THE BULLETIN OF SYMBOLIC LOGIC
Volume 12, Number 3, Sept. 2006

REVIEWS

The Association for Symbolic Logic publishes analytical reviews of selected books and articles in the field of symbolic logic. The reviews were published in ***The Journal of Symbolic Logic*** from the founding of the JOURNAL in 1936 until the end of 1999. The Association moved the reviews to this BULLETIN, beginning in 2000.

The Reviews Section is edited by Alasdair Urquhart (Managing Editor), Steve Awodey, John Baldwin, Lev Beklemishev, Mirna Džamonja, David Evans, Erich Grädel, Denis Hirschfeldt, Roger Maddux, Luke Ong, Grigori Mints, Volker Peckhaus, and Sławomir Solecki. Authors and publishers are requested to send, for review, copies of books to *ASL, Box 742, Vassar College, 124 Raymond Avenue, Poughkeepsie, NY 12604, USA.*

In a review, a reference "JSL XLIII 148," for example, refers either to the publication reviewed on page 148 of volume 43 of the JOURNAL, or to the review itself (which contains full bibliographical information for the reviewed publication). Analogously, a reference "BSL VII 376" refers to the review beginning on page 376 in volume 7 of this BULLETIN, or to the publication there reviewed. "JSL LV 347" refers to one of the reviews or one of the publications reviewed or listed on page 347 of volume 55 of the JOURNAL, with reliance on the context to show which one is meant. The reference "JSL LIII 318(3)" is to the third item on page 318 of volume 53 of the JOURNAL, that is, to van Heijenoort's *Frege and vagueness*, and "JSL LX 684(8)" refers to the eighth item on page 684 of volume 60 of the JOURNAL, that is, to Tarski's *Truth and proof*.

References such as 49***5*** or 280***1*** are to entries so numbered in *A bibliography of symbolic logic* (the JOURNAL, vol. 1, pp. 121–218).

David Hilbert's lectures on the foundations of geometry 1891–1902. edited by Michael Hallett and Ulrich Majer, David Hilbert's Lectures on the Foundations of Mathematics and Physics, 1891–1933, vol. 1. Springer, Berlin, Heidelberg and New York, 2004, xviii + 661 pp.

Hilbert's book ***Grundlagen der Geometrie*** of 1899 has played a peculiar role in the development of the foundations of mathematics. It has been hailed as the dawn of a new era, as in Weyl's 1944 paper *David Hilbert's mathematical work*. However, Hilbert's topic, the study of axiomatic geometry in the synthetic tradition of Euclid, has been an absolutely marginal field in 20th century mathematics. The ***Ω-Bibliography of Mathematical Logic*** listed just a few dozen papers up to 1986, many of which are of no concern. The importance of Hilbert's geometry is to be searched elsewhere than in its influence on research in geometry, so outdated after a century that journals dedicated to geometry would not even publish such material anymore. Hilbert's role was rather, as Weyl emphasizes, to show the possibility of formalization of mathematical arguments to the extent that one can check their correctness without the need to think of their intuitive content. Thus, geometry gave a model for how to formalize mathematics and how to pose the central foundational questions, such as those of consistency, completeness, and decidability. Yet, anyone who has started to really formalize Hilbert's proofs will soon have noticed that claims to rigorous formalization in Hilbert, such as Weyl's, are exaggerated: Consider Hilbert's very first axiom I1 which reads, from the popular seventh edition of 1930: "For two points A, B there exists always a line a such that

1079-8986/06/1203-0004/$2.10

both of the points A, B are incident with it." A careful reading of Hilbert's proofs shows that he never mentions axiom I1 in them, but a notation such as AB for a line just drops out of the blue into the proofs. A reading of the first edition of 1899 gives axiom I1 as: "Two points A, B distinct from each other determine always a line a; we shall set $AB = a$ or $BA = a$." There was no need to mention axiom I1 in proofs because its application was contained in the notation for the construction AB. A comparison of the various editions shows that Hilbert gradually changed his axiomatization from one with geometric constructions into one with purely existential axioms. The proofs in the book were not reworked, and there remains forever a basic incongruity between axioms and proofs in Hilbert's geometry; a real trap for those who might want to formalize Hilbert's geometry. Defects in formalization of proofs are not restricted to axiom I1 in later editions, but appear in various forms from the first one on. More examples can be found in *Formalization of Hilbert's geometry of incidence and parallelism*, ***Synthese***, vol. 110 (1997), pp. 127–141 by the present reviewer.

The present volume is the first of a projected six on Hilbert's foundational lectures, published under the general editorship of William Ewald, Michael Hallett, Ulrich Majer, and Wilfried Sieg. The next two volumes contain lectures on logic and arithmetic, and the rest material on foundations of physics and general lectures on scientific matters. The level of the lectures in this volume varies from notes prepared by Hilbert for giving the lectures into finished texts that were made available in mimeographed form. The latter were usually prepared on the basis of the actual lectures by assistants assigned to the task, with the final result corrected and approved by Hilbert. The original texts are divided into six chapters with dates from 1891 to 1902. The German original is faced with a translation into English and numerous footnotes comment the text from various points of view. A critical edition of the first edition of Hilbert's book is included.

The main interest of Hilbert's lectures stems, of course, from their relation to the ***Grundlagen der Geometrie***, either as anticipations or, in one case, as afterthoughts. The lectures of 1894 and of 1902 are especially interesting. To give a couple of samples, the former introduces (p. 74) a notation for incidence $a > A$ ("line a contains point A") and gives axioms conceived as rules of inference, such as

$$\frac{a = AB,\ a > C,\ A \neq C}{AC = a}$$

(The first premiss is stacked above the other two in the original.) Compare this uniqueness axiom to the work of a contemporary, Mario Pieri of the Peano school who notes explicitly in 1894–95, in his *Sui principi che reggono la geometria di posizione*, that a substitution axiom for equality is needed. Hilbert never arrived at a formal treatment of the equality relations of his geometry, even if he understood that these relations were ones of geometric congruence, not identity.

In the lectures of 1902, three years after the publication of the ***Grundlagen der Geometrie***, Hilbert repeats the axiom of parallels of 1899 as: "Through a point A outside a line a only one line b can be drawn that does not intersect a." He then notes the possibility of an alternative formulation: "If two lines a and b do not meet a third line c, they do not meet each other, either." There is no construction of a parallel present in the latter, but still Hilbert states that it "is really equivalent with our formulation, as one should immediately recognize" (p. 562). It is clear that Hilbert here takes geometric constructions for granted, irrespective of the type of axiomatization chosen.

The editing of the volume has involved a great deal of scholarly work. Hilbert's lecture texts have been studied and edited with an attention to details up to the last comma. The introductions to the various chapters record the circumstances of the lecture series and list to some extent their contents, often with comparisons between different lecture series. These introductions are short, except for the published version of 1899 and the 1898–99 lecture series that preceded it. We find that at crucial points such as the very formulation of Hilbert's geometrical axioms in the various editions of the ***Grundlagen der Geometrie*** and lectures directly relating to it, a systematic approach would have been motivated and useful. On pp. 419–423, great attention is given to details of the formulation of the axioms and theorems in the second edition, up to revisions of German orthography, but the editors seem not to have recognized the changes of wording and content in the incidence axioms and what these changes mean. With the later editions, no comment is made about the change of axiomatization from one with constructions into one with existential axioms, or about the defective formulation of the notion of parallelism up to the 1930 edition. The editing of this volume of lectures would have profited from a critical attitude towards Hilbert's achievement in geometry, such as in the review by Hans Freudenthal of 1957 that is mentioned on p. 411 as a "(very) critical commentary."

Hilbert's book of 1899 is often referred to as "the *Festschrift*." The reason lies in its appearance as a part of a volume celebrating the inauguration of the Gauss-Weber memorial in Göttingen. The editors state that "from its second edition (1903) on, it became a self-standing monograph" (p. 408). However, copies of the ***Grundlagen der Geometrie*** with publication at B. G. Teubner of Leipzig in the year 1899 exist at least at the mathematics library of Princeton University. Jean van Heijenoort referred to this edition in his ***From Frege to Gödel***, so there should be such copies also at Harvard. What the provenance of these copies is, remains to be cleared.

There is much to be praised in this volume that will certainly become a standard source for work on the development of the foundations of mathematics. There is just one point of serious criticism: The editors have not included any index of names, nor any index of subjects. Scholarly use of the volume is therefore made unnecessarily laborious. Where does Hilbert mention Frege or Peano? Is there any other consideration on decidability than the one about the basic relations in the 1894 lectures (p. 75)? Perhaps the editors could consider publishing indices for this volume in the next volume of the series to appear, to make finding answers to questions such as the above easier.

