Chapter 1

Introduction

Omne ignotum pro magnifico.

Tacitus

The great advances in logic in the last century and a quarter saw a turn from its historical preoccupation with arguing and reasoning in favour of quite particular contributions to mathematics. It made possible important gains in both the *foundations* and the *methodology* of mathematics. The foundational contribution was largely of philosophical interest. It sought to establish a basis for logicism, for the reduction of mathematics to logic. The methodological contribution also has its philosophical significance, but it threw its net more widely, capturing the interest of those who thought that mathematics could only benefit from the rigour and the standards of exact proof that the new logic was in process of articulating.

It is difficult to overestimate the significance of the mathematical turn in logic. Not only did the new logic greatly narrow logic's former range of interests, it was able to do so only after determining that the traditional syllogistic approach to logic was inadequate for logic's new ambitions. Ever since its inception, 2500 years thence, logic had been in all essentials the logic of the syllogism. The mathematical turn brought a surprisingly abrupt end to Aristotle's long-lived hegemony.

Given the venerability and sheer persistence of that influence, it is perhaps not wholly inexplicable that mathematical logicians did not entirely break with the traditional line that logic is about reasoning and about arguing. There are plenty of textbooks on mathematical logic, including some of the best and most senior, in which we find it said, without a shred of irony or embarrassment, that mathematical logic is the most general, or the basic

theory of reasoning. Those of greater circumspection would claim that the new symbolic logic was the theory of *mathematical* reasoning.

It would be quite wrong to overlook the fact that mathematical logicians have been quick to recognize various respects in which the claim of logic to be a theory of (mathematical) reasoning is implausible. To that end, various distinctions have been invoked:

- *process/product*
- *descriptive adequacy/normative legitimacy*
- *actual circumstances/ideal conditions*

What these distinctions were thought to have had in common was that (a) while mathematical logic misdescribed the left side and properly described the right side, nevertheless, (b) left side circumstances could be thought of as *approximating* to right hand conditions in ways that would make it accurate to say that logic makes fruitful provision for the left side too.

Ever since its inception, and throughout the mathematical revolution, logic has been conceived of as a highly specialized investigation of *language*. In Aristotle's hands, the language of logic was Greek; in the hands of Frege, the language was the stylized notation of the *Begriffsschrift*. We see in this passage from natural to ideal languages a not inconsiderable development. But here, too, there were common constants. One was that all the target properties that a logic would seek to elucidate were represented as properties of linguistic structures. As Quine would say, with characteristic verve, 'Logic is linguistics on purpose'.

If modern mathematical logic attaches its findings to languages that no one speaks, or could, the complaint recurs that logic can't be about reasoning and arguing. Here, too, distinctions were invoked. Chief among them was that between

an actual sentence of a real language/its logical form in an ideal language

Considerable effort was expended to show that when conditions are right, *some* at least of the properties of ideal linguistic structures map to certain natural language structures in a principled way [Woods, 2002c, sec. 6]; for sober reconsideration, see [Woods, 2003, chapter 15].

We might refer collectively to these myriad efforts to support the claim that mathematical logic is a theory of reasoning and arguing as the *Standard Defence*. The Standard Defence is not lightly dismissible. It is closely patterned on widely accepted methods for showing that the empirical inaccuracies of our best scientific theories are discountable under the appropriate

approximation relations. No one dismisses the physics of frictionless surfaces just because its laws fail in nature, even as regards the pre-game, freshly Zambonied ice of Maple Leaf Gardens. All the same, the Standard Defence of mathematical logic has come under scrutiny from two largely unconnected sources, *computer science* (including AI) and *informal logic* and *argumentation theory*. A common reservation is captured by this question: Are the approximations postulated by the Standard Defence sufficiently intimate to justify its claim that logical theory may be seen as overriding empirical inaccuracy on the ground? Their answer, severally and jointly, is No. Informal logicians would observe that mathematical logic isn't particularly adept at modelling fallacious reasoning; computer scientists would point out the difficulties in getting plausible AI models out of standard logic. Some AI theorists would also note that certain features of reasoning and cognition generally are *sublinguistic* and thus lie exposed to systematic misdecription by theories that concentrate on investigating various properties of linguistic structures.[1]

Out of this welter of criticism certain themes have come to dominate. The authors of the present volume have particular interest in the following two:

1. Mathematical logic makes inadequate provision for the investigation of *practical reasoning*;

2. In its decontextual preoccupation with language, mathematical logic makes inadequate provision for the analysis of *cognitive structures*.

It is not to our purpose in this Introduction to adjudicate these claims; we want rather to *motivate* the book that follows. But we say in passing that much of the work in mainline logic itself these past thirty years has been to modify the standard or classical expression of logic in ways that take such criticisms seriously into account. The sheer scope and intensity of these adjustments is discernible in the fecund pluralism of the present-day research programme. Suffice it here to note developments in modal, deontic and epistemic logic; relevant and linear logic; dynamic and temporal logic; logics of action and labelled deduction; adaptive and preservationist logics; dialethic logic; dialogue and interrogative logic; and many more. To the extent possible, our approach in this book is to preserve the spirit of this collective attempt at logical self-reform in the cause of 'user-friendliness'. But we also wish to emphasize what many of these otherwise attractive

[1] Alternatively, some theorists take subdoxastic processes to involve symbol manipulation, but in a different representational system than that in which doxastic reasoning occurs.

systems of logic do not. We wish to respond positively and constructively to the challenges implied by the two basic complaints noted just above. Accordingly, what we expressly seek for is

1. a logic of *practical reasoning*; and

2. a logic of *cognitive systems.*

The present book is the first volume of *A Practical Logic of Cognitive Systems* (*PLCS*), of which three further volumes are forthcoming. One is in an advanced state of readiness, *The Reach of Abduction: Insight and Trial*, and a second is well underway, *Seductions and Shortcuts: Fallacies in the Cognitive Economy.* Following these will be a volume provisionally entitled *Formal Models of Practical Reasoning.* In each case our choice has been motivated by the conviction that these matters are of essential importance to practical logic, and that they are in need of further theoretical attention than they have hitherto received (and so cannot be thought of as closed parts of the research programme).

In most approaches, practical reasoning is distinguished in one or other of two ways. One sees its distinctive mark in the *content* of the reasoning; the other sees it in its *standards of rigour.* On the content side, practical reasoning is often said to be reasoning about what to do or how to solve problems; on the standards side, practical reasoning is thought of as governed by standards both less theoretical and less strict than those of 'pure' or 'formal' logic. We do not dispute these conceptions of the practical, but we do favour an alternative. We find it both intuitively attractive and theoretically fruitful to conceive of practical reasoning as reasoning done by practical agents, and in turn to conceive of practical agency in terms of the degree of access to key cognitive resources such as *information, time* and *computational capacity.* Given that such access is a matter of degree, practical agency is a comparative concept. As access enlarges, practicality recedes in favour of the theoretical, as we shall say. Intuitively, individual agents are paradigms of practical agency, whereas institutional agents such as NASA or Italian physics in the 1930s are theoretical agents par excellence.

This, the *resource-bound* approach to agency gives a conception of the practical that while different from, is not hostile to, either the subject matter or standards approach. It may be that practical agents in our sense deal rather more with matters of common or everyday interest to human beings than theoretical agents in our sense do; it may also be true that, since individual agents usually operate under press of scarce resources, the standards against which to assess their cognitive performance would be less rigorous and exacting as those required in retrofitting the Concorde. Even

so, it is clear that the subject matter, standards and resources approaches to practical agency are disjoint.

We have it, then, that a logic of practical reasoning is a certain kind of aspects of description of a practical agent. But not everything a practical agent does or is capable of doing is grist for the mill of practical logic. We shall therefore say that a practical logic is a description of certain aspects of the behaviour of practical agents under conditions that qualify it broadly as cognitive. Accordingly, we shall also find it useful to deploy the notion of *cognitive system.*

A cognitive system is a 3-tuple of a cognitive agent, cognitive resources, and cognitive tasks performed dynamically in real time. A cognitive agent is a being capable of perception, memory, belief, desire, reflection, deliberation, decision and inference. A practical cognitive system is a cognitive system whose cognitive agent is a practical agent in our sense, that is, an individual. A practical logic of the sort we are describing gives 'a certain kind of description' of a practical cognitive system. It is necessary to say something more about this.

Writing as logicians, we are interested in those aspects of cognitive behaviour for which a logician's more or less standard repertoire of target properties are instantiable in illuminating ways. In addition to properties such as *inference, consquence, consistency* and *validity,* we shall in due course add to the list notions such as *revision,* and, of course, *relevance.* Writing as logicians who have an interest in theories of reasoning that score well on the score of empirical adequacy, we seek descriptions of the behaviour of logial agents that deploy our logical vocabulary systematically and unsuperficially, but not in ways that take us to distant idealizations for which plausible approximation relations are hard to find.

On the face of it, our conception of a practical logic echoes a conviction of Bacon, who took logic to be a part of rational psychology. Although we stop well short of Bacon, ours is avowedly an approach to logic that could be called *psychologistic.* This will offend purists who, entirely correctly, have been quick to appreciate that model theory, proof theory, set theory and recursion theory have nothing to do with psychology [Barwise, 1977]. But there is more to our conception than is to be found in the four central domains of mathematical logic. In as much as we want our logic to give an account of aspects of the cognitive behaviour of practical agents, it is essential that psychological parameters not be overlooked entirely. In consequence, we find ourselves in agreement with those for whom the distinction between logic and psychology is neither exact nor exhaustive (see, e.g. Thagard [1982]).

There is an important sense, therefore, in which the logic of practical cognitive systems is not psychology. The relevant distinction is characterized best in *operational* terms, concerning which an analogy with mathematical logic is revealing. Mathematical logic gives an account of various properties (such as entailment, deducibility and consistency) of linguistic structures. Recall here Quine's quip: 'Logic is linguistics on purpose'. This should trigger an obvious question. *Why isn't logic linguistics*? Although some logicians have attempted to meet this question head-on (e.g. Quine [1960]), the answer for the most part is to be found by examining the different things that logicians and linguists actually *do* with the common matters that bind them. In each case the boundary between logic and linguistics is operationally discernible in the different things that logicians and linguists are *interested in* and *good at*.

