DAVID MAKINSON. ***Bridges from classical to nonmonotonic logic.*** Text in Computing, vol. 5. King's College, London, 2005, xvi + 216 pp.

The study of non-monotonic reasoning aims at modeling the way "intelligent" agents, human or otherwise, draw reasonable conclusions from information which falls short of being definitive. From the logico-mathematical point of view, this translates into devising more generous notions of logical consequence than the "classical", Tarskian one. In particular agents reasoning non-monotonically should be in a position of arriving 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. This attention to the flexibility of reasoning under imperfect information puts non-monotonic logics in prominent position among the disciplines investigating "intelligent reasoning" about "complex" and dynamic situations.

Although it is sometimes remarked that the formal study of non-monotonic reasoning, starting with Aristotle, is just as old as its monotonic counterpart, there is little doubt that the explosion of interest in this subject arose as a consequence of the development of logical techniques in artificial intelligence. It is no historic coincidence that some very influential studies in non-monotonicity were published on a special issue (number 13, 1980) of the ***Artificial Intelligence Journal***. Indeed, apart from a few important exceptions—notably the legal reasoning community—the initial motivation fueling the interest in non-monotonicity came essentially from computer science—especially from the database community. To some minds this is not a noble origin. At least not as noble as the root that certain branches of mathematical logic might happen to have, say, in Frege's foundationalist work. To such minds, non-monotonic logics may be worth no more than the cold shoulder or at least some "justified" skepticism 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. Further than that, the book under review helps us realising in full how the subject of non-monotonic logics is today in its logico-mathematical and logico-philosophical maturity.

Bridges is a text derived from Makinson's courses and pays special attention to the reader's learning process. Therefore the exposition, which is lively and stimulating throughout, is highly structured. Each section ends with a number of exercises which show 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.

Despite the topics discussed in chapters 5–6, the focus of the book is somewhat narrow. It amounts, in fact, to the development of the logico-mathematical apparatus required to construct non-monotonic 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 the purely logico-mathematical interest, the more general import of the topic developed in the book. However, as the author himself remarks, the only necessary prerequisite is a firm, mathematical command of classical propositional logic.

Chapter 1—*Introduction*—sets the stage. Its key message is that non-monotonic logics (as many other so-called non-standard logics) are not to be taken as alternatives to the classical one—in the way, for instance, intuitionistic logic is regarded as opposed to classical logic (p. 13). Engaging in non-monotonic logics means aiming at extending classical logic, rather than replacing it *tout court*. The book under review puts 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 non-monotonic logics could have been helpful to the reader. From the point of view—essential to the motivations of non-monotonic logics—according to which logical systems are directed toward capturing (some aspects of) the way intelligent agents reason, classical inference is to be seen as formalising a very special case of non-monotonic inference, namely the situation in which agents reason under certain and complete information. Whenever this rather demanding condition holds, the unconstrained form of monotonicity satisfied by classical logic is entirely justified. Hence the project of *extending* logical reasoning to those domains 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. As ***Bridges*** helps us realising in full non-monotonic logics are, by design, as conservative and truth-preserving as they can possibly be. (It might be noted at this point that the subject surely owes some of its "mystery" to the misleading nature of the term *non-monotonic* which suggests rejecting monotonicity rather than constraining it.)

Such a fundamental interaction between classical and non-monotonic logics takes place at the object-level. As to the meta-logic there is no surprise whatsoever: it is classical logic. Therefore non-monotonic logics are deductive. Regrettably Makinson occasionally refers to classical monotonic logic as "deductive logic", meaning of course that given the "truth" of the premises no conclusion can be "false". This terminology, however, can be unfortunate to the extent that one might fail to resist the temptation of opposing non-monotonic logics to deductive, "valid" logic.

Chapters 2–4 constitute the mathematical core of the book in which the author presents three canonical constructions to obtain non-monotonic consequence relations out of monotonic ones. The author follows a general pattern to introduce such constructions. The starting point is always the mathematical machinery provided by classical consequence. Hence the non-monotonic consequence relations discussed in the book all rest on the very same language as classical propositional logic. The second step consists in introducing a (monotonic) *bridge consequence relation* which is used to reach a *non*-monotonic one. This pattern leads to three families of non-monotonic consequence relations that arise from considerations of syntactic (Chapter 2—*Using additional background assumptions*) semantic (Chapter 3: *Restricting the set of valuations*) and structural (Chapter 4: *Using additional rules*) nature. Each chapter closes with a discussion on the "Specializations, Variants and Generalizations" of the main construction presented.

In Chapter 5—*Connections between Nonmonotonic and Probabilistic Inference*—the author applies the above strategy to shed some light on the connections between probabilistic and non-monotonic reasoning. The main aspect on which the author focuses for his comparison is the relation between the notion of conditional probability on the one hand, and the probability that an agent assigns to a conditional on the other. This is a natural choice given that conditional probability can be seen to correspond to a "limited kind of revision" (p. 116) and therefore relates prominently to the study of non-monotonicity.

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 non-monotonic inference operations (pp. 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 non-monotonic consequence relations.

On a very general level, this volume could be read as an exhortation to practice logical pluralism. If one has a 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 toward 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 non-monotonic *logics* right in the title). At the macroscopic level, he presents three general patterns to construct non-monotonic consequence relations out of the classical, monotonic one. Magnifying, however, more details can be found. Each of these general patterns admits of several variations and interpretations which clearly contribute toward 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 acknowledge the mathematical interest lying behind such a plurality provided that *the right logic* is pointed out to them. Although Makinson dismisses the reasonableness of pursuing this line of thought from the very beginning (pp. 13–14), we find that an exhaustive reply to such an attitude is provided along with the discussion of the various ways of generating non-monotonic 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 non-monotonic inference operations. [...] From a practical perspective, they should be regarded as elements of a toolbox to be applied as suitable and convenient. (p. 54)

Whenever transparent connections, clear transitions and general patterns can be pointed out precisely, plurality equals richness. The book under review constitutes a very important step toward achieving this in the field of non-monotonic logics. For this very reason it is somewhat regrettable that the author did not include a comprehensive discussion of the *plausibility measures* framework 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 toward *explaining* why, for instance, so many distinct model-theoretic approaches to non-monotonic inference turn out to be equivalent in practice, that is, why the constraints on non-monotonic 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.

HYKEL HOSNI

School of Mathematics, The University of Manchester, UK. hykel@maths.man.ac.uk.

Entry for the Table of Contents: