Bridges from Classical to Nonmonotonic Logic

David Makinson

Department of Computer Science
King's College London

Contents

Preface

Who are we writing for?

This book is directed to all those who have heard of nonmonotonic reasoning and would like to get a better idea of what it is all about. What are its driving ideas? In what ways does it differ from classical logic? How does it relate to probability? More generally, how does it behave and how can it be used? We will try to answer these questions as clearly as possible, without undue technicality on the one hand, nor vagueness or hand waving on the other.

It is written for the student in a classroom, the instructor teaching a course, but also for the solitary reader. The lone reader faces a particularly difficult task, for there is nobody to turn to when the going gets tough. In sympathy, we try to help as far as possible: explaining all points fully, saying some things in more than one way to get the point across, providing recapitulations at the end of each chapter, giving exercises and selected solutions. These will be useful to all readers, but especially for those working alone.

What does the reader need to know already?

Tackling a subject such as this does require a minimal background, and in honesty we must make it clear from the beginning what is needed. To those without any grounding in classical logic at all, who may out of curiosity have picked up the volume, we say: put it down, and go to square one. Get the rudiments of classical propositional logic (alias truth-functional logic) under your belt, and then come back. Otherwise you will be like someone trying to learn about algebra before knowing any arithmetic, counterpoint without sol-fa, or style without the elements of grammar.

For that purpose, students of computer science, mathematics, philosophy or linguistics should take the introductory course in logic offered by their department, and read whatever textbook the instructor recommends. If working alone with an interest in computation or mathematics, try chapters 6-8 of James Hein *Discrete Structures, Logic, and Computability* (Boston: Jones and Bartlett, second edition 2002). Those whose background is in the humanities should try L.T.F. Gamut *Logic, Language and Meaning: Volume I. Introduction to Logic* (Chicago University Press, 1991). And for those lucky enough to find it, there is the author's out-of-print *Topics in Modern Logic* (London: Methuen 1973; chapters 1 and 3).

In the Introduction of this book we do review some features of classical propositional logic, but with a special purpose. We bring out certain aspects that tend to be passed over in most presentations, but are essential for unfolding our theme. In particular, we explain some very general concepts that reveal themselves not only in classical consequence but also in other contexts. Among these are the general notion of a closure operation or relation, the property of compactness, and the idea of a Horn rule.

The reader should also be familiar with a few simple mathematical tools. These are of two kinds. First: some 'working set theory' – basic operations on sets, such as finite intersection, union and relative complement, indexed families of sets, infinite intersection and union; ordered pairs, Cartesian products, the treatment of relations and functions as sets of ordered pairs, and the notion of a well-ordering. Second: mathematical induction as a method of proof, in both its simple and course-of-values forms over countable sets. This is needed in order to establish routine properties of the inductively defined sets that will be encountered, beginning with the set of formulae of propositional logic.

These tools are often taught along with elementary logic, and also as part of first courses in computing and discrete mathematics. For the reader who lacks them, there is no alternative to going back to an elementary course or to reading the first few chapters of an introductory book. A good old standby is Seymour Lipschutz *Set Theory and Related Topics* (New York: McGraw-Hill Education, second edition 1998), especially the first five chapters. Alternatively, chapters 1-4 of the book of Hein mentioned above. For those who would like to do it in elegance and style, the first twelve chapters of Paul Halmos' *Naïve Set Theory* (New York: van Nostrand, 1960) are unbeatable.

The proofs of some results in this book, for example Theorem 2.2–2, appeal to maximality principles such as Zorn's Lemma. Familiarity with these principles, as set out say in chapter 9 of Lipschutz or chapters 14-17 of Halmos, is needed for full understanding of those proofs, and is in any case a premium investment. However, even without them the student may get a partial understanding: read the theorem for the finite case only, and recall that for a finite set, if some subset of it has a property then it is always included in at least one maximal subset with the same property.

In the author's opinion, all these things should be taught in high school, and in some privileged places they are. They are fun to learn, and indispensable in the electronic world. But sadly, there are also students who are allowed to go through an entire undergraduate education without ever being exposed to them.

So, for those who do not have the necessary minimal background we say goodbye – and hope to see you again later. For those who are ready, we say: welcome and read on!

Main themes

From the outside, nonmonotonic logic is often seen as a rather mysterious affair. Even from the inside, it can appear to lack unity, with multiple systems proposed by as many authors going in different directions. The few available textbooks tend to perpetuate this impression.

