Ways of Doing Logic: What was Different about AGM 1985?

David Makinson

Department of Computer Science, King's College London

This brief review does not give any new results, nor any new proposals for modelling belief change. In principle, everything that it says, readers familiar with the logic of belief revision will know already – or disagree with already. My intention is to stand back and reflect on what, in 1985, was new or different about AGM belief revision as away of doing logic, so as to encourage discussion on what we might profitably be doing now.

A bit of history

The logic of belief change and nonmonotonic logic are both creatures of the 1980s. As everyone knows now, they have close relations to each other, as well as relations, rather less close, to updating and to counterfactual conditionals. But the exploration of the connections came *after* the separate formulations. This, I think, is as it should be. I am rather suspicious of new subjects being created by translation from another. The best order is: self-sufficient creation, then maps with neighbours, possibly ending with full translations.

The development of these subjects has followed a pattern of overlapping sequences that is normal in such affairs. First, there were journal papers on one or the other of the two subjects.

Then books appeared, collecting together conference papers. For belief revision, these began with a volume edited by André Fuhrmann and Michael Morreau in *The Logic of Theory Change* (1991), followed a year later by one edited by Peter Gärdenfors, *Belief Revision*. The most recent such collection, which appeared this year 2001 is *Frontiers in Belief Revision*, edited by Mary-Anne Williams and Hans Rott.

A number of journals devoted special issues to belief change– *Notre Dame Journal of Formal Logic* (1995), *Theoria* (1997), *Journal of Logic, Language and Information* (1998), and *Erkenntnis* (1999). Overviews of work on the subject appeared in 1995 and 1998 in two of the *Handbooks* that Dov Gabbay edited, following an overview of nonmonotonic reasoning in a *Handbook* of 1994.

At the same time, individually authored books began to appear in the two areas. Perhaps the first was Peter Gärdenfors' *Knowledge in Flux* of 1988, on belief contraction and revision. In the 1990's, the focus moved to nonmonotonic logic - in 1991, Gerhard Brewka's book *Nonmonotonic Reasoning*, then Karl Schlechta's *Nonmonotonic Logic* and Grigori Antoniou's *Nonmonotonic Reasoning*, both in 1997. Then in 1997, returning to belief change nearly a decade after Gärdenfors' first book, we have André Fuhrmann's *An Essay on Contraction*, followed in 1999 by Sven Ove Hansson's *Textbook of Belief Dynamics* - the first on either of these subjects to be written as a regular textbook for students, with exercises and answers.

I must have left some books out, and I hope that nobody is offended by their absence, as I am not trying to be exhaustive.

In 2001 two further single-authored volumes have appeared, and they seem to be announcing a new way of approaching the area. Each book mentioned so far dealt either with belief change, or with nonmonotonic reasoning; the focus was squarely on one even if touching tangentially the other. But now, there is a deliberate attempt to bring them together, treating them side by side and as far as possible, with the same concepts - even common formal structures. I am thinking of Alexander Bochman's *A Logical Theory of Nonmonotonic Inference and Belief Change*, and Hans Rott's *Change, Choice and Inference: A Study of Belief Revision and Nonmonotic Reasoning*.

Thus at present we have two streams of action: continued research on specific problems and concepts, and books that try to form broad pictures with a coherent perspective and an orderly internal development. In these, belief revision and nonmonotonic logic are coming to be seen as chapters of a general theory of belief management.

Evidently, this broad perspective is helpful, but I am not sure yet whether the two areas should be presented as instances of a common *formalism*. That may hide interesting differences of gestalt and intuition hiding underneath formal similarities – I will mention one later. But perhaps I will be proven wrong – perhaps a general theory will emerge, covering both, of sufficient elegance and transparency to make the abstraction worthwhile.

Back in 1985

Thus far is recent history and current affairs, but let's go back to the paper that first put the formal logic of belief change into the public arena, AGM - "On the logic of theory change: partial meet contraction and revision functions" written by Carlos Alchourrón, Peter Gärdenfors and myself and published in 1985.