JAN VON PLATO

Department of Philosophy, P. O. Box 24, 00014 University of Helsinki, Finland
jan.vonplato@helsinki.fi.

MARTIN DIETZFELBINGER. ***Primality testing in polynomial time—from randomized algorithms to "PRIMES is in P".*** Lecture Notes in Computer Science, vol. 3000. Springer-Verlag, 2004, x + 147 pp.

One of the most interesting issues in computer science is the use of randomness (that is, coin-tossing) in computation. It seems clear that there are instances where randomness is absolutely necessary, for example when choosing a secret key in cryptography, or when breaking symmetry in a communications setting. More interesting is the question of whether randomness can be useful when one simply wants to perform a computation. The most common use of randomness in computation is in the simulation of a probabilistic situation in physics or engineering; the goal is to estimate the probability of an event occurring, or, more generally, to estimate the expected value of some random variable. Although it may seem natural to use randomness to do a probabilistic simulation, we know of no way of proving that it is necessary. (In fact, if $\mathbf{P} = \mathbf{NP}$, then such a simulation can be done almost as efficiently without using randomness.) Nonetheless, it is intuitive that it should sometimes help to use randomness when estimating an expected value.

Much less intuitive is the question of whether randomness can be useful when simply computing a "yes" or "no" answer, that is, when computing membership in some fixed "language" (that is, a set of finite strings).

The language that is the topic of this book is the set PRIMES consisting of those strings of 0's and 1's that are the binary representation of a prime number. For many years, until 2002, PRIMES was the most famous example of a language where the only efficient membership algorithms known used randomness. Both the Solovay-Strassen algorithm and the Rabin-Miller algorithm have the following characteristics: for some constant c, on an input of n bits: they toss coins, they run in time $O(n^c)$ (that is, they run in time bounded above by some constant times n^c), and they give the right answer with high probability. (Note that on *every* input there is a high probability of getting the right answer.) We usually denote by **P** (for polynomial time) the class of languages for which there is a *deterministic* membership algorithm with running time $O(n^c)$ for some c. No such algorithm for PRIMES was known until the appearance in 2002 of a paper by M. Agrawal, N. Kayal, and N. Saxena entitled *PRIMES is in P* (***Annals of Mathematics***, vol. 160 (2004), pp. 781–793). Even though the c used here is bigger than that of the randomized algorithms, this paper completely changed the way we think about primality testing.

Martin Dietzfelbinger attempts to do something in this monograph which is, as far as I know, rather unusual. This book contains a complete presentation of a proof of the "Primes is in P" theorem. It assumes that the reader understands mathematical induction and is familiar with the simplest algebraic manipulations from high school, but it assumes *absolutely no other knowledge whatsoever*. We are taken on a tour through number theory and algebra leading up to the final chapter. This chapter gives us a description of the algorithm, a proof of its running time, and a proof of its correctness. The version of the proof given in this book results from simplifications of the original proof that are due to the original authors, H. W. Lenstra, and D. G. Bernstein. The most important simplification is that a difficult theorem by E. Fouvry on the density of primes is no longer needed, allowing the presentation to be both reasonably short and completely self-contained. I liked this book very much, and I'll speculate below on who the ideal readers of it might be.

The book begins with some background material on algorithms and their running times. For example, it discusses the running times of the usual integer multiplication and division algorithms. Some very basic things about algorithms are described, and an informal model of computation is discussed, but I think it is probably necessary that the notion of an algorithm is not totally foreign to the reader. In order to get the absolutely best running time out of the AKS algorithm, it is necessary to use very sophisticated methods for integer multiplication and division, but these methods aren't necessary to understand AKS or to implement it as a polynomial time algorithm, and so these methods are only briefly mentioned in the text.

The next chapter contains some basic number theory and some basic number theory algorithms. We learn how to do modular exponentiation efficiently. We learn that for every two integers n and m (not both 0) there exist integers x and y such that $nx + my = \gcd(n, m)$. We then learn the Extended Euclidean Algorithm for efficiently finding x and y and see a proof that the algorithm works. But wait a minute—can't we use this algorithm and proof to show the existence of x and y, rather than doing this in a separate proof? This unwillingness to use an algorithm as an existence proof is one of my few complaints about this monograph. After some more fundamental number theory, the chapter closes with a very important lower bound on the density of prime numbers: the product of the prime numbers $\leq 2n$ is greater than 2^n (for integers $n \geq 2$). This theorem will be used to bound the running time of the AKS algorithm. As with other topics in the book, Dietzfelbinger gives a somewhat broader coverage than is actually needed, which is very nice. For example, he also proves an upper bound on the density of primes, although this is (I believe) unnecessary for AKS.

