L&PS – Logic & Philosophy of Science
Vol. III, No. 1, 2005

A note on *Bridges from Classical to Nonmonotonic Logic*

Hykel Hosni
School of Mathematics, The University of Manchester
e-mail: hykel@maths.man.ac.uk

D. Makinson, *Bridges from Classical to Nonmonotonic Logic*,
King's College Text in Computing, volume 5, 2005
ISBN 1-904987-00-1. Published February 2005. Paperback
Available from Amazon, 10, USD 19.50 (multiple copies also direct from King's College Publications).
`http://www.dcs.kcl.ac.uk/staff/makinson`

First and foremost, the study of nonmonotonic logics aims at constructing logical systems capable of modelling the way "intelligent" agents, human or otherwise, draw reasonable conclusions from information which falls short of being definitive. The key idea pursued here is to allow agents to derive "more consequences" than granted by the classical (Tarskian) consequence relation, namely by permitting them to arrive at defeasible conclusions. These are consequences which despite being supported by the information currently available to the agent, could nonetheless be rejected in the light of new, or more refined, information. In a fundamental sense then, nonmonotonic logics occupy and undoubtedly prominent position among the disciplines investigating "intelligent reasoning" about "complex" and dynamic situations.

Thus it can hardly be surprising, as Gabbay, Woods and co-authors remarked on several occasions, that the formal study of nonmonotonic reasoning, starting with Aristotle, is at least as old as its monotonic counterpart (see e.g. Gabbay et al., 2002, p.4). It is certainly beyond question, however, that the explosion of interest in this subject arose as a more or less straightforward consequence of the development of logical techniques in artificial intelligence. It is no historic coincidence that some of the first influential studies in nonmonotonicity, among which we must mention at least Reiter's default logic and McCarthy's circumscription, were published on a special issue of the Artificial Intelligence Journal (Bobrow, 1980). Indeed, apart from a few exceptions, like the legal reasoning community, the initial motivation of the present interest in nonmonotonicity came essentially from computer science, and especially from the database community. To some minds this is not a noble origin. At least not as noble as the root that other branches of mathematical logic happen to have, say, in Frege's work. To such minds, nonmonotonic logics may be worth no more than the cold shoulder or at least the "justified" scepticism arising from their mysterious nature (and origins). In the words of its author, *Bridges from Classical to Nonmonotonic Logic* aims at taking this mystery out of the subject.

In the recent history of the discipline, of which the author of *Bridges* is a major actor, the second *annus mirabilis* is 1985 (after, that is, the publication of the above mentioned special volume in 1980). While Gabbay (1985) puts down the basis for a Genzen-style characterization of nonmonotonic inference, Alchourrón et al. (1985) present the soon-to-become logical paradigm of belief revision. Towards the end of the decade, Makinson introduces a very influential model theoretic framework relating to Gabbay's proof-theoretic intuitions and paving the way to the third *annus mirabilis*, 1990, when (Kraus et al., 1990) came to be published. There the soundness and completeness of the Gabbay-Makinson inferential system is proved with respect to the so-called preferential semantics. Subsequent investigations (culminating in (Friedman and Halpern, 2001)) will extend this soundness and completeness result to a variety of semantics, leading thus to the identification of the (refinement of the) Gabbay-Makinson one as the 'core system' of nonmonotonic reasoning. As the book under review helps us to realise in full, the subject of nonmonotonic logics is today in its logico-philosophical and logico-mathematical maturity.

Bridges, is a text written "out of the blackboard". Therefore the read, which is lively and stimulating throughout, is highly structured. Each section ends with a number of exercises which help see the fundamental concepts put to work. More challenging exercises are proposed as "problems" . Some fully detailed solutions to both exercises and problems are included. Interestingly, the author suggests also some "projects", which often encourage the reader to undertake some individual research. To complete the structure, each chapter is equipped with a "Review and Explore" section which includes a "Recapitulation" of the main points made, a "Checklist of Concepts and Definitions for Revision" and a list of "Further Readings" which often point to alternative accounts of the material presented in the corresponding chapters. A glossary of the special symbols, most of which will be unfamiliar to the reader who explores nonmonotonic logics for the first time, is included as an appendix.

