leq \gamma$	transitivity	(EE1)
$\alpha \vdash \beta \Longrightarrow \alpha \leq \beta$	dominance	(EE2)
$\alpha \leq \alpha \wedge \beta$ or $\beta \leq \alpha \wedge \beta$	conjunctiveness	(EE3)
$\alpha \notin A \iff \alpha \leq \beta,$ if $\perp \notin A$	minimality	(EE4)
$\alpha \leq \top$	maximality	(EE5)

Observe that $\leq$ is connected,

$$\alpha \leq \beta \text{ or } \beta \leq \alpha$$

(by EE1–EE3), that

$$A = \{\alpha\colon \bot <_A \alpha\}, \text{ if } A \text{ is consistent}$$

(by EE2 and EE4), and that the maximal elements under $\leq$ are exactly the logical truths, i.e., $\mathrm{Cn}(\varnothing)$ (by EE2 and EE5). (The strict relation $<$ is defined as usual: $\alpha < \beta$ is short for: $\alpha \leq \beta$ but not conversely, $\beta \leq \alpha$.)

To obtain the contraction of some theory by a sentence α it seems natural to let α function as a cutoff point and to discard, along with α, all sentences that are no more entrenched than (“below”) α. This is, *cum grano salis*, the principal case in the following definition:

$$\text{OC} \qquad \beta \in A - \alpha \iff \beta \in A \text{ and } \begin{cases} \alpha <_A \alpha \vee \beta & \text{(the principal case), or} \\ \alpha \notin A & \text{(the vacuous case), or} \\ \top \leq_A \alpha & \text{(so } \alpha \in \mathrm{Cn}(\varnothing)\text{)} \end{cases}$$

(For technical reasons—to do with the verification of recovery—one needs the condition $\alpha < \alpha \vee \beta$ rather than the expected $\alpha < \beta$.)

We shall see below how the EE-relation for a theory may be unveiled by observing the theory’s dispositions to change.

2.3 Other modelings There are many more prima facie different approaches to modeling contractions (or revisions). Under suitable and natural assumptions all these approaches are essentially equivalent in that they pick out the same class of contraction functions. Some of the relevant representation results are surveyed in Rott [33],[34].

The method of *safe contraction* (Alchourrón and Makinson [2], Fuhrmann [7]) is, in a way, a mirror image of the partial meet approach to be introduced below. Instead of maximizing on the property of not entailing a particular sentence (as in the partial meet approach), it prunes the minimal subsets of a theory that entail the sentence to be retracted.

The *spheres models* of Grove [15] generalize Lewis’s [25] semantics for counterfactual conditionals in terms of systems of spheres of possible worlds. This approach is compared with constructions in terms of epistemic entrenchment in the contribution by Peppas and Williams.

The *minimal models* approach of Katsuno and Mendelzon [18],[19] is inspired by Shoham’s [36] modeling of reasoning from deductively insufficient premises in terms of preference relations on models (so-called preferential reasoning).

Ryan and Schobbens in their contribution below first observe a new connection between the verisimilitude orderings of models and the kind of preference orderings used in the theory of preferential reasoning. They then use a well-known link between preferential reasoning and belief revision to connect the latter with the notion of verisimilitude. Britz and Brink lend their algorithm for computing verisimilitude to belief revision, thereby providing a computational account of the latter.

Models of theory change can also be obtained from accounts of how to adjust in the light of new evidence quantitative measures of belief such as probabilities or Shackle measures (cf. Shackle [35]). Epistemic entrenchment relations can be shown to be the qualitative part of Shackle measures. This connection is much employed—independently of Shackle—in the work of Dubois and Prade [5]. Spohn [37] follows another quantitative approach. He uses conditionalizations of *ordinal conditional functions* to represent revisions of theories. In the present issue, Boutilier presents a theory of belief change using Popper functions.

2.4 The AGM postulates Given a sentential language and a consequence operation, Cn, on that language we have fixed a space of theories bounded at bottom by $\mathrm{Cn}(\varnothing)$ (the weakest theory, the set of theorems of the underlying logic) and at the top by the set of all formulas (the strongest theory, classically identical with $\mathrm{Cn}(\bot)$). Contractions and revisions are moves in the space of theories governed—according to AGM—by two families of postulates as detailed in the table below.