Our main purpose is to take some of the mystery out of the subject and show that it is not as unfamiliar as may at first sight seem. In fact, it is easily accessible to anybody with the minimal background that we have described in classical propositional logic and basic mathematical tools - provided certain misunderstandings and a tenacious habit, both signalled in the first chapter, are put aside.

As we shall show, there are logics that act as natural bridges between classical consequence and the principal kinds of nonmonotonic logic to be found in the literature. These logics, which we call *paraclassical*, are very simple to define and easy to study. They provide three main ways of getting more out of your premises than is allowed by strict deduction, that is, by good old classical consequence. In other words, they are principled ways of creeping, crawling or jumping to conclusions. Like classical logic, they are perfectly monotonic, but they already display some of the distinctive features of the nonmonotonic systems that they prefigure, as well as providing easy conceptual passage to them. They give us, in effect, three main paths from the classical homeland to nonmonotonic shores.

The book examines the three one by one. We begin with the simplest among them, whose moving idea is to use *additional background assumptions* along with current premises. Then we consider ways of getting the same result by *excluding certain classical valuations*. And finally we examine a third means to the same end, by *adding rules* alongside the premises.

In each case we obtain a monotonic bridge system that may be rendered nonmonotonic in a natural and meaningful way. In fact, each of the bridge systems leads to quite a range of nonmonotonic ones, and this is where the diverse systems of the literature show themselves. But they are no longer a disorderly crowd. They fall into place as variants, natural and indeed to be expected, of three basic ideas.

The book then turns to the subtle question of the relation between logic and probability – more specifically, between classical logic, probabilistic inference, and nonmonotonic reasoning. On the one hand, there are several different ways of characterizing classical consequence in probabilistic terms. On the other hand, we can also use probability to define supraclassical consequence operations, both monotonic and nonmonotonic. They turn out to be quite different in certain important respects from the qualitative operations of earlier chapters; but as we show there are also ways of bringing them closer together.

There are obvious resemblances between some of the ways of generating nonmonotonic inference relations and constructions that have been used for other purposes – for example, the logics of belief revision and update, as well as counterfactual conditionals and conditional directives. The last chapter of the book discusses the links between these different kinds of ‘belief management’ and their residual differences. It also presents sample representation theorems for some of the principal systems of nonmonotonic logic. We do not attempt anything like a complete coverage of all the different representations that may be found in the literature; there are too many, and still growing. Two strategic examples are selected, one with a complete proof and the other with its central construction merely sketched.

Topics Not Dealt With

We owe it to the reader – and especially to the instructor who may be thinking of using the book – to say clearly what topics, associated with nonmonotonic reasoning, are *not* covered.

All inference relations studied are defined over a purely propositional language using only Boolean connectives. They could be extended to first-order classical logic, but most of the interesting questions appear to arise already at the propositional level. In this subject, to work in a first-order language is to increase overheads for little added value in understanding.

We will not be examining logics in which classical connectives are reinterpreted in a non-Boolean manner. Systems of logic programming with negation using various forms of ‘negation as failure’ are sometimes presented in this way although, as we will suggest in chapter 4, they are more transparently seen as rules about propositions in a fragment of the classical language. The study of logic programming has now become a world of its own, with an emphasis on questions of computational complexity, which are not at the centre of our concern. For a recent review of logic programming using the ‘answer set semantics’ see Gelfond and Leone (2002); for an overview of various approaches to logic programming with negation see Brewka, Dix and Konolige (1997) chapters 6-7; and for an encyclopaedic account see the book of Baral (2003).

Nor will we be examining languages that add non-classical connectives to the usual classical ones. In particular, we will not be considering autoepistemic logics. These are formed by adding to the Boolean connectives a special kind of modal operator whose ‘introspective’ reading engenders nonmonotonic behaviour. Despite their correspondence with certain maverick default-rule logics under translations due to Konolige and others, they are quite different from the mainstream ones, for reasons that we will explain in chapter 4. The reader interested in studying autoepistemic logic is referred to chapter 4.2 of Brewka, Dix and Konolige (1997) or, in more detail, the overview of Konolige (1994).

There are also two approaches to nonmonotonic reasoning in which logical connectives have little role to play. They are the theory of defeasible inheritance nets, and the abstract theory of argument defeat. In both cases the focus is on the notion of one path (in a net) or argument (in a discussion) defeating or undercutting another, and on ascertaining the final effect of complex patterns of iterated defeat. Here too, the issues at stake are rather special, and not discussed in this volume. For defeasible inheritance nets, the reader is referred to the overview of Horty (1994). For argument defeat, consult the recent survey of Prakken and Vreeswijk (2001) and the papers of García and Simari (to appear) and Bench-Capon (to appear).