Three other chapters give a good introduction to the algebra of groups, rings and fields, with again a slightly broader coverage than is strictly needed for AKS. These chapters culminate in the following important theorem needed for AKS: if p and r are distinct primes, then $X^{r-1} + \cdots + X + 1$ is the product of irreducible polynomials of degree $\mathrm{ord}_r(p)$.

Dietzfelbinger takes a detour to give complete descriptions (with proofs) of the randomized algorithms of Solovay-Strassen and Rabin-Miller.

The last chapter deals with the AKS algorithm and proof. The most difficult part is showing that if the algorithm claims that the input number is prime, then it truly is prime. This proof is done carefully in a sequence of many small lemmas, and that's good. But it would also be nice if the author could give us a better overview of this proof that makes it seem at least somewhat intuitive.

The book has a few errors and typos, but most of these have been listed at the books's web site (given in the text).

Who is the target audience for this book? Top high-school students might be fascinated by the book. Many computer scientists who have forgotten much of their math background might find it easier to read the AKS proof here rather than in a journal. But the ideal target is first-year university mathematics students. The only math subjects these students are taking are calculus and maybe some linear algebra, and here is an opportunity for them to see a very recent, exciting result in mathematics. They will see more algebra in later years. But when they do study algebra, not only will the previous exposure to the topics in this book be useful, it will also serve as a strongly motivating introduction to a subject that can often seem abstract and pointless.

CHARLES RACKOFF

Department of Computer Science, University of Toronto, 10 King's College Road, Toronto Ontario Canada, M5S 3G4. rackoff@cs.toronto.edu.

DANIELLE MACBETH. ***Frege's logic.*** Harvard University Press, Cambridge, Massachusetts, 2005, xii + 206 pp.

This is a book about Frege's logic; it thus concerns an important chapter in the history of logic. Macbeth's goal is "to learn to read in a new way" Frege's *logical* notation and his writings about it. She rejects the common reading of Frege that she calls "the standard reading" that she traces back to Russell, and that considers the notation of modern quantificational logic, including the polyadic quantification, a variant of that of Frege's *Begriffsschrift*. Macbeth offers another way reading of Frege's conventions for the expression of "logical generality" that requires of us to "learn to read" the *Begriffsschrift* notation without transcribing it into another notational system as we usually do.

Macbeth finds at least four features of Frege's logical notation that invite doubt about the standard reading and poses the following questions about them.

1. Why is Frege's *Begriffsschrift* notation written in two dimensions with all logical signs (with the exception of identity) displayed in an array on the left? Here it would seem crucial to take Frege's logicist purposes into account since for his purposes he considers his bidimensional notation more perspicuous than Peano's unidimensional one.
2. Why does Frege require that the premises of an inference be "acknowledged to be true"? This is a stringent requirement; his notation includes a special sign for it (the so-called judgment stroke). Modern quantificational logic does without it, using the stroke instead to denote tautology or the deducibility relation.
3. Why does Frege use two different sorts of letters, both Latin italic and Gothic, for the expression of generality? Modern quantificational logic uses one set of letters for both free and bound variables and ignores the fact that there is no single sign for the existential quantifier in *Begriffsschrift* as a "minor omission".
4. Why does Frege claim (as of 1890) that the notion of "sense" (*Sinn*) is a properly logical notion, and that a sentence can express it even when it fails to have a truth-value? There is no room for such a notion in standard quantificational logic.

If by "logic" we mean "quantificational logic", says Macbeth, then none of these features of Frege's logic would be logically motivated; to explain them, she adds, we may have to consider the purposes underlying the design of his notation. He never expressed any misgivings about it. Whenever he faced objections against it he insisted upon its virtues. For his purposes, he insists, the *Begriffsschrift* combines *perspicuity* and brevity with "rigor of demonstration and maximal logical *precision*". Frege's sense of "perspicuity" and "logical precision" seems completely foreign to modern quantificational logicians.

Great innovations in the formal sciences often rest upon the design and development of a new notation that comes out of deep insights and much intellectual effort. Macbeth is not playing down Frege's originality, nor does she call into question his position in the history of logic as one of the great innovators. Her rejection of the view of Frege as one of the discoverers of modern polyadic quantificational logic derives from her view of his "*conceptual* notation" as combined with his understanding of logic is "something radically *different*" from what Peano and Russell developed,—"something essentially *new*" that has thus far been regrettably overlooked.