Despite the topics discussed in chapters 5-6, the focus of the book is somehow narrow. It amounts, in fact, to the development of the logico-mathematical apparatus required to construct nonmonotonic inference operations out of the classical, monotonic one. As a consequence, one might think of it a dish better served *after* some appropriate appetizers. Indeed, scholars who never felt the urge to consider more generous notions of consequence than the one offered by classical logic might fail to appreciate, beyond mathematical logic, the import of the topic developed in the book. Still as the author himself remarks, the only necessary prerequisite is a firm, mathematical, command of classical propositional logic.

Chapter 1 - *Introduction* - sets out the stage. Its key message is that nonmonotonic logics, as in fact many other so-called non-standard logics, are not to be taken as alternatives to the classical one, in the way, that is, in which intuitionistic logic is regarded as opposed to classical logic (13). Both at the object- and at the meta- level the importance of classical logic is beyond dispute in nonmonotonic logics. Engaging in nonmonotonic logics means aiming at extending classical logic, rather than replacing it *tout court*. Similar remarks, of course, apply e.g. to probabilistic logics and many-valued logics. It is the main purpose of the book under review to put forward the mathematical details of how such an extension programme can be carried out in the case of reasoning with qualitatively incomplete information.

Perhaps a stronger emphasis on the fundamental interaction between classical and nonmonotonic logics could have been helpful to the reader. From the point of view - essential to the motivations of nonmonotonic logics - according to which logical systems are directed towards capturing (some aspects of) the way intelligent agents reason, classical inference is to be seen as formalising a very special case of nonmonotonic inference. Namely the situation in which every conclusion drawn by an agent is being performed under complete information. If this condition holds then the unconstrained form of monotonicity satisfied by classical logic is entirely justified. It is widely agreed that Mathematics constitutes the privileged example of a context in which unconstrained monotonicity is fully acceptable, and so uses monotonic logic. After all, we want to be able to utilise our theorems to derive other theorems. And this concatenation, once the corresponding axiomatic base is fixed (and consistent!), must never lead us to derive false conclusions from true premisses. Note that there is an important sense in which the monotonicity of classical logic is a necessary condition for the development of mathematics itself. What advance could there be in the subject if *everything* had to be proved from scratch *every time*? Hence for us limited agents, the conservativeness of mathematical reasoning can be seen as a condition of its development. This is why the project of *extending* logical reasoning to those domains (outside Mathematics) in which monotonicity cannot be accepted does not assume the form of a complete rejection of monotonicity altogether. Instead, as the author shows in the subsequent three chapters, the programme moves on by putting monotonicity under suitable constraints. In this sense, nonmonotonic logics are, by design, as conservative as they can possibly be.

This fundamental interaction takes place at the object-level. As to the meta-logic, there is no surprise whatsoever: it is classical logic. Therefore nonmonotonic logics are deductive. Regrettably Makinson refers here and there to classical monotonic logic as “deductive logic”, meaning of course that given the “truth” of the premisses no conclusion can be "false". This terminology, however, can be unfortunate to the extent that one might fail to resist the temptation of opposing nonmonotonic logics to deductive, "valid" logic. As *Bridges* helps us to realise in full, nonmonotonic logics are as conservative and truth-preserving as they can be.

The property of extending classical inference is known in the trade as *supraclassicality*. An important, and very often overlooked, observation is that supraclassicality does not come for free. Indeed the price paid for extending classical inference to include defeasible conclusions is summed up in Theorem 1.1 which states that supraclassical closure relations fail, except for two limiting cases, to be closed under uniform substitution. The upshot of this being that nonmonotonic consequence relations which are built 'on top' of such closure relations will, in the most interesting cases, inherit this failure. To require closure under substitution is therefore "a habit to suspend" if we want to understand nonmonotonic logics (16).