Of course, that paper was not born out of nothing. It was preceded by several earlier papers of its three authors, writing in various combinations. One of us, Gärdenfors, came to the area from epistemology and the philosophy of science, more specifically the study of the evolution of scientific theories, influenced by work of William Harper and Isaac Levi going back to the 1970's. Alchourrón and I came to it from problems in the philosophy of law, and in particular of abrogating an item from a legal code. As a result, Gärdenfors focussed primarily on revision, while Alchourrón and I thought foremost in terms of contraction. The encounter took place when carlos and I submitted a paper on what is now known as choice contraction to the journal *Theoria*, which was then edited by Peter. Of course, we were all influenced by - and at the same time attempting to escape from - the possible-worlds approach to non-classical logics, and in particular the account of counterfactual conditional given by David Lewis in his 1973 book *Counterfactuals*.

Work in nonmonotonic logic was already going on for several years before AGM was written, though on the level of rather specific systems. Remember that the notion of circumscription, a parent of the later and more general concept of preferential inference, was introduced by John McCarthy in a paper published as early as 1980. Remember also that Ray Reiter's extraordinary paper "A logic for default reasoning" was published in the same year, 1980. As far as Peter and I can remember, when writing AGM 1985 none of us was aware of those two important papers.

Anyway, the point is that AGM 1985 came into existence following and alongside seminal work in nonmonotonic reasoning. And what I am going to say about AGM applies, to a large degree, to early work on nonmonotonic reasoning as well, for they broke with the past in similar ways. I will try to put my finger on some of the shifts of perspective.

Very roughly, I will group them into changes of a philosophical nature, heuristic ones, and finally some important technical ones. This is simply for convenience, no weight is put upon the division. I will try to indicate each shift by a slogan, and then explain and discuss it.

Philosophical Perspectives

Logic is not just about deduction

Not all inference is deductive. Except for formal logicians, everybody has known that for a long time – mathematicians working with probabilities, as well as detectives, lawyers and garage mechanics. But only with the development of nonmonotonic logic did this evident fact become a subject of serious study by logicians.

Nor can all knowledge management be called inference. To begin with, belief revision operations $K{*}a$ are not inference operations, although their projections $*_K(a)$ on the input argument are, as noticed by Gärdenfors and myself in 1991. Evidently, belief contraction operators $K\text{-}a$ are even further removed from inference: they take information away rather than add it in.

Another example may be found in the logic of conditional goals or obligations. We may entertain a condition without it becoming automatically a goal. The passage from condition to goal is not an inference, and the rule of identity or inclusion fails. Deontic logic has had difficulties coping with this, as it has tended to be under the influence of the inference paradigm and to offer constructions that make anything a goal when it is taken as its own condition. To deal with conditional goals, Leendert van der Torre and I recently introduced what we call input/output logics. But that is another story…. The important point here is that both AGM and nonmonotonic logic both participated in opening the horizons of logic beyond the limits of deductive inference, and over the borders of inference itself.

There is nothing wrong with classical logic

It is important to bear in mind that when we devise logics for belief change or nonmonotonic reasoning, we are not objecting to any classical principles. In this, the enterprise is quite different from that of the relevantists, or the intuitionists. We do not see ourselves as fabricating non-classical logics, but rather as offering a more imaginative use of classical logic

Indeed, as is evident from the first few pages of AGM, the logic of belief revision is formulated assuming a monotonic consequence operation in the background. The same is true of most approaches to nonmonotonic logic, for example that of preferential inference. In some instances, authors have accepted classical logic in a rather different way. For example Daniel Lehmann and some other authors have constructed systems for nonmonotonic inference in which classical consequence does not figure explicitly, but emerges implicitly from the nonmonotonic principles.

In the AGM presentation, the background monotonic consequence operation may be taken to be classical consequence itself, but not only so. It can be any supraclassical consequence operation satisfying all the Tarski conditions (inclusion, monotony, idempotence) plus disjunction in the premises plus compactness. Equivalently: any operation formed from classical consequence by adding a fixed set of extra premises.