It is the same way with the distinction between logic and psychology. Here, too, the difference is an operational thing. Even when, as in our case, the logician and the psychologist share a good many interests, our respective methodologies (what we are respectively *good at*) will serve to preserve the distinction non-trivially. If a logician has been mathematically trained, or has imbibed something of what goes on in computer science, he will bring to the table a competency in *formal modelling*. If the logician has been philosophically trained, he will bring to the table competency in *conceptual analysis*. In our approach, the two are systematically linked. In giving 'a certain kind of description' of aspects of the cognitive behaviour of practical agents, we do the following two things in order. First we give an analysis of the concepts that are central to the identification and basic description of such behaviour. A conceptual analysis may be interesting in its own right, but on our approach it is also input to a process of formal modelling. The logic in question is a linked partnership between *conceptual models* and *formal models*.[2]

We note in passing that there is nothing in what we are proposing with which to reprove, still less ignore, the extraordinary success of the modern logic of linguistic structures. What it may lack in psychological reality or applicability, it more than compensates for in results that are both indispensable in describing a cognitive agent's resources (for example, his ability to draw *consequences* or his partiality for *consistency*), and of obvious help to the theorist who describes such behaviour. So we disavow entirely the anti-formalist apostasy indulged in by some members of the informal logic community.

[2]So we do not cast our lot with John Cohen: 'if there is such a thing as psychology, it should consist (to paraphrase Bertrand Russell) of propositions which do not occur in any other discipline.' [Cohen, 1972, 9].

We have, in effect, re-pledged ourselves to the proposition that the laws of logic are the laws of thought. We are not alone in this:

> This is a doctrine which was popular in the last [=19th] century, but is now [=1979] very much out of favour. Nevertheless, I think it is true ... My thesis is that laws of logic are like [... scientific laws]. They are laws governing the structure of ideally rational belief systems ... They can be used to explain at least some of the features of ordinary belief systems, and the theory of rational belief systems in which they are embedded provides a framework for determining what remains to be explained about of belief systems. It thus defines a research programme. Ellis, [1979, v]

A logic that is practical in our sense falls within the ambit of the *pragmatic*. Historically, pragmatics is that branch of the theory of signs in which there is irreducible and non-trivial reference to agents, to entities that receive and interpret messages. By an easy extension, a pragmatic theory of reasoning is a theory in which there is express irreducible and non-trivial reference to cognitive agents. If in turn a cognitive agent is conceived of as a certain kind of information-processor, then a pragmatic theory of cognitive agency will provide descriptions of processors of information. Given that a logic is a principled account of certain aspects of practical reasoning, logic too is a pragmatic affair. If we ask, 'which aspects of practical reasoning are the proper province of logic?', we say again that the answer lies in operational arrangements. Practical logic is that part of pragmatics that investigates practical agency from the point of view of properties the logician finds interesting and is adept at analysing and modelling. Thus, again, properties such as implication, deducibility, generalization, relevance, analogy, plausibility and hypothesis, as studied by the methods of conceptual and formal analysis. The present work, *Agenda Relevance*, is an exercise in pragmatics in this sense. Given that the pragmatic enquiry that it triggers is subject to the methods of formal modelling, it may also be said that the book is an exercise in *formal pragmatics*; hence the work's subtitle.

As understood by a number of theorists, pragmatics is always a branch of the investigation of language. In the approach we take here, the importance of language can hardly be gainsaid. But since our emphasis is on cognitive systems, and since there are aspects of cognition that occur sublinguistically (or anyhow, subdoxastically), we are faced with a decision. One option is to reserve the *logic* of cognitive systems for those aspects of cognition that are linguistically manifest and to leave all else to the other branches of cognitive science. The alternative is to include the pre- or sublinguistic in logic's

reach. We do not suppose that this is a knockdown argument that decisively dismisses either of these two possibilities. Even so, the choice need not be arbitrary. Counting for the first option is the comparative *manifestness* of language, and the efficiencies engendered by this fact. Counting for the second option is the fact (or apparent fact) that the logician's target properties are also definable for structures that are not in the requisite ways linguistic. So, for example, it appears that some of our inferences are sublinguistic (or subdoxastic) and that, for beings like us, evasions of irrelevant information are largely automatic. Our own inclination, therefore, is to embrace (with appropriate caution) the more generous option. Accordingly, a practical logic is that part of a pragmatic theory that deals with the requisite aspects of practical cognitive agency at both linguistic and sublinguistic levels, and for which a suitably flexible notion of information will prove necessary.

It is well to emphasize that, in taking logic into a practical turn, we are not alone. Our approach, although developed independently, also shows a certain affinity to work done under the rubric of 'the dynamic turn', an approach to logic that emphasizes the 'interfaces with cognitive science, and the experimental study of how information and cognition works in humans once we set ourselves to study the psychological and neurological realities underneath ...' [van Benthem, 2001, p. 5].