Chapters 2–4 constitute the core of the book in which the author presents three canonical constructions to obtain nonmonotonic consequence relations out of monotonic ones. The author follows a general pattern for introducing such constructions. The starting point is always the mathematical machinery provided by classical consequence. It is fundamental to notice that the nonmonotonic consequence relations discussed in the book all rest on the very same language as classical propositional logic. The second step consists in introducing a *bridge consequence relation*. Such a (monotonic) bridge is then used to construct a *non*monotonic consequence relation. Each chapter closes with a discussion on the "Specializations, Variants and Generalizations" of the main construction presented.

The main idea of Chapter 2 - *Using additional background assumptions* - is simple and powerful: distinguishing among local and background assumptions. It is reasonable to believe that when making inferences we, intelligent reasoners, do not consider all the information available to us, that is the premisses of our logical argument, as "equally conspicuous" (23). In particular it seems appropriate to distinguish between the information embodying 'the current premisses' of our argument - the set of local assumptions - from the set of assumptions which we commit to somehow tacitly - the set of background assumptions (or expectations). As discussed in (Datteri et al., 2005), this distinction is extremely pertinent when modelling computational learning systems. The bridge consequence relation discussed in this chapter, the *pivotal assumption consequence*, formalises this intuition. More precisely given a set of local assumptions the pivotal assumption consequence is defined as the usual classical consequence *modulo* a set of background assumption. This gives rise to supraclassical, yet monotonic consequence relations which are completely represented by a small set of syntactic properties. Note that each set of background assumptions defines a distinct consequence relation or, to put it in the other way round, local and background assumptions do not interact in any dynamic way. But of course, in the "real world" these kinds of assumption *do* interact dynamically: we very often *change* our expectations in the light of the "local information" we happen to possess. (See, again (Datteri et al., 2005) for some examples of this from artificial intelligence). Allowing the set of background assumptions to vary - essentially in a consistency preserving way - with the set of local assumption amounts to enabling nonmonotonic reasoning, as captured by the *default assumption* consequence relation, fully investigated in Section 2.2.

Default assumption consequences behave well as far as the standard properties of nonmonotonic reasoning are concerned. In particular they satisfy Cautious Monotonicity, the key property requiring that any system of nonmonotonic inference should be "as monotonic as possible" thus minimising the unnecessary revisions of accepted conclusions. However, when considering default assumption consequence we face, in the author's terminology, a basic Dilemma (34). Makinson proves, in fact, that if the set of background assumptions is *not* closed under classical consequence, then it is *not* language invariant. The unwelcome upshot of this being that logically equivalent sets of background assumptions might give rise to distinct consequence relations. On the other hand, if the set of background assumptions *does* happen to be closed under classical consequence, than the corresponding default assumption consequence collapses into classical consequence. The unwelcome upshot in this case being that we made no progress whatsoever into modelling intelligent reasoning beyond what can be captured by classical logic. Makinson suggests dealing with the first unwelcome consequence by putting constrains on the syntax which help reducing the negative effects of the phenomenon without losing expressive power. One wonders, however, whether this is indeed plausible when it comes to modelling actual agents' reasoning. As to the latter unwelcome consequence, a number of solutions are discussed. Indeed, default assumption consequence, which relates closely to historically significant ideas, like the Closed World Assumption (Reiter, 1978) and Poole's abductive framework (Poole, 1988), admit of a number of variations on the main theme, which Makinson categorises as *partial meet operations*, *submaximal operations* and *intersection-free operations*. All these variations allow the set of background assumptions to be closed under classical consequence yet avoiding the collapse of default assumption consequence into the classical one. Section 2.3 (and the corresponding exercises and problems) illustrates this at length.