These operations do not have a standard name. They could be called *paraclassical* consequence operations. Mathematically, their investigation is quite trivial given knowledge of the behaviour of classical consequence, but tactically they play an interesting intermediate role between classical consequence and say preferential inference relations.

Heuristic Orientations

<u>*Look at Relations, not sets*</u>

It is interesting to compare the shifting objects of attention of logicians from the mid-nineteenth century to the present day.

For Boole, logical principles were typically formulated as *identities* or *equivalences* between propositions. Frege made a major innovation when he centred study on a distinguished *set* of propositions – those that are necessarily true. In the early twentieth century, some logicians began to think of logic as the study of a *relation* between propositions - one of deducibility or inferability. It is not clear to me to whom this perspective should be credited, if indeed to any one person, but I suspect that C.I. Lewis was involved. As is well known, Tarski introduced the idea of looking at inference as an *operation* on sets of propositions, which gathers together the results of inference. And Gentzen was the first to focus attention on *multiple-conclusion* relations with the multiplicity of the premise sets read conjunctively and that of the conclusion sets read disjunctively, giving a full symmetry between left and right. These were in turn considered as operations by Dana Scott.

I think that it is fair to say that Frege's approach dominated mathematical logic for the first half of the twentieth century, and it still plays in centre court. The use of a relation or an operation of consequence of course dates back to the 1930s, but it was usually seen as a special-purpose or alternative presentation, rather than the official one.

Of course, in classical logic we can move freely between the distinguished set of necessarily true propositions and the corresponding inference relation, by putting a |- x iff the material conditional proposition $a \rightarrow x$, i.e. $\neg a \lor x$, is in the distinguished set. But when we get into nonmonotonic reasoning, it turns out that this reduction is no longer available. In belief contraction and revision, if one asks for a distinguished set, one could take it to be whatever belief set is being contracted or revised, but these are arbitrary; what we are really interested in are the *operations* that take them into contracted or revised sets.

Thus once we go beyond classical logic, there is a substantive difference between a logic formulated in terms of a distinguished set of propositions and one presented through a relation between propositions. On the other hand, the difference between relations between formulae and operations on sets of formulae seems to remain essentially one of notational convenience – at least, so long as the structures investigated are compact, so that whatever we want to say about a infinite set of premises can be expressed in terms of its finite subsets and thus in terms of conjunctions of its elements.

However, compactness does not always hold for, say, preferential models, and this forces a procedural choice when we treat preferential inference from infinite sets of premises. We can apply the preferential definition to all cases, finite or infinite, or else apply it only to the finite case and then define the operation in the infinite case by the compactness biconditional. I am not sure whether one procedure is better than the other, but something makes me prefer the former as more principled.

There remains a question on which, I believe, discussion is still open. Is there any real advantage in working with *multiple-conclusion* consequence operations as the classical backdrop in belief revision or nonmonotonic reasoning? Bochman's recent book, mentioned above, is exceptional in that it takes multiple-conclusion consequence as its official framework. This permits him to develop a certain approach to iterated revision; but he also shows that for one-shot belief change and nonmonotonic inference, the single-conclusion Tarski consequence operations suffice. The benefits of working with multi-conclusion consequence operations may thus be closely related to the approach to iterated revision that one adopts.

Don't internalise too quickly

It was a tradition in philosophical logic for much of the twentieth century to represent, whenever possible, a logical notion as a propositional connective in the object language, alongside the truth-functional ones. Examples include: strict implication (and modal logics formulated with necessity and possibility as connectives), entailment and relevance logics, counterfactual conditionals, most multi-valued logics, and linear logic.

The view of AGM is quite different: treat contraction and revision as *operations in the metalanguage*. Likewise, most presentations of nonmonotonic inference treat it as as a *relation* between propositions, or as an *operation* on propositions. The idea is: do not prematurely internalise the key relations and operations of the metalanguage, to become connectives (often called operators) of the object language. In brief: operations *sí*, operators *no*!

