


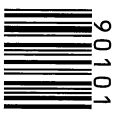
palgrave
macmillan

The jacket reproduces a photograph of rail tracks on the Chichi Line (© 1999 by Winsin Publishing Company, printed in Taiwan).

ISBN 1-4039-0287-9



9 0 1 0 1



9 781403 902870

Printed in Great Britain

palgrave
macmillan

Norms, Naturalism and Epistemology Jonathan Knowles

Jonathan Knowles

Norms, Naturalism and Epistemology

The Case for Science without Norms



Naturalism is winning ever-increasing popularity in philosophy. In this book, concerning epistemology, the author seeks to argue that naturalism brings with it certain costs. In the form of limitations on our philosophical ambitions. On the conception of naturalism defended, there are and can be no *a priori* norms for guiding our belief-formation: we must start our inquiries *in situ*, assuming some beliefs and the general reliability of our basic cognitive practices to justify others. Naturalised epistemology seeks to build on, but also goes beyond naturalism: to motivate and justify specifiable epistemic norms using naturalistic materials. The author argues that, whilst naturalism must be embraced, this more ambitious project is in vain. A systematic taxonomy of the possible varieties of naturalised epistemology is presented, followed by critique of each in turn. The underlying theme of the discussion is that to the extent one can genuinely justify naturalistic norms, they are not needed for optimal rational belief-formation.

For a note on the author, please see the back flap

Also by Jonathan Knowles
Vitenskap, rasjonalitet og common sense (SCIENCE, RATIONALITY AND
COMMON SENSE)


Norms, Naturalism and Epistemology

The case for science without norms

Jonathan Knowles

*Philosophy Department
Norwegian University of Science and Technology, Norway*

palgrave
macmillan



© Jonathan Knowles 2003

All rights reserved. No reproduction, copy or transmission of this publication may be made without written permission.

No paragraph of this publication may be reproduced, copied or transmitted save with written permission or in accordance with the provisions of the Copyright, Designs and Patents Act 1988, or under the terms of any licence permitting limited copying issued by the Copyright Licensing Agency, 90 Tottenham Court Road, London W1T 4LP.

Any person who does any unauthorised act in relation to this publication may be liable to criminal prosecution and civil claims for damages.