Thus this book does not pretend to be anything near a comprehensive treatment of all the investigations that have gone under the name of nonmonotonic logic. But it does present some very central ideas, structuring them in a coherent perspective, and seeks to explain them as clearly as possible.

Strategy of Presentation

When the author first began planning this text, he followed engrained habits and sought for maximal generality. "We have available a multitude of different formal approaches to nonmonotonic reasoning", he thought. "So in order to put some order into the affair, we need to find the most general schema possible, under which they all fall. We can then present the different accounts as so many special cases of the general schema".

However, it quickly became clear that this strategy is hopeless. It does work quite well when focussing on the *kinds of consequence relation that are generated*, for these may be classified according to the sets of regularity conditions that they satisfy. One may thus begin with a minimal set of conditions, defining a very broad class of consequence relations, and gradually add to the conditions, defining narrower classes. Indeed, this was the strategy followed in the author’s paper ‘General theory of cumulative inference’ (1989) and survey ‘General Patterns in Nonmonotonic Reasoning’ (1994). But if, as in this book, one is primarily interested in the principal *modes of generation* of consequence relations, the strategy makes little sense. The modes differ abruptly from each other; and when one abstracts upon them in an attempt to obtain a ‘most general mode’, one comes up with formulations that are opaque to intuition, clumsy to work with, and mathematically close to empty.

For this reason, a totally different principle guides the organization of this book. It is the principle of *theme and variations*. For each of the three main ways of ‘going nonmonotonic’ – namely, by adding background assumptions, restricting the set of valuations, and adding background rules – we begin by providing a paradigm formulation, and then sketch some of the many specializations, variations, and generalizations. A few variations already appear at the level of the monotonic bridge systems, and they multiply at the nonmonotonic level. They are most easily understood as an open family of relatives of a central exemplar, rather than as particular cases of an all-embracing definition.

For the mathematician this may go rather against the grain, but in the present enterprise it is the only reasonable way to proceed. And after all, even mathematicians are accustomed to this procedure in certain cases. For example, when giving an overview of different approaches to axiomatizing set theory, nobody at the present state of play would try to give a most abstract definition of what in general constitutes a set theory, and then fit in the Zermelo-Fraenkel version, Quine's New Foundations, and so on as special cases. The standard procedure is to delineate a few broad lines of approach in conceptual terms, present exemplars of each, and finally sketch some of their many variations. That is the way in which we proceed in this book.

Review and explore sections

Each chapter ends with a 'Review and explore' section. It contains three packages to help the reader review the material covered and to go further into the literature. First, a *recapitulation* of the essential message of the chapter. Second, a *checklist* of the key concepts, both formal and informal, that were introduced there. The concepts are not explained all over again, but are simply listed for checking off at review time. The reader can relocate the formal definitions and informal explanations by looking for the corresponding italicised terms in the main text. Third, a short selection of texts for *further reading*. Some of these cover essentially the same ground as the section itself, albeit from a different angle, in more detail or with particular attention to aspects that we have not dwelt on. Some go further. In general, these will be 'bite size' – short papers or chapters. Occasionally, a more extended text is mentioned.

What is the best way to read this book?

Do it pencil in hand; scribble in the margins. Take nothing on faith; check out the assertions; find errors (and communicate them to the author); pose questions.

Notoriously, the best test of whether one has understood a definition is to be able to identify positive and negative examples. Nor has one understood a theorem very well if unable to apply it to straightforward cases and recognize some of its more immediate consequences. Without this one may have the illusion of understanding, but not the real thing.

For this reason, the book includes *exercises*. Many of them ask readers to verify claims made in the flow of the text. To help the brave loner maintain the discipline of working out the exercises, and also to take some of the load off instructors, selected exercises, marked by asterisks, are given answers at the end of the text.

As well as exercises, there are some *problems*. They are more demanding. In general they require more than checking positive and negative instances of a definition or straight applications with a couple of steps of argument. Their solutions may need perceptive guessing plus the ability to prove or disprove the guesses. Both the art of making good guesses and the ability to check them out are acquired skills, and grow with practice.

Finally, *projects* are longer-term tasks. Whether to attempt any of them depends on the interests and time of the reader, and the goals of the instructor. In general, they involve consulting items from the literature and working on them. When suggesting a project we will usually indicate a specific text as entry-point.

