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

DAVID MAKINSON, *Department of Computer Science, King's College London, Strand, London WC2R 2LS, UK.*
E-mail: makinson@dcs.kcl.ac.uk

Abstract

The logic of belief change and nonmonotonic logic both broke with tradition in significant ways when they were developed in the 1980s. On a philosophical level they rose to the challenge that logic is not just about deduction. They also showed how one may go beyond classical logic without objecting to any of its principles, nor creating 'non-classical logics', by using the classical apparatus more imaginatively. On a heuristic level, they confirmed a growing trend to centre logic on inference relations rather than on a distinguished property of individual propositions. They also cut free from the reflex, when exploring a new logical notion, of first representing it as a connective of the object language. On a technical level, three novel features stand out: acceptance of an essential multiplicity in the object of study; suspension of the requirement of uniform substitutability for elementary letters; and use of non-Horn rules in defining families of operations.

Keywords: AGM, belief change, contraction, revision, nonmonotonic reasoning.

1 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. One can be rather suspicious of new subjects being created by translation from others. 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 modes of publication that is normal in such affairs. First, journal papers on one or the other of the two subjects; then books collecting together conference papers; special issues of journals; rather later overview chapters in handbooks, and individually authored books.

For belief revision, the conference collections began with a volume edited by André Fuhrmann and Michael Morreau on *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 in 2001 is *Frontiers in Belief Revision*, edited by Mary-Anne Williams and Hans Rott.

Special journal issues included *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 with Oxford University Press, following an overview of nonmonotonic reasoning in a *Handbook* of 1994.

The first individually authored book on belief contraction and revision was Peter Gärdenfors' *Knowledge in Flux* of 1988. In the 1990s, the focus tended to move to nonmonotonic logic, but in 1997, 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 to be written as a regular textbook for students, with exercises and answers.

In 2001 two further single-authored volumes appeared, which seem to be announcing a new way of approaching the area. Each book mentioned above dealt with belief change, and others that have not been mentioned dealt with nonmonotonic reasoning; the focus was squarely on one even when 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 Nonmonotonic Reasoning*. In this way, belief revision and nonmonotonic logic are coming to be seen as chapters of a general theory of belief management.

Evidently, this broad perspective is valuable, but I am not sure whether the two areas should be presented as instances of a common *formalism*. That may hide interesting differences of gestalt and intuition 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.

2 Back in 1985

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 papers by its three authors, writing in various combinations, and was also influenced by earlier work of others.

We were all influenced by — but at the same time attempting to escape from — the possible-worlds approach to nonclassical logics, and in particular the account of counterfactual conditionals given by David Lewis in his 1973 book *Counterfactuals*.

Peter Gärdenfors was also influenced by work of William Harper and Isaac Levi in epistemology and the philosophy of science. He wanted to show that we can understand counterfactual conditionals without committing ourselves ontologically to a realm of possible worlds. His strategy, initiated in a paper of 1979, was to use the Ramsey test to provide conditions for the acceptability of a counterfactual in terms of the minimal change of belief upon imagined acceptance of the antecedent. Ironically, it was also Peter who nearly a decade later in 1986, showed that the reduction was impossible — although as shown subsequently by Katsuno and Mendelzon, it can be carried out with an operation of update in place of revision.

On the other hand, Alchourrón and I came to belief change from a problem in the philosophy of law, of determining the result of abrogating an item from a legal code, and published our first paper on the subject in 1981. Here it was perfectly natural to proceed without bothering at all with possible worlds, using instead maximal non-implying subsets of a given belief set.

The encounter took place when Carlos and I submitted to the journal *Theoria* a paper on what is now known as maxichoice contraction, i.e. an operation selecting a maximal subset of the initial belief set that fails to imply the item to be contracted. *Theoria* was then edited by Peter. Collaboration developed into the paper AGM 1985, drafted through exchanges of longhand airmail between Buenos Aires, Lund and Beirut, plus the occasional telegram. Those were the days before email.