Thus in nonmonotonic reasoning, the closure conditions on snake are formulated in English. For example, cumulative transitivity says: whenever $a \mid\sim x$ and $a \wedge x \mid\sim y$ then $a \mid\sim y$. It is not presented as a distinguished object-language formula $((a \mid\sim x) \wedge (a \wedge x \mid\sim y)) \rightarrow (a \mid\sim y)$. And no iterated snakes are considered.

This is not for philosophical reasons. The rationale is not like that of Quine's strictures against modal logic, where he doubted the possibility of giving any coherent meaning to iterated modalities. Our reason is a methodological one: get the flat case right and understand it well. Then the first-degree case (where we treat the relation as a connective, to which Boolean operators may be applied but which cannot be iterated) should more or less look after itself. The iterated case should be left to last, as it is likely to introduce all sorts of complications and multiple options, some of which may be rather baroque, artefacts of our ingenuity without much practical significance. Get the flat picture quite clear before going into more complex formulae, particularly with iterated applications.

Such was the strategy of AGM. But it must be granted that in that very case, i.e. the logic of belief change, it leaves a big gap to be filled. For it is very natural to think of belief change as a process that in ordinary life we routinely iterate, with the product of the last change becoming the origin for the next one. In recent years, there has been considerable effort to fill this gap, and I think that this has been timely. The strategy of AGM was intended as an initial, not an eternal one. However, I am not sure whether we yet have a consensus on the way to handle iterated belief change, or still a lot of disparate proposals in the literature without a clear idea of what is preferable.

Curiously, in the case of nonmonotonic reasoning there is much less motivation for iteration, despite the existence of a formal map between the two domains in the non-iterative case. The reason seems to be a pragmatic feature - what we want to hold constant and what we want to vary differ between the two domains. Bochman has described nonmonotonic inference as "static", in the sense that the premise entertained need not lead to any loss of background beliefs with which they are inconsistent. No background beliefs are abandoned, but some of them are *left unused* in drawing the inference. On the other hand, contraction and revision are "dynamic" in that no belief (beyond the tautologies) need be held through successive changes. Isaac Levi has made essentially the same point in more philosophical terms. This is an example, I think, where over-attention to formal translations can lead us to ignore underlying non-formal differences.

Even on a formal level, differences appear to arise once we get into iteration. An iterated revision is typically something like $(K{*}x){*}y$, while an iterated nonmonotonic inference is typically something like $(x \mid\!\sim y) \mid\!\sim z$ or $x \mid\!\sim (y \mid\!\sim z)$. But these don't seem to have much to do with each other under the Makinson-Gärdenfors translations of 1991, even when for simplicity $K = Cn(k)$ for some individual proposition k. The translation of $z \in (K{*}x){*}y$ into the language of nonmonotonic inference would be $y \mid\!\sim_J z$ where $J = C_K(x) = \{u\colon x \mid\!\sim_K u\}$. On the other hand, the translation of $(x \mid\!\sim y) \mid\!\sim z$, say, into the language of belief revision would be $z \in K{*}(x \mid\!\sim y)$. It is not clear to me, at least, how these relate to each other. In sum, iterated belief revisions do not seem to correspond neatly to iterated nonmonotonic inferences, and the Makinson-Gärdenfors translations appear to be appropriate only for the non-iterative case.

<u>*Do some logic without logic*</u>

What I mean is that in logical investigations there are often a number of interesting questions that arise *before* we begin to consider the presence of any connectives at all in the object language. For example, if we work with inference operations (rather than relations) we can formulate principles like cumulative transitivity and cautious monotony in a way that does not refer to any object-language connectives – not even conjunction.