Chapter 3 - *Restricting the Set of Valuations* - has a prominently model-theoretic flavour. The main intuition, of which the author is an early contributor, consists in weakening the Tarskian notion of logical consequence by allowing the agents to choose some subset of the set of all possible interpretations on which to define consequence. It is a transition from

$$\text{“}\phi \text{ follows from } \theta \text{ iff } \phi \text{ is satisfied in every model of } \theta\text{”} \tag{1}$$

to

$$\text{“}\phi \text{ follows from } \theta \text{ iff } \phi \text{ is satisfied in every preferred model of } \theta.\text{”} \qquad (2)$$

The transition from the classical, monotonic consequence defined in (1) to the *preferential*, nonmonotonic consequence captured by (2) takes place through the bridge consequence which Makinson calls *pivotal valuations*. The pattern is essentially the same as the one followed by the pivotal assumptions consequence except for the fact that this time the bridge consequence is classical consequence modulo a set of distinguished valuations on the language.

Pivotal valuations consequence turn out to fail compactness. It is precisely the property of compactness that distinguishes pivotal valuations from Pivotal assumptions consequence operations. As Makinson is able to prove in Theorem 3.4, pivotal assumption consequence operations are exactly the compact pivotal valuations ones. It is interesting to note that this failure is widespread in logics for uncertain reasoning, an important case in point being Maximum Entropy reasoning as characterized by the Paris-Vencovská probability logic (Paris and Vencovská, 1997).

In the finite case, pivotal valuations consequence admits of an elegant representation theorem pointing to the properties that suffice to characterise it completely. Then, as shown in full detail in section 3.2, the bridge from the notion of consequence captured by (1) to the one described by (2) above is completed by modifying pivotal valuation consequences so as to the allow the set of background valuations to vary with the set of local assumptions. This results in the nonmonotonic inference operations which Makinson calls *default valuations*. As they correspond (in finite, or more generally *smooth*, models) to the preferential consequence relations of (Kraus et al., 1990; Makinson, 1994), they turn out to occupy a spot of fundamental importance in the general field of nonmonotonic logics. Thus, default valuations consequence relates with a great variety of nonmonotonic logic systems discussed in the literature. For those already discussed in chapter 2, Makinson provides punctual connections. The others, from circumscription (McCarthy, 1980) to the more abstract algebraic approaches based on ideals are reviewed in Section 3.3.

Chapter 4 - *Using Additional Rules* - brings back the discussion from the model-theoretic level to the syntactic, indeed structural, one. The bridge therein presented, *Pivotal rules* consequence, is analogous to the pivotal assumption one except for the fact that instead of distinguishing a set of background assumptions (propositions), we distinguish a set of pivotal *rules*. Intuitively then, Pivotal rule consequence relations are defined in the following way: a sentence x is a Pivotal rule consequence of a set of local assumptions A modulo a set of rules R (amounting to a binary relation over the set of sentences) if and only if x is contained in every superset of A which is closed under classical consequence and under the relation R. The class of Pivotal rule consequences include that of pivotal assumption ones. In fact (Theorem 4.2) the pivotal assumption operations are exactly those Pivotal rule ones which satisfy the familiar property of *disjunction introduction in the premisses*. The way to go nonmonotonic from this monotonic bridge consequence relation, as the reader might well have guessed by now, consists in allowing the set of rules R to vary with A. Again, this dynamism must be subject to constraints, which - being the account syntactical - turn out to be consistency-based. The *Pivotal Rule* consequence relation is thus reached. In practice, this gives rise to the familiar notion of *extension*, pioneered in (Reiter, 1980). From a metamathematical point of view, however, this way of reaching the notion of an extension of a set of default rules is a great deal more satisfactory. Instead of the quasi-inductive, fixpoint construction familiar to those accustomed with Reiter's seminal paper, the definition of Pivotal rule consequence only resorts to purely inductive methods. The details and the consequences of this approach are discussed at length in section 4.2. Given the strong connection with Reiter's approach, the generalisations of Pivotal rule consequence mirror the variations on "default logic". An interesting discussion of complexity related issues (104-105) compares again Reiter fixpoint construction to the inductive constructions of Pivotal rule consequence. Given that much of the motivation for the study of nonmonotonic logics in general arises within computational domains, it would be of extreme interest to see how the distinct ways of arriving at (possibly distinct) families of nonmonotonic consequence relations discussed in the previous chapters might affect the complexity of the corresponding systems.