In this work, and subsequently, the different paths that had led us to belief change gave a difference of perspective. Peter thought primarily in terms of revision, while Alchourrón and I thought mainly in terms of contraction. But somehow these came together in the paper itself. Contraction was taken as basic, and revision defined from it by the Levi identity. On the syntactic level, the properties of revision were developed in detail, and it was observed that contraction could be recuperated from it via the Harper identity. But on the semantic level, the working construction was partial meet contraction, i.e. the intersection of a selected subfamily of the family of all maximal subsets of the initial belief set that fail to imply the item to be contracted. A semantics designed directly for revision had to wait until the work of Grove, a few years later.

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 paper 'A logic for default reasoning' was published in the same year.

As far as Peter and I can remember, when writing AGM 1985 neither of us was aware of those two important papers; nor as far as we know was Carlos Alchourrón (who died in 1996). Nevertheless, AGM came into existence following and alongside seminal work in nonmonotonic reasoning. And what I am going to say about it applies, to a large degree, to 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.

3 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 sustained formal study by logicians.

Indeed, in my view, it is not helpful to refer to all knowledge management as inference. Terminology is not very settled in this area, but a convenient minimal condition on an inference operation would seem to be inclusion $A \subseteq C(A)$; in other words the identity condition $a \mid\!\sim a$, if we are expressing inference as a relation between individual propositions.

Belief revision operations $K * a$ are not inference operations in this sense, for the trivial reason that with their two argument places they have the wrong arity. Nor are their left projections $*_a(K)$ inference operations, since we do not have the inclusion $K \subseteq *_a(K)$ except in the limiting case that a is consistent with K, in which case $*_a(K) = Cn(\{a\} \cup K)$. On the other hand, the right projections $*_K(a)$ on the input argument are inference operations on singletons, as noted by Gärdenfors and myself in 1991, since we have $a \in *_K(a)$. Belief contraction operations $K - a$ are even further removed from inference. They take information away rather than add it in, and neither of their two projections satisfy inclusion.

Another example of a non-inferential logical operation may be found in the logic of conditional goals or obligations. We may entertain a condition without it automatically becoming a goal. The passage from condition to goal, modulo an assumed code of conditional goals, is not an inference, for the rule of identity 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 (see the paper of the same name, *Journal of Philosophical Logic* 29 (2000) 383–408).

But that is another story. . . . The important point here is that the logics of both belief change and nonmonotonic reasoning opened the horizons of logic beyond the limits of deductive inference and, in the case of the former, beyond 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 should 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 — typically classical consequence itself. The same is true of most approaches to nonmonotonic logic, for example that of preferential inference. In a few 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.

4 Heuristic orientations

Look at relations, not sets

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 property* of propositions — that of being 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. I am not sure 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 premiss 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 dates back to the 1930s, but it is usually seen as a special-purpose or alternative presentation, rather than the leading one — a point of view which, personally, I find unfortunate.

Of course, in classical logic we can move freely between the distinguished set T of necessarily true propositions and the corresponding inference relation. In one direction, we put $a \vdash x$ iff $a \rightarrow x \in T$, where $a \rightarrow x$, i.e. $\neg a \vee x$, is the material conditional proposition. In the other direction, we can put $x \in T$ iff $p \vee \neg p \vdash x$. But when we get into nonmonotonic reasoning, it turns out that the reduction is no longer available. We do not have the equivalence $a {\mid\!\sim} x$ iff $p \vee \neg p {\mid\!\sim} a \rightarrow x$. The left implies the right for suitable nonmonotonic inference relations ${\mid\!\sim}$, but notoriously the right does not imply the left. Indeed, as is well known, such an equivalence would permit us to derive a form of monotony: $a \wedge b {\mid\!\sim} x$ whenever $a {\mid\!\sim} x$.

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.

At first sight, the difference between *relations* between formulae and *operations* on sets of formulae is essentially one of notational convenience — and it is so, provided the structures investigated are compact, so that whatever we want to say about a infinite set of premisses 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 for nonmonotonic inference, and this forces a procedural choice when we treat preferential inference from infinite sets of premisses. 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 bi-conditional. I am not sure whether one procedure is better than the other: the former appears to be more principled, the latter better behaved.

There is another question on which 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 for 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 internalize 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 a *relation* between propositions, or as an *operation* on propositions. The idea is: do not prematurely internalize 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 behind Quine's strictures against modal logic, doubting 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 before getting into more complex formulae. Once the flat case is well understood, 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.