This raises at least two interesting historical questions.

First, did Russell learn quantificational theory from Frege? Macbeth says, Russell developed the full polyadic predicate calculus after he learned Peano's notation and the logical insights that underlie it. He was earlier familiar with Peirce's work on logical relations (based upon Boole's algebra). Peano and Peirce offered all he needed for his results. Their students will be happy to see their claims supported. By the time Russell studies Frege's work, Macbeth suggests, he "discovers" that Frege's notation and his own "could be read as a notational variant" of each other; and ever since then we have taken "Frege's notation as notational variant of quantificational logic". Macbeth rejects this view; her concern is not priorities but the understanding of the logical system that Frege had in mind when speaking of "a language of pure thought modelled upon the formula language of arithmetic"—as the subtitle of *Begriffsschrift* announces.

Second, what brought about the remarkable division of labour between logic (early writings 1879–1884) and philosophy of language (late writings beginning with 1892)? Macbeth argues that this remarkable division of labour is an artifact of the standard reading of Frege's logic. Frege himself, she says, came fully to understand the nature of his logic (first developed in 1879) only towards the end of his career: his logic "is brought to maturity" only in 1892, with the introduction of the distinction between "Sinn" and "Bedeutung" and of his claim that the notion of "Sinn" is properly *logical*. Her reading turns on taking this notion as fundamental for the proper historical understanding Frege's logic and thereby is opposed to that of Dummett, who assumes that towards the end of his career Frege wanted "to push his inquiries further" and "give a general account of the workings of language" not "qua logician but qua philosopher". Contrary to this she argues that Frege himself considered the *Sinn/Bedeutung* distinction as "a thoroughgoing development" of his logical views.

One may argue that there is nothing clearly defined that could be called "*the* standard reading" of Frege's logic against which to set Macbeth's reading. There are indeed many readings of Frege available, but they all share the underlying conception of language as functioning in a certain way, namely, by way of meaningful primitives that are then combined to form sentences the meanings ("Bedeutung") of which are to be understood by appeal to truth-conditions. It is what all other readings share that Macbeth wishes to take issue with. In *Begriffsschrift*, on her reading, all signs, whether simple or complex, designate only relative to an analysis of a sentence that expresses a thought and designates a truth-value. The two-dimensionality of Frege's notation does already suggests something more like an arithmetical table with its columns and rows, which is not to be read left to right, but instead looked at in the way you look at tables and arrays in calculus. What an analysis of *Begriffsschrift* sentences then gives you is a reading, and many are possible.

Our logic begins with a fundamental distinction between syntax and semantics; on Macbeth's reading Frege begins with thoughts expressed by sentences. Our logic begins with meaningful primitives signs (constants) and combines them. Frege begins with judgments and analyzes them into concept expressions: he comes by concepts only by analyzing *Begriffsschrift* sentences. In Frege's notation then one can form "concept words" ("complex predicates" in Dummett's terminology) of arbitrary complexity that in the context of a whole sentence can be variously analyzed. On the standard reading of Frege a sign for a concept that is the result of an analysis is "not really an expression—a bit of language—in its own right" (Dummett, ***Frege: Philosophy of Language***, p. 31). This difference relates to one of the most striking aspects of Frege's logic: a "*conceptual* notation" ought to reveal the inferential potential of concepts, an idea which will be completely lost in the different notational systems proposed after Frege. It may be objected that this is not to be reconciled with his mathematical practice and that some *Grundgesetze* derivations are difficult to account for non-quantificationally. Even if this were the case in some instances, Frege's guiding ideas about notation that proved most fruitful seem to point in a completely different direction. Macbeth argues that Frege held fast to these ideas throughout of his career.

After many failed efforts in presenting Frege as more than a predecessor to Russell, this book may finally succeed and so it is going to be discussed extensively. Macbeth takes notation as the starting point. Notations are clothing, but when tight they tell us something about the body they clothe. Of the three notations available, that of Frege, that of Peano and Russell, and the Polish, Russell's became canonical. The author intends to present Frege's notation as he himself understood it and as it contrasts with Russell's.

NORMA B. GOETHE

School of Philosophy, University of Cordoba, Cuidad Universitaria, 5000 Cordoba, Argentina. ngoethe@ffyh.unc.edu.ar.