Chapter 5 - *Connections between Nonmonotonic and Probabilistic Inference* - tackles the issue in a relatively unusual way. Specifically, the author intends to apply the same strategy used so far to address the following questions: (i) how does probability relate to classical consequence; (ii) how can we go about constructing probabilistic supraclassical consequence operations and how they would compare to the qualitative ones; (iii) whether there are any interesting probabilistic *bridges* to be investigated.

According to de Finetti, probability is "the logic of the uncertain". By this, he meant that any individual facing uncertainty should reason according to the calculus of probability, or face the consequences of her "irrationality". So, by definettian analogy, one could answer to (i) by saying that classical consequence captures "the logic of the certain". That is, classical reasoning is a limiting case of the more general probabilistic one. This view, the reviewer believes, is only partly shared by Makinson. Surely, there are many mathematical points of divergence between the two notions of inference, chief among them, the lack of truth-functionality of the probabilistic one. The main aspect on which the author focusses for his comparison, however, is the relation between the notion of conditional probability on the one hand, and on the other hand the probability that an agent assigns to a conditional. This choice is very interesting with respect to the fact that conditional probability corresponds to a "limited kind of revision" (116), and hence relates prominently to the study of nonmonotonicity (more about this below). The author insists on there being a fundamental difference between conditional probability and the probability of a conditional, or in short $CP \neq PC$, thus exhorting us not to undertake the project of finding a connective adequate to represent conditionalization itself. In particular, he argues against the viability of the path suggested by de Finetti himself, namely reconciling the two notions of inference by means of a three-valued logic. De Finetti never developed such a logical system, in which truth values would be True, False and Null (*Nullo*, in Italian). Makinson tries to dissuade us from embarking on the enterprise by saying that

> to give up the attractiveness of two-valued logic [for the classical connectives] in favour of a weaker three-valued logic is an unpalatable price to pay for a dubious gain (120).

It is the reviewer's belief that qualifications should be made about this point. De Finetti's reasons for introducing such a three-valued logic cannot be separated from his general view on probability: the subjectivistic one. Probability, for de Finetti, corresponding to an agent's disposition to bet on a certain fact taking place (or not) can only be defined on *events*, that is to say well-determined facts of which we can fix in advance the circumstances under which it will obtain *or not*. In the betting interpretation, this latter situation results in the bets being called off:

> Indeed, if we are to gamble on a conditional bet, we must specify exactly under which circumstances the bet would be called off. When writing $P(E \mid H)$, in fact, by H we mean an *event* and not a *fact*. There could be very many facts that correspond to a single event. (de Finetti, 1995)

This important restriction on the domain of "meaningful" probability functions clearly supports a the idea of there being a strong correspondence between the probability of an event taking place under certain circumstances, and the corresponding 'conditional'.

Further elements of comparison between classical and probabilistic inference are provided in Section 5.2 where the author illustrates several probabilistic characterizations of classical consequence, nicely summarized on page 124. Point (ii), the investigation of supraclassical probabilistic inference, is addressed in Section 5.3. There the failure of monotonicity of the latter is somehow relegated to the background by its failing the very familiar property of *conjunction in the conclusions* or AND. Kyburg's *lottery paradox* as well as Makinson's own *preface paradox* (which can be seen as a qualitative version of the former) are adduced as reasonable occurrences of such a failure. Hence, in addressing point (iii) above, discussed in the subsequent Section 5.4, the author tries to identify intermediate consequence relations that whilst nonmonotonic, satisfy AND. The main two approaches discussed are the *limiting probabilities* and the *possibility functions* ones. The former occupy a prominent position in the study of nonmonotonicity. In its most mathematically elegant formulation, introduced in (Lehmann and Magidor, 1992) this approach formalises expressions of the type "if θ then typically ϕ", as "the probability of ϕ *given* θ is *very high*", where "very high" is equated to $1 - \varepsilon$, for infinitesimal ε. The resulting ε-semantics yields soundness and completeness results for the so-called *rational consequence relations*, an extension of the preferential inference recalled above. In particular then, the corresponding consequence relation satisfies AND. The picture emerging from these considerations is one in which AND can be seen as a demarcation property between quantitative (for which it fails) and qualitative (for which it is satisfied) inference *independently* of the fact that the semantics of the consequence relations is given in quantitative or qualitative terms. Failure to appreciate this point generated countless debates of dubious profitability.