Such was the strategy of AGM for belief change. But it must be granted that in that very case it left a big gap to be filled. For it is very natural to think of belief change as a process that in ordinary life we iterate routinely, 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 noniterative case. The reason seems to be pragmatic: we have different conceptions of what we want to hold constant and what we want to vary. In recent publications Isaac Levi, Bochman and Rott have all described contraction and revision as 'dynamic' in that beliefs are modified as far as needed to maintain consistency with the input. On the other hand, nonmonotonic inference is 'static', in the sense that the premiss entertained need not lead to any loss of background beliefs with which it is inconsistent. No background beliefs are *abandoned*, although some of them may be *left unused* in drawing a particular inference. When the inference is over we revert to the whole set of background beliefs, suspending application of different ones among them as needed when we consider another inference.

This is an example, I think, where over-attention to formal translations can lead us to ignore underlying nonformal differences. But even on a formal level, differences seem to arise once we get into iteration. An iterated revision is typically something like $(K * y) * x$, 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. The translation of $z \in (K * y) * x$ into the language of nonmonotonic inference would be $x \mid\!\sim_J z$ where $J = C_K(y) = \{u : y \mid\!\sim_K u\}$. On the other hand, the translation of $x \mid\!\sim (y \mid\!\sim z)$ into the language of belief revision would be $(y \mid\!\sim z) \in K * x$; and the translation of $(x \mid\!\sim y) \mid\!\sim z$ would be $z \in K * (x \mid\!\sim y)$. It is not clear how these relate to each other — nor what they mean individually. In sum, even on a purely formal level, 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 noniterative case.

Do some logic without logic

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 cheating, on the ground that we are in effect making the process of grouping premisses 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 premiss sets that are closed under classical consequence simulate the set of disjunctions of their respective elements. A well-known example is the rule of disjunction in the premisses for a supraclassical inference operation C, which may be written as $C(A) \cap C(B) \subseteq C(Cn(A) \cap Cn(B))$. 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 was the perspective taken 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 classic paper of 1990, thus taking the work much further. This is also what 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, compared with the 'connective-rich' Reiter default logic. The same spirit, of working in a connective-poor language first, underlies the treatment of entrenchment in the logic of belief change given by Hans Rott in his recent book *Change, Choice and Inference*.

5 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 logic 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 we are reasoning about it in the metalanguage. Even intuitionists and relevantists, who do not accept all of classical logic, feel the same way about their own systems; they have some difficulties, however, in reconciling this 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 similar postulates), or equivalently, all those that can be constructed as partial meet contractions (or by similar means) where the selection function or the minimizing relation is allowed to vary freely so long as it satisfies certain abstract conditions. This intrinsic non-uniqueness was, I think, a rather new feature of AGM belief revision, and perhaps contributed to the initial slowness of assimilation.

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 when 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 boundaries 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. Even when the preference relation is assumed to be linear, so that there is no more than one minimal state satisfying any given set of premisses, 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 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 nonmonotonic inference, are closed under substitution. More precisely, these operations are typically not closed under uniform substitution of arbitrary formulae for elementary ones. In this respect, our operations are quite unlike classical consequence, or of any of the usual subclassical logics — intuitionistic, relevantist, many-valued, etc.

When $\mid\!\sim$ is a nonmonotonic inference relation, $a \mid\!\sim x$ does not in general imply $\sigma(a) \mid\!\sim \sigma(x)$ for substitution functions σ. For example, suppose that $\mid\!\sim$ is the preferential inference relation determined by a linear model $\{1,2,3\}$ with the natural order, and $v(p,i) = 1$ iff $i = 1$ while $v(q,i) = 1$ iff $i \in \{1,2\}$.

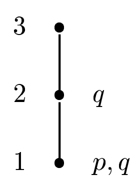

Then $p \mid\!\sim q$ since q is true in the least world 1 satisfying p but, $\sigma(p) \not\mid\!\sim \sigma(q)$ where σ is the substitution that puts, say, $\sigma(p) = \neg p$ and $\sigma(q) = \neg q$ since q is still true in the least world failing p, namely world 2.

Likewise for belief contraction and revision: $x \in K - a$ does not imply $\sigma(x) \in K - \sigma(a)$ or anything like that.