Of course, this might be regarded as self-delusion or cheating, on the ground that we are in effect making the process of grouping premises into sets simulate the act of conjunction. With multiple-conclusion logic, we can likewise simulate disjunction. Even without the multiple-conclusion format, we can make the intersection of premise sets that are closed under classical consequence simulate the disjunction of their respective elements. But even if there is an element of self-deception, the effort is still an interesting one, in so far as it helps to get a picture clear in the purest context, without any distractions at all.

That is the perspective that I took in my first paper on nonmonotonic reasoning, "General theory of cumulative inference", published in 1989, and I think it was useful. Kraus, Lehmann and Magidor then incorporated truth-function connectives into the object of analysis, in their well-known paper of 1990, thus taking the analysis much further.

This is what also makes the theory of defeasible inheritance, and the more complex theory of argumentation, so fascinating. We have some quite intricate problems of iterated defeat arising before we even begin to think about object-language connectives. Much the same is true of the "connective-poor" inference using logic programs with negation (LPN), under the semantics devised by Gelfond and Lifschitz in 1988, or equivalently the justification-based truth maintenance systems (JTMS) of Doyle back in 1979, each compared with the "connective-rich" Reiter default logic.

Technical Differences

I do not want to get lost in details, but there are also some technical matters that make an enormous difference to our gestalt.

There is no unique object of attention

This is very important, because it is something that disorients many of those coming to nonmonotonic reasoning or belief change for the first time, especially when they have become used to classical and its well known sublogics.

For deduction, we are familiar with the idea that there is just one core logic, up to notational differences and matters like choice of primitives. That core is classical logic, and it is also the logic that we use when reasoning ourselves in the metalanguage. Even intuitionists and relevantists, who do not accept all of classical logic, feel the same way, but about their own systems. They have some difficulties, however, in reconciling this view with their own practice in the metalanguage.

So it is natural for the student to ask, which is *real* nonmonotonic inference? Which is the *correct* operation of belief contraction? What is the one that we use in practice, even if we can study others? The early literature on nonmonotonic reasoning tended to pose the same questions.

The answer is that there is none. There is not a unique belief contraction operation on a belief set, but indefinitely many of them. They are all those operations satisfying certain syntactic conditions such as inclusion and recovery (i.e. the AGM or other postulates), or equivalently, all those that can be constructed as partial meet contractions (or similar means) where the selection function or the minimizing relation is allowed to vary freely so long as it satisfies certain formal conditions. This intrinsic non-uniqueness was, I think, a rather new feature of AGM belief revision, and perhaps responsible for some of the initial difficulties of assimilation by logicians.

It is not like saying that there are a lot of different modal logics, K, S4, S5 etc, according to the list of postulates that we wish to place on the necessity operator. We are saying more than that. Even if we agree to work with a fixed set of conditions for contraction – say the extended AGM postulates or the corresponding family of transitively relational partial meet contractions – we still have not one but infinitely many such operations. The postulates, and likewise the construction, do not determine a unique operation, but designate the limits of a class of them.

Exactly the same is true of preferential inference relations: what may be inferred from what depends on the particular preference relation chosen for the model. This is not the familiar point that we can have many extensions of a premise set, determined by the many minimal states that may satisfy them, just as in Reiter default logic a premise set can have many extensions. For even when the preference relation is assumed to be linear, so that there is no more than one minimal state satisfying a given set of premises, the identity of that state will still depend on the particular preference relation chosen, and so the set of conclusions that are supported will also depend on it.

Moreover, if one tries to get away from this kind of non-uniqueness by intersecting all the many relations or operations, the result is just classical logic. The intersection of all preferential snakes is classical gate; the intersection of all AGM contractions of x from a belief set K closed under classical consequence is just $K \cap Cn(\neg x)$.

Not closed under substitution

Another feature that made logicians uncomfortable is that neither the operations of belief change, nor those of non-monotonic inference, are closed under substitution. To be precise, they are typically not closed under substitution of arbitrary formulae for elementary ones.