		Basic Postulates for Contractions
(C1)	closure	$A-\alpha=\mathrm{Cn}(A-\alpha)$
(C2)	success	$\alpha\in A-\alpha\Longrightarrow\vdash\alpha$
(C3)	inclusion	$A-\alpha\subseteq A$
(C4)	vacuity	$\alpha\notin A\Longrightarrow A\subseteq A-\alpha$
(C5)	congruence	$\alpha\dashv\vdash\beta\Longrightarrow A-\alpha=A-\beta$
(C6)	recovery	$A\subseteq\mathrm{Cn}((A-\alpha)\cup\{\alpha\})$
		Supplementary Postulates
(C7)	intersection	$(A-\alpha)\cap(A-\beta)\subseteq A-(\alpha\wedge\beta)$
(C8)	conjunction	$\alpha\notin A-(\alpha\wedge\beta)\Longrightarrow A-(\alpha\wedge\beta)\subseteq A-\alpha$
		Basic Postulates for Revisions
(R1)	closure	$A*\alpha=\mathrm{Cn}(A*\alpha)$
(R2)	success	$\alpha\in A*\alpha$
(R3)	inclusion	$A*\alpha\subseteq\mathrm{Cn}(A\cup\{\alpha\})$
(R4)	preservation	$\neg\alpha\notin A\Longrightarrow\mathrm{Cn}(A\cup\{\alpha\})\subseteq A*\alpha$
(R5)	consistency	$\neg\alpha\in A*\alpha\Longrightarrow\vdash\neg\alpha$
(R6)	congruence	$\alpha\dashv\vdash\beta\Longrightarrow A*\alpha=A*\beta$
		Supplementary Postulates
(R7)	conjunctive inclusion	$A*(\alpha\wedge\beta)\subseteq\mathrm{Cn}(A*\alpha\cup\{\beta\})$
(R8)	conjunctive preservation	$\neg\beta\notin A*\alpha\Longrightarrow\mathrm{Cn}(A*\alpha\cup\{\beta\})\subseteq A*(\alpha\wedge\beta)$

For a detailed motivation of the postulates, the reader may consult e.g., Gärdenfors [12]. Here we just give some “plausibility-enhancing” paraphrases of the conditions.

The two *closure* conditions require contractions and revisions to stay within the space of theories. Moving outside this space incurs a serious and unnecessary quality deterioration.

The *congruence* conditions require that the result of revision or contraction should not depend on syntactic properties of the sentences to be revised or contracted by; only their logical content should count.

If a sentences follows from the empty set, then it follows from any set (by the monotonicity of Cn) and is thus part of every theory. Such sentences cannot be removed. Otherwise contractions are *success*ful: a sentence to be contracted by will not be in the contracted theory. For revisions, there are two aspects to success. First, the sentence to be consistently added must be in the revised theory. Second, this is the requirement of *consistency*, the resulting theory must be consistent—again, logic permitting.

Contractions usually remove sentences from a theory; in any case, so the *inclusion* condition, they do not enlarge a theory. However, if a sentence to be retracted is not part of the theory, then the contraction operation is *vacuous*.

As with revisions, here the *inclusion* condition gives expression to the idea that revision has both an addition and a subtraction component. But sometimes the subtraction component need not be exercised: *preservation*.

A contraction is successful if the sentence α to be contracted by has been removed. However, if α follows from the empty set (i.e., is a logical truth), then it follows from *any* set (by the monotonicity of consequence). In particular, α follows in this case from any theory contracted by α. Hence, any attempt at contracting α is doomed to fail—in this sense, a contraction by α will be *vacuous*.

The vacuity condition for contractions gives already some expression to the maxim of minimal mutilation: no (proper) contraction without necessity. But the maxim really resides in the *recovery* condition. Recovery requires a contraction of A by α to be such that it can be revoked. Enough must be left in the contracted theory, $A - \alpha$, so as to recover the original theory A once we add α to $A - \alpha$.