The connection between probabilistic representations of uncertainty and nonmonotonic reasoning has attracted numerous scholars over the past few years, producing very interesting results which connect solidly nonmonotonic logics with the Maximum Entropy approach (see, e.g. Kern-Isberner, 2001; Kern-Isberner and Lukasiewicz, 2004; Hill and Paris, 2002). What is noticeable about this connection is that Maximum Entropy reasoning is axiomatized by a number of so-called commonsense principles and many of those bear a very close resemblance to the Gabbay-Makinson rules/conditions for nonmonotonic inference. Although a detailed discussion of this connection would in many ways exceed the intended focus of the book under review, it would have surely contributed to make the general picture more compelling.

Finally, Chapter 6 - *Some Brief Comparisons* - gives a concise idea of the "many faces of monotonicity" to paraphrase one of the author's felicitous expressions, with special attention to its "minimality". Makinson - one of the pioneers of the logical approach to *theory change*, or as it is more commonly known, belief revision - guides us across the tight parallel between the AGM and KM postulates for revision and update on the one hand, and nonmonotonic inference operations (141–147) on the other. Moreover an account of the connection between those and the logical formalizations of conditionals and counterfactuals is provided in Section 6.2. The chapter, and with it the main body of the book, appropriately ends up by recalling the fundamental representation, completeness and soundness results for nonmonotonic consequence relations.

On a very general level, this volume could be read as an exhortation to practicing logical pluralism. There is no question about the fact that modelling the patterns of reasoning followed by "intelligent" agents is an extraordinarily complex task, and that only wild illusion can feed the belief that a single logical system will ever suffice to accomplish it. To make a long story short: if one has the disposition to believe that intelligent ways of reasoning can ever be captured logically, then one must be ready to admit, from the very outset, a plurality of logics, each one doing its own part in the achievement of the grand goal. Makinson provides many examples of this attitude. (Incidentally, we note that the author could have emphasised this intrinsic pluralism by speaking of nonmonotonic *logics* right in the title). At the macroscopic level, he presents three general patterns to construct nonmonotonic consequence relations out of the classical, monotonic one. Magnifying, however, more details can be found. Each of these general patterns admit of several variations and interpretations which clearly contribute towards multiplying the final outcomes. No wonder that this must be a dazing scenario for logical monists. Monists of an ecumenical sort, however, could be happy to admit the mathematical interest behind such a plurality if *the right logic* is pointed out to them. Although Makinson dismisses the reasonableness of pursuing this line from the very beginning (13–14), we find that an exhaustive reply to such an attitude is provided along with the discussion of the various ways of generating nonmonotonic consequence relations by using additional assumptions:

> We would suggest that none of them is *the* right one, and none of them is *always* the best to use. From a theoretical perspective, they are all interesting ways of generating nonmonotonic inference operations. [. . .] From a practical perspective, they should be regarded as elements of a toolbox to be applied as suitable and convenient. (54)