Thus when $|\sim$ is a nonmonotonic inference relation, $a \mid\sim x$ does not in general imply $\sigma(a) \mid\sim \sigma(x)$ for substitution functions σ. Likewise, $x \in K - a$ does not imply $\sigma(x) \in K - \sigma(a)$ or anything like that. Moreover, as long known as part of the folklore, the same failure already arises for all of the paraclassical consequence operations that I mentioned earlier, except for the total one under which every formula is a consequence of every formula. In this respect, our operations are quite unlike the consequence operation of classical logic, or of any of the better-known subclassical logics – intuitionistic, relevantist, many-valued, etc.

Thus these operations are not *structural*, in the sense of the term used in Wojcicki's 1988 book on *Theory of Logical Calculi*. Depending on one's terminology, one might not even like decorating them with the words "formal", "calculi" or "logical" at all – such is the weight traditionally given to substitution and the notion of logical form. The role of substitution was not always explicit, but it was often presumed to be constitutive of what deserves to be called an object of formal logic.

Failure of substitution has a number of consequences. In particular, it means that the algebraic method of investigation, so successful in the analysis of subclassical, modal and even relevantist logics (as in Rasiowa and Sikorski's 1963 book *The Mathematics of Metamathematics*, for example), is not likely to have much of a run in the analysis of belief revision and nonmonotonic reasoning.

The essence of the algebraic method was to define consequence as preservation of distinguished values under valuations, understood as homomorphisms on the algebra of formulae into algebras of a variety. Such an approach will validate substitution, for substitution is itself an endomorphism on the algebra of formulae, so that a valuation of a substitution instance of a formula is the composition of two homomorphisms, and thus itself a valuation of the initial formula.

There sense in which we can recover substitution, and at the same time a unique object of attention. If in the end we do internalise our operation or relation, allowing unlimited iteration, then we can present the logic of belief change, or the logic of nonmonotonic inference, as a logical system of the traditional kind. That is one of the seductions of internalisation.

It should be noted however that even then, the unique object of attention will not be some specific preferential inference relation, or a distinguished contraction or revision operation, but rather an algebra of formulae, essentially a free algebra in a variety. Moreover, if we define a relation R to hold between boolean formulae a and x iff the non-boolean formula $a \mid\sim x$ is a theorem of that logic, then R will be just classical consequence. So the unique relation that we get is not one that was wanted.

Out through the door, back in the window

This is not to say that substitution plays no role at all in the logic of belief change or nonmonotonic reasoning. It intervenes at a higher level – just as do monotony and compactness.

To explain this as simply as possible, consider the case of preferential inference relations. Let R be an inference relation between sets A of propositions and individual propositions x, and let R^+ be the intersection of all preferential inference relations R' that include R. Then R^+ is itself a preferential inference relation, since the conditions defining such relations are all Horn. Moreover, the operation taking R to R^+ is a closure operation in the standard mathematical sense of the term (i.e. it satisfies inclusion, idempotence and monotony), although its values R^+ are not closure relations, since they are not in general monotone.

It is not difficult to show that this operation, commonly called preferential closure, is also compact. It is also closed under substitution, in the sense that $\sigma(R^+) \subseteq (\sigma(R))^+$. In other words, whenever $(A,x) \in R^+$ then $(\sigma(A), \sigma(x)) \in (\sigma(R))^+$. Care should be taken when reading this property: the last term in it is not R^+ but $(\sigma(R))^+$. If we wrote the former we would be saying that the value R^+, which is a preferential inference relation, is itself closed under substitution – which in general is not the case.

Thus, on the level of individual preference relations, monotony, compactness and substitution all fail, but the function that takes an arbitrary inference operation to the least preferential operation including it, is nevertheless a compact closure operation satisfying substitution. What we threw out the door comes back in the window the next floor up. To speak in riddles, these properties can reappear even when they are not there – but not always, as we shall see below.

Non-Horn rules

In the study of belief change and in nonmonotonic reasoning, we find another phenomenon emerging more saliently than it did in the past – the existence of interesting non-Horn rules. This can be illustrated from either of the two areas.