Intersection expresses the idea that whatever survives removal of both α and β must also survive removal of $\alpha \wedge \beta$. For to remove $\alpha \wedge \beta$, it suffices to remove α or β—and if we cannot make up our minds, we shall at most remove both. Conjunction says: if we remove α (along with $\alpha \wedge \beta$), then $A - \alpha \wedge \beta$ can be no stronger than $A - \alpha$.

The corresponding supplementary postulates for revisions are motivated similarly. It will here suffice to note that *conjunctive inclusion* generalizes inclusion (let β be α and use success) and that *conjunctive preservation* would generalize preservation (let α be $\top$), if we added the condition that $A * \top = A$ for consistent A. (That latter condition is, however, a special case of preservation.)

Whenever one is confronted with a set of postulates to characterize some problematic notion, the question naturally arises: why these and not other postulates? In the present case the answer consists in showing that the chosen sets of postulates provide a *stable* characterization of contractions and revisions. That is to say, there are a number independently motivated approaches (some of them have already been outlined above) which turn out to agree in their pronouncing a mapping from theories *cum* formulas into theories a contraction, respectively a revision operation. Such representation results will be cited in a moment. But there is a second sense in which the two sets of postulates may be called stable: they support each other.

It turns out that through the Levi and the Gärdenfors identities the postulates for contractions on the one hand and the postulates for revisions on the other hand are interderivable. If a function $-$ (of the appropriate type) satisfies the AGM postulates (C1–5, 7, 8) for contractions, then the function $*$ defined from $-$ by (LI) satisfies the AGM postulates (R1–8) for revisions; and if a function $*$ (of the appropriate type) satisfies the AGM postulates (R1–8) for revisions, then the function $-$ defined from $*$ by (GI) satisfies the AGM postulates (C1–8) for contractions.

Note that the contraction postulate (C6) of recovery, $A \subseteq A - \alpha + \alpha$, is not needed for the derivation of the revision postulates from the contraction postulates via the Levi identity. However, recovery does follow from the revision postulates via the Gärdenfors identity. For further observations on the peculiar status of the recovery postulate the interested reader is referred to Makinson [28].

2.5 Representations In AGM [3] the following result is proved: for each theory A, a function $-$ (of the appropriate type) is a partial meet contraction over A if and only if it satisfies the basic contraction postulates (C1–6) for A. If one assumes that the partial meet contraction results as the meet of maximal elements under some relation, then the correspondence extends to the supplementary condition (C7) of intersection. If, in addition, the relation is assumed to be transitive, then the modeling also validates the condition (C8) of conjunction. (Warning: some commentators confuse the relations that underly the relational partial meet modeling with relations of epistemic entrenchment.)

A similar representation result holds for contraction operations as generated from epistemic entrenchment relations. Makinson and Gärdenfors [30] have shown that if $\leq_A$ is an EE-relation for a theory A, then the function $-$, defined according to (OC), satisfies conditions (C1–8) for contractions of A.

The EE-relation of a theory may be uncovered by observing how the theory behaves under contraction: β can be no more entrenched than α if one decides against α—and possibly also against β—when either α or β needs to be removed. This idea gives rise to the following connection between epistemic entrenchment and contraction:

(CO) $\alpha \leq_A \beta \iff \alpha \notin A - \alpha \wedge \beta.$

Then the reverse of the result just cited can be proved (as in [30]): if $-$ is a contraction function over a theory A, satisfying (C1–8), then the relation $\leq_A$, defined according to (CO), satisfies (EE1–5).

According to the first result, there exists a mapping which, by (OC), takes us from the class of all of EE-relations for A into the class of all contraction functions over A; according to the second result, there exists a mapping which, by (CO), takes the reverse route. The result can be strengthened to the observation (made in [30]) that the two mappings establish a bijection between the two classes of operations.

3 Extensions, generalizations and applications The classical (AGM) version of the logic of theory change is the common ancestor for a large variety of descendants. These descendants are naturally grouped according to how they depart from the AGM view.