ALAN BUNDY, DAVID BASIN, DIETER HUTTER and ANDREW IRELAND. ***Rippling: meta-level guidance for mathematical reasoning.*** Cambridge Tracts in Theoretical Computer Science, vol. 56. Cambridge University Press, 2005, xiv + 202 pp.

The central problem of automated reasoning is that there are so many different potential proof paths that it is hard to know what to do next. However, in the situation where a given is being used to prove a goal, such as in the step case of an inductive proof, the technique of rippling can be used to constrain the search and guide the way to a successful proof. This technique is the subject of this book, which aims to summarize the current state of knowledge in one coherent presentation. In the opinion of this reviewer the authors have achieved this aim, resulting in a useful volume that would be ideal as an introduction for a research student or as a reference for a more seasoned academic.

The book is well written, with all the techniques explained in detail, and most importantly for a practical automated reasoning text, well motivated with clear examples. The technique of rippling has been refined, analyzed, and applied so much over the last twenty years that the organization of a book such as this is a major challenge. The authors have grouped the topics into six main chapters, which will be considered in turn.

Chapter 1: *An introduction to rippling.* This chapter introduces rippling, situating it within the broader contexts of automated reasoning and formal methods. To get the reader up to speed as quickly as possible, the basic rippling technique is applied to a simple example, illustrating how wave fronts expand across the formula. There is also an extremely interesting section on the history of rippling, beginning with the origin of the term in Aubin's 1976 Ph.D. thesis at the University of Edinburgh (unpublished), and continuing in chronological order with a list of the significant extensions and crediting the researchers involved.

Chapter 2: *Varieties of rippling.* The simple picture drawn in the introduction is refined by introducing different varieties of rippling. In addition to outwards rippling, the analogy of ripples on a pond is stretched with the introduction of inwards and sideways rippling, and the picture becomes even more complicated with higher order rippling. However, despite the difficulty of taking it all in on a first reading, there is a real benefit in having all the different techniques collected together in a single chapter, easily accessible for future reference. In addition, the techniques are introduced carefully, and each one is motivated with a simple example that cannot be dealt with by the other techniques.

Chapter 3: *Productive use of failure.* This chapter describes a way that rippling can be used to improve the automation of induction proofs. If rippling fails to prove the goal from the given, the failure can be analyzed by a *proof critic* to suggest a possible *proof patch*. These patches can take various forms, from setting up the initial wave fronts differently to generalizing the conjecture or suggesting helpful lemmas. Although the ideas are made clear through well chosen examples, this reviewer found the level of detail to be insufficient for a thorough understanding of the methods. However, the complexity is perhaps inherent to the task, and in common with the rest of the book the chapter contains pointers to the literature to aid further study.

Chapter 4: *A formal account of rippling.* Rippling is a rewriting technique, and can be analyzed using the formal language of term rewriting. In addition to confirming that rippling's rewriting with wave fronts is well behaved with respect to substitutions, the question of whether rippling terminates can be addressed by defining a suitable reduction order. Although rippling can be understood without knowing the formal basis, studying this material gives a new perspective on the technique that repays the effort.

Chapter 5: *The scope and limitations of rippling.* The precise scope of rippling is pinned down by looking at twelve boundary examples: the first seven of which can be proved by a technique of rippling; and the other five cannot. In addition to providing more examples of rippling in action, the failures clearly show the limits of the technique and point the way for further research in the field.

Chapter 6: *From rippling to a general methodology.* Many techniques in automated reasoning annotate terms to carry out their task, just as rippling does with wave fronts. This chapter presents a general language for annotating subterms, and shows how many techniques can be encoded as inference rules on formulas in this language, including rippling, basic superposition and window inference. Annotated formulas can then be mapped to abstract formulas to constrain the search for proofs—enforcing the basic restriction, rippling rules, etc. In common with all the techniques presented in the book, this framework has been implemented in an automated reasoning tool.

In addition to the six main chapters, the book has a short conclusion summarizing the different aspects of rippling covered by the book and suggesting avenues of future work, and also two appendices. Appendix 1 presents the annotated term language of Chapter 6 in more detail, and gives a sound and complete unification algorithm, while Appendix 2 defines all the functions used in the book.

JOE HURD

Oxford University Computing Laboratory, Wolfson Building, Parks Road, Oxford OX1 3QD, UK. joe.hurd@comlab.ox.ac.uk.

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.