Indeed, as known in the folklore since at least the 1930s, the same failure already arises for all supraclassical closure operations, except for classical consequence itself and the total operation under which everything is a consequence of everything. For if $Cn < C$ then there are A, x with $x \in C(A)$ but $x \notin Cn(A)$. Hence there is a classical valuation v with $v(A) = 1$, $v(x) = 0$. Substituting tautologies for elementary letters that are true under v, and contradictions for those that are false under v, we have a set A' of tautologies and a contradiction x'. If C is closed under substitution, $x \in C(A)$ gives $x' \in C(A')$. But $A' \subseteq Cn(B)$ for arbitrary B, and $y \in Cn(x')$ for every y. Putting these together and using supraclassicality, monotony and idempotence of C, we have $y \in C(B)$.

Thus none of these operations are *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 honouring 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.

There is a way in which we can recover substitution and create a unique object of attention. If in the end we do internalize 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, with a set of distinguished formulae that is closed under substitution. That is one of the seductions of internalization.

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. It will be a set of formulae, the 'theorems' of the logic. Moreover, if we try to extract an appropriate relation R by defining it to hold between formulae a and x iff the formula $a \mid\!\sim x$ is a theorem of that logic, then R, restricted to Boolean formulae, 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 arbitrary relation between sets A of propositions and individual propositions x (both Boolean), 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 is not R^+, nor $\sigma(R^+)$ but $(\sigma(R))^+$. If we wrote the first we would be saying that R^+, which is a preferential inference relation, is itself closed under substitution — which in general is not the case. If we wrote the second, we would be saying only a triviality that holds by the definition of substitution on any relation over formulae.

Thus, on the level of individual preference relations, monotony, compactness and substitution all fail, but the function that takes an arbitrary inference relation to the least preferential relation 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 $(K-a)\cap(K-b) \subseteq K-(a\wedge b)$. This is a Horn condition, as it says: whenever $x \in (K-a)$ and $x \in (K-b)$ then $x \in K-(a\wedge b)$. But the other AGM supplementary postulate (K-8) tells us that if $a \notin K-(a\wedge b)$ then $K-(a\wedge b) \subseteq (K-a)$. In other words, whenever $x \in K-(a\wedge b)$ and $a \notin K-(a\wedge b)$ then $x \in (K-a)$. This is not a Horn condition, as it has a negative premiss. Equivalently, of course, it can be expressed with a disjunctive conclusion: whenever $x \in K-(a\wedge b)$ then either $x \in (K-a)$ or $a \in K-(a\wedge b)$.

The same sort of thing arises notoriously for rational monotony, i.e. the rule that says: whenever $a|\!\sim y$ and $a \not\mid\!\sim \neg x$ then $a \wedge x|\!\sim y$ — 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 Michael Freund, Daniel Lehmann, Ramón Pino Pérez, Hassan Bezzazi and myself.

Such quasi-closure conditions, which resemble Horn conditions except that they have a negative premiss (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.

There is a certain similarity of spirit with the 'rules of rejection' that were introduced by Jan Łukasiewicz in 1951 in his book *Aristotle's Syllogistic from the Standpoint of Modern Formal Logic*, which he and others have also applied to characterize classical propositional logic, intuitionistic logic, and other systems. But Łukasiewicz' rules of rejection do not appear to be quite the same, whether in purpose or in form. They are used to characterize a unique set or relation (e.g. the set of all classical tautologies) by simultaneous induction on that set and its complement, whereas we are looking at a family of relations with a family of corresponding complements. Moreover, rules of rejection are typically of the form 'If X is in and Y is out then Z is out', whereas rational monotony, say, is of the form 'If X is in and Y is out then Z is in'. The former is merely a contrapositive of a Horn rule, but the latter is not.

Non-Horn 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 it is better to stop here.

Acknowledgements and apologies

This paper is based on a talk at the Third Augustus de Morgan Conference on *Logic: Belief Revision and Contraction*, held at King's College London, 5–6 November 2001, and at the Workshop on Nonmonotonic Reasoning of the KR'02 conference held in Toulouse, 19–21 April 2002. The author would like to thank Alexander Bochman, Peter Gärdenfors, Daniel Lehmann, and Hans Rott for their penetrating and helpful comments on a draft. Allard Tamminga drew attention to the possible connection between non-Horn rules in nonmonotonic logic and Łukasiewicz rules of rejection.

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

Received January 2002