One such departure antedates the AGM theory. In their first papers on theory change, Alchourrón and Makinson [1],[2] had not yet decided to concentrate on deductively closed sets of sentences. Their theory also covered open sets, i.e., *bases* for generating a theory by closure. Fuhrmann [6], Hansson [16], and Nebel [31] have independently pointed out that there are many advantages in retaining the more general viewpoint of a theory that also covers changes of theories as generated from a fixed base. A prominent victim of such a generalized perspective is the condition of recovery. For, if a sentence α can be removed from a theory only by removing a deductively stronger sentence (or set of sentences) from its base, then the original theory cannot be recovered by putting α back.

Whereas theories of base change pertain to the "left-hand-side" of a contraction $A - \alpha$, theories of *multiple change* loosen AGM strictures on the right by allowing for contractions by sets of sentences rather than single sentences. Two ways of multiply contracting need to be distinguished: a package contraction removes all elements of a set from a theory (or base), a choice contractions just makes sure that a set will no longer be fully contained in a theory (or base). Multiple contractions are not in general reducible to singleton contractions (not even in the finite case) and are thus operations *sui generis*. A theory of multiple contractions is presented in Fuhrmann and Hansson [9].

Lindström and Rabinowicz [26] have proposed a departure from the AGM theory in a very different direction. They waive the functionality of belief change and present a *relational theory* according to which the contraction or revision of a theory may issue in more than one offspring.

Yet a different departure concerns the representation of belief. In the AGM framework sentences are either accepted or rejected. A more fine-grained approach would allow acceptance or rejection to be a matter of degree. If theories are viewed as credal functions (either in a probabilistic or in a more general sense, like the one advocated in Shackle [35]), a theory of theory change studies, accordingly, how such degrees of belief should be changed in response to new evidence. In the present issue, Boutilier investigates this kind of belief change, covering also iterated changes.

There are also different formats in which a logic of theory can be cast. It has repeatedly been suggested (Fuhrmann [8], van Benthem [4], de Rijke [32]) that theory change may be described in the language of dynamic logic. After all, the crucial question is: which sentences hold after a fixed theory has been subjected to a certain "change program." Jaspars below takes the modal logic perspective on theory change.

Finally, an important area of application needs to be mentioned. To accept β after some background theory A has been revised by α can be seen as an inference from α to β given the background information contained in A. But note that the inference is not monotonic: revising A by α in conjunction with some other sentence may undermine the inference to β. Indeed, the nonmonotonicity is exactly of the kind that has been at the center of attention in AI. For in likening $\beta \in A * \alpha$ to an inference from α to β the theory A functions like a reservoir of default information that is used for ordinary deductions from α as far as consistency permits. This connection between belief revision and nonmonotonic inference (Gärdenfors and Makinson [13], for a slightly different connection see Fuhrmann and Levi [10]) is closely related to a connection noted much earlier (Gärdenfors [11]) between belief revision and conditionals: someone who accepts a conditional in a given state of belief undertakes a commitment to accept its consequent when he revises his state of belief by the antecedent (the so-called Ramsey test). This application of belief change is represented below in the contributions by Wansing and Wobcke.

4 Guide to the perplexed The literature on formal theories of belief change has exploded since the mid-1980s. The novice is now confronted with a dazzling stream of papers in journals and conference proceedings. In addition an increasing number of papers circulate virtually published on the Internet before they find their place in print.

There are, however, a some contributions that can act as sign-posts and should not be overlooked by anyone new to the field. Let me begin by mentioning Makinson's [27] informal introduction to the topic. In the same year AGM published their paper on partial meet modelings of theory change, which became a point of departure for the whole enterprise thereafter. The paper is fairly technical but requires no prerequisites apart from the usual. Gärdenfors [12] is the first monograph on theory change. It presents in eminently readable fashion the basic theory as well as applications that will interest philosophers more than computer scientists. Since the publication of Gärdenfors's book the subject has progressed rapidly. Gärdenfors and Rott [14] is again a self-contained survey of the field which records particularly more recent research. That research has tended to focus on the common core of a family of theories: one of them concerning theory change, others concerning conditionals, nonmonotonic reasoning, or conditional obligations. Makinson [29] studies this common core but also emphasizes relevant differences. Levi has critically commented upon mainstream research on theory change in many of his writings. He has summarized his views in two books, [23] and [24].