The AGM supplementary postulate (K-7) for contraction tells us that $(K\text{-}a)\cap(K\text{-}b) \subseteq K\text{-}(a\wedge b)$. This is a Horn condition, as it says: whenever $x\in (K\text{-}a)$ and $x\in (K\text{-}b)$ then $x\in K\text{-}((a\wedge b)$. But the other AGM supplementary postulate (K-8) tells us that if $a \notin K\text{-}(a\wedge b)$ then $K\text{-}(a\wedge b) \subseteq (K\text{-}a)$. In other words, whenever $x \in K\text{-}(a\wedge b)$ and $a \notin K\text{-}(a\wedge b)$ then $x \in (K\text{-}a)$. This is not a Horn condition, as it has a negative premise. Equivalently, of course, it can be expressed with a disjunctive conclusion: whenever $x \in K\text{-}(a\wedge b)$ then either $x \in (K\text{-}a)$ or $a \in K\text{-}(a\wedge b)$.

The same sort of thing arises notoriously for rational monotony, i.e. the rule that says: whenever $a \mid\!\sim y$ and $a \mid\!\not\sim \neg x$ then $a\wedge x \mid\!\sim y$ – not surprisingly, for there are close relations between (K*8) and rational monotony under the Makinson-Gärdenfors translations. Moreover, there is a whole family of related non-Horn conditions on inference relations that have been studied by various authors, including Ramón Pino Pérez, Hassan Bezzazi and myself in a paper of 1997.

Such quasi-closure conditions, which resemble Horn conditions except that they have a negative premise (or equivalently, a disjunctive conclusion), have played an important and even defining role in the logic of belief change since AGM 1985 itself, as also in nonmonotonic logic. In contrast, I don't remember earlier logical systems giving much attention to such conditions, and certainly not a central place.

These conditions also give rise to quite difficult mathematical questions. The best known is that of defining the rational "closure" of a preferential inference relation, i.e. the "least" among the preferential relations that include the first one and at the same time satisfy the non-Horn condition of rational monotony. I put the word "closure" between pincers because this is not a closure operation in the standard mathematical sense of the term, since the intersection of relations satisfying the condition of rational monotony does not in general satisfy that condition. Indeed, the intersection of all rational relations that include a given non-rational one is never rational, as Lehmann and Magidor have shown. Thus also the word "least" cannot mean least with respect to set inclusion.

Nevertheless, Lehmann and Magidor have given an interesting proposal for what exactly such a rational closure should be, in their 1992 paper "What does a conditional knowledge base entail?". Many researchers find this account convincing, but others do not; some would say that there is not much point in searching for the "right" definition of rational closure, because in general there is no such thing.

Evidently, the same questions can be raised for operations of belief revision, passing from an operation satisfying the basic AGM postulates, say, to one also satisfying the two supplementary postulates (K*7) and (K*8), of which the latter is non-Horn. Presumably the same construction can be made, although I am not sure whether it would be equally convincing that it is the "right" one, nor whether the question itself is of equal interest in the context of belief change, despite the formal maps between the two.

I have a vague feeling that the subject of non-Horn conditions, and of "right" extensions of a relation or operation satisfying them, still need investigation, both in the present context and in quite general terms, so as to give us a clearer picture of what is desirable and what is possible. But facts are one thing and vague feelings another, so I stop here.

Acknowledgements and Apologies

This paper is based on an talk at the Third Augustus de Morgan Conference on Logic: Belief Revision and Contraction, held in King's College London, November 5-6 2001. The author would like to thank the participants for the discussion and others who were not present at the conference but whose comments helped in the preparation of the final version, especially Alexander Bochman and Daniel Lehmann.

The author apologises for the informal referencing of the paper. But with one or two of the most recent books in hand, e.g. those of Rott, Bochman, or Williams & Rott mentioned in the first section, the reader can easily locate in their bibliographies all publications mentioned here.

Draft, last revised: 26.11.01

Word count: 5,500