Those strapped for time and wanting to cover essentials only, could skim or even omit section 2.3 on specializations and generalizations of the basic pattern of that chapter, and likewise sections 3.3, 4.3 – provided they promise the author to come back some day later! They could also omit the rather technical section 5.4. Below that, they risk not getting the general picture.

For the instructor

The author's experience is that in a graduate class, teaching can take place at average of one section per hour, without counting section 1.1 and the brief 'review and explore' sections. This adds up to about 20 hours of instruction time for the entire book, or about 15 hours for a minimal version like the one mentioned above.

Of course, this will vary with the level and background of the students. In particular, those whose background in classical logic did not include much exposure to the notion of a consequence relation/operation or concepts such as compactness may need considerably more than one hour for the revision in section 1.2. With them in mind, there are particularly many exercises in that section. On the other hand, the sections entitled 'specializations and generalizations' provide instructors with a starting point for more detailed treatments of their favourite special topics, with open schedules.

Conventions

Theorems are numbered by their section, followed by a dash and a digit. For example, the first explicitly displayed theorem in the book occurs in section 1.3 and is numbered Theorem 1.3–1. However, only major results are displayed and numbered in this way. Many lesser facts are simply stated in the flow of the text, and the attentive reader should mark them as they are found. The same applies to definitions. Only a few key ones are given special display; the others are within the text, with their central terms italicized.

Generally speaking, the notation used is fairly standard. One exception: we use plain brackets rather than angular ones for ordered pairs – thus (x,y) instead of $<x,y>$. In chapter 4 we also simplify the usual way of writing default rules. Symbols special to nonmonotonic logic are recalled in a glossary at the end.

In this preface, we have given references in full. Likewise in the suggestions for further reading at the end of each section. But in the main text, they are called by author name and date, e.g. Gabbay (1985), and are answered with full details in the reference list at the end of the book.

Acknowledgements

As the basic ideas underlying this book were developed, they were articulated in workshop and conference presentations, journal publications, and the classroom.

The author would like to thank the organisers of the following meetings and occasions, large and small, for the opportunity to air ideas that went into this book. They include: Workshop on Cognitive Economics, Porquerolles, September 2001; KR-02 Special Session on Belief Change, Toulouse, 19-21 April 2002; ASL European Summer Meeting Special Session on Nonmonotonic Logic, Münster, 4-8 June 2002; Prague International Colloquium on Formal Epistemology, 10-13 September 2002; King's College London Department of Computer Science Colloquium, 9 October 2002; Manchester University Mathematics Department Colloquium, 18 November 2002; University of Leipzig Department of Computer Science, 14 May 2003; Dresden University Department of Computer Science, 19 May 2003; University of Buenos Aires Department of Computer Science, 17 July 2003; University of Bahía Blanca Department of Computer Science, 07 August 2003; Augustus de Morgan Conference on Agents and Knowledge Representation (King's College, 03 November 2003) ; University of New South Wales Department of Computer Science Colloquium, 02 April 2004.

The general vision of nonmonotonic logic developed here was outlined, in a form suitable for professional logicians, in Makinson (2003a). A thumbnail sketch, adapted for an audience of economists interested in logic, was given in Makinson (2003b). A stripped-down version of the present text, without such things as exercises, answers and 'review and explore' sections, will appear as the overview Makinson (to appear).

Much of the text was written while the author taught an MSc course on Nonmonotonic Reasoning in the Department of Computer Science, King's College London in the spring semesters of 2003 and 2004. Thanks to all the students and auditors for their questions, which helped make the presentation more transparent; particularly to Audun Stolpe and Robert Schubert for their perceptive queries, Xavier Parent and Simon Speed for finding bugs in exercises and providing some model answers; and David Gabelaia for comments and important observations.

Many other people made helpful comments on particular points during the gestation process. They include Gerd Brewka, Björn Bjurling, Marc Denecker, Michael Freund, Donald Gillies, Dov Gabbay, Jörg Hansen, Daniel Lehmann, Joao Marcos, Wlodek Rabinowicz, Hans Rott, Karl Schlechta, and an anonymous referee for the *Logic Journal of the IGPL*.

The book is dedicated to the memory of my former co-author Carlos Alchourrón, who died in January 1996. Despite his fascination with the logic of belief revision he distrusted its neighbour nonmonotonic logic, and I could never convince him otherwise. While writing this book, I realized that it is a continuation of our conversations.