REFERENCES

[1] Alchourrón, C. E., and D. Makinson, "Hierarchies of regulations and their logic," pp. 125–148 in *New Studies in Deontic Logic*, edited by R. Hilpinen, Reidel, Dordrecht, 1981. MR 83a:03014 3

[2] Alchourrón, C. E., and D. Makinson, "On the logic of theory change: contraction functions and their associated revision functions," *Theoria*, vol. 48 (1982), pp. 14–37. Zbl 0525.03001 MR 85g:03023 2.3, 3

[3] Alchourrón, C. E., P. Gärdenfors and D. Makinson, "On the logic of theory change: partial meet contraction and revision functions," *The Journal of Symbolic Logic*, vol. 50 (1985), pp. 510–530. Zbl 0578.03011 MR 87c:03020 1, 2, 2.1, 2.5

[4] van Benthem, J., *Language in Action*, North-Holland, Amsterdam, 1991. Zbl 0743.03018 MR 92g:03002 3

[5] Dubois, D., and H. Prade, "Belief change and possibility theory," pp. 142–182 in *Belief Revision*, edited by P. Gärdenfors, Cambridge University Press, Cambridge, 1992. MR 1173427 2.3

[6] Fuhrmann, A., *Relevant Logics, Modal Logics, and Theory Change*, Ph.D. thesis, Australian National University, Canberra, 1988. 3

[7] Fuhrmann, A., "Theory contraction through base contraction," *Journal of Philosophical Logic*, vol. 20 (1991), pp. 175–203. Zbl 0723.03010 MR 92g:03013 2.3

[8] Fuhrmann, A., "On the modal logic of theory change," pp. 259–281 in *The Logic of Theory Change*, edited by A. Fuhrmann and M. Morreau, LNAI 465, Springer-Verlag, Berlin, 1991. Zbl 0925.03108 MR 1096799 3

[9] Fuhrmann, A., and S. O. Hansson, "A survey of multiple contractions," *Journal of Logic, Language and Information*, vol. 3 (1994), pp. 39–76. Zbl 0791.03014 MR 95i:68129 2, 3

[10] Fuhrmann, A., and I. Levi, "Undercutting and the Ramsey Test for conditionals," *Synthese*, vol. 101 (1994), pp. 157–169. Zbl 0976.03016 MR 95i:03037 3

[11] Gärdenfors, P., "Conditionals and changes of belief," *Acta Philosophica Fennica*, vol. 30 (1978), pp. 381–404. Zbl 0413.03011 3

[12] Gärdenfors, P., *Knowledge in Flux*, MIT Press, Cambridge, 1988. MR 89k:00015 2.4, 4

[13] Gärdenfors, P., and D. Makinson, "Nonmonotonic inference based on expectations," *Artificial Intelligence*, vol. 65 (1994), pp. 197–245. Zbl 0803.68125 MR 95d:03045 3

[14] Gärdenfors, P., and H. Rott, "Belief revision," pp. 35-132 in *Handbook of Logic in Artificial Intelligence and Logic Programming*, vol. 4, edited by D. M. Gabbay et al., Clarendon Press, Oxford, 1995. MR 97h:03053 4

[15] Grove, A., "Two modellings for theory change," *Journal of Philosophical Logic*, vol. 17 (1988), pp. 157–170. Zbl 0639.03025 MR 89e:03027 2.3

[16] Hansson, S. O., "New operators for theory change," *Theoria*, vol. 55 (1989), pp. 114–132. Zbl 0722.03028 MR 92i:03010 3

[17] Hansson, S. O., "Reversing the Levi identity," *Journal of Philosophical Logic*, vol. 22 (1993), pp. 637–669. Zbl 0788.03033 MR 94i:03066 2