Whenever transparent connections, clear transitions and general patterns can be pointed out precisely, plurality equals richness. The book under review, taken in conjunction with the earlier (Makinson, 1994), provides us with all that. For this very reason it is somehow regrettable that the author did not include a comprehensive discussion of the *plausibility measures* framework (Friedman and Halpern, 2001; Halpern, 2003), which is just mentioned in passing. This latter, encompassing several of the formalisms with the most appreciated currency in the uncertain reasoning trade, permits of very general representation results which contribute towards *explaining* why, for instance, so many model-theoretic approaches to nonmonotonic inference turn out to be equivalent in practice. That is, why the constraints on nonmonotonic inference introduced by Gabbay and Makinson, and subsequently refined by Kraus, Lehmann and Magidor result in logical systems which are sound and complete for a variety of intuitively distinct semantics. On a second edition of the volume, the addition of this topic (maybe among the representation results of section 6.3) would surely be extremely welcome.

To conclude, there is something noticeable about the way the book is being published. To begin with the end, the book is not available in bookshops. It can only be purchased on Amazon (or, for multiple copies, directly from King's College Publications). This potential drawback is well balanced, in our own view, by the resulting mutual benefits for readers and authors. The former take advantage of small production costs, which the paperback, print-on-demand strategy surely contributes to reducing, as they end up facing a very affordable volume (the book under review is between six and ten times cheaper than a volume from the Oxford Logic Guides). The latter, on the other hand, benefit from the fact that many intermediate steps between their LaTeXing and the final product are avoided. Therefore it will hardly be a surprise if this approach to scientific publication will result in a long lasing equilibrium between the readers' and the authors' payoffs.

References

C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50: 510–530, 1985.

D.G. Bobrow. Special issue on non-monotonic logics. *Artificial Intelligence*, 13 (1-2), 1980.

E. Datteri, H. Hosni, and G. Tamburrini. Machine learning from examples: A non-inductivist analysis. *Logic and Philosophy of Science*, 1(3), 2005. `http://www.univ.trieste.it/∼episteme/`.

B. de Finetti. *Filosofia della probabilità*. Il Saggiatore, Milano, 1995. English translation forthcoming.

N. Friedman and J. Halpern. Plausibility measures and default reasoning. *Journal of the ACM*, 48(4):648–685, 2001.

D.M. Gabbay. *Theoretical foundations for non-monotonic reasoning in expert systems*, pages 439–457. Proceedings NATO Advanced Institute on Logics and Models of Cuncurrent Systems. Springer, Berlin, 1985.

D.M. Gabbay, R.H. Johnson, H.J. Ohlbach, and J. Woods. *Handbook of the Logic of Argument and Inference*. North-Holland, 2002.

J. Halpern. *Reasoning About Uncertainty*. MIT Press, 2003.

L.C. Hill and J.B. Paris. When maximising entropy gives the rational closure. *Journal of Logic and Computation*, 12(1):1–19, 2002.

G. Kern-Isberner. *Conditionals in Nonmonotonic Reasoning and Belief Revision - Considering Conditionals as Agents*, volume 2087 of *Lecture Notes in Computer Science*. Springer, 2001.

G. Kern-Isberner and T. Lukasiewicz. Combining probabilistic logic programming with the power of maximum entropy. *Artificial Intelligence*, 157(1-2):139–202, 2004.

S. Kraus, D Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative inference. *Artificial Intelligence*, 44(1):167–207, 1990.

D. Lehmann and M Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55:1–60, 1992.

D. Makinson. General patterns in nonmonotonic reasoning. In D.M. Gabbay, C. Hogger, and J. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, pages 35–110, 1994.

J. McCarthy. Circumscription: a form of non-monotonic logic. *Artificial Intelligence*, (13):27–39, 1980.

J.B. Paris and A. Vencovská. In defence of the maximum entropy inference process. *International Journal of Approximate Reasoning*, 17:77–103, 1997.

D. Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36(1): 27–47, 1988.

R. Reiter. On closed world databases. In H. Gaillaire and J. Minker, editors, *Logic and Data Bases*, pages 55–76. Plenum Press, New York, 1978.

R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1–2):81–132, 1980.