[18] Katsuno, H., and A. O. Mendelzon, "Propositional knowledge base revision and minimal change," *Artificial Intelligence*, vol. 53 (1991), pp. 263–294. Zbl 0792.68182 MR 93h:68133 2.3

[19] Katsuno, H., and A. O. Mendelzon, "On the difference between updating a knowledge base and revising it," pp. 183–203 in *Belief Revision*, edited by P. Gärdenfors, Cambridge University Press, Cambridge, 1992. Zbl 0765.68197 MR 1142168 2.3

[20] Levi, I., *Gambling with Truth*, Knopf, New York, 1967. 1

[21] Levi, I., "Subjunctives, dispositions and chances," *Synthese*, vol. 34 (1977), pp. 423–455. Zbl 0364.02002 2

[22] Levi, I., *The Enterprise of Knowledge*, MIT Press, Cambridge, 1980. 1

[23] Levi, I., *The Fixation of Belief and Its Undoing*, Cambridge University Press, Cambridge, 1991. 2, 4

[24] Levi, I., *For the Sake of the Argument*, forthcoming from Cambridge University Press. Zbl 0935.03001 MR 98a:03030 4

[25] Lewis, D., *Counterfactuals*, Blackwell, Oxford, 1973. MR 54:9979 2.3

[26] Lindström, S., and W. Rabinowicz, "Epistemic entrenchment with incomparabilities and relational belief revision," pp. 93–126 in *The Logic of Theory Change*, edited by A. Fuhrmann and M. Morreau, LNAI 465, Springer-Verlag, Berlin, 1991. Zbl 0925.03128 MR 1096792 3

[27] Makinson, D., "How to give it up," *Synthese*, vol. 62 (1985), pp. 347–363. MR 86m:03024 4

[28] Makinson, D., "On the status of the postulate of recovery in the logic of theory change," *Journal of Philosophical Logic*, vol. 16 (1987), pp. 383–394. Zbl 0632.03008 MR 89b:03019 2.4

[29] Makinson, D., "Five faces of minimality," *Studia Logica*, vol. 52 (1993), pp. 339–379. Zbl 0790.03026 MR 95b:03029 4

[30] Makinson, D., and P. Gärdenfors, "Revisions of knowledge systems using epistemic entrenchment," in *TARK II*, edited by M. Vardi, Morgan Kaufmann, Los Altos, 1988. Zbl 0711.03009 MR 1011079 2.5, 2.5, 2.5

[31] Nebel, B., "A knowledge level analysis of belief revision," pp. 301–311 in *Principles of Knowledge Representation and Reasoning*, edited by R. Brachman et al., Morgan Kaufmann, San Mateo, 1989. Zbl 0709.68091 MR 1011061 3

[32] de Rijke, M., *Extending Modal Logic*, Ph.D. thesis, Institute for Language, Logic and Information, Amsterdam, 1993. 3

[33] Rott, H., "Two methods of constructing contractions and revisions of knowledge systems," *Journal of Philosophical Logic*, vol. 20 (1991), pp. 149–173. Zbl 0723.03009 MR 92g:68136 2.3

[34] Rott, H., "On the logic of theory change: more maps between different kinds of contraction functions," pp. 122–141 in *Belief Revision*, edited by P. Gärdenfors, Cambridge University Press, Cambridge, 1992. MR 1173426 2.3

[35] Shackle, G. L. S., *Decision, Order and Time in Human Affairs*, second edition, Cambridge University Press, Cambridge, 1969. 2.3, 3

[36] Shoham, Y., *Reasoning About Change*, MIT Press, Cambridge, 1988. MR 90b:68081 2.3

[37] Spohn, W., "Ordinal conditional functions: a dynamic theory of epistemic states," pp. 105–134 in *Causation in Decision, Belief Change, and Statistics*, vol. 2, edited by W. L. Harper and B. Skyrms, Reidel, Dordrecht, 1988. 2.3

Philosophische Fakultät
Postfach 5560
78434 Konstanz
Germany
email: Andre.Fuhrmann@uni-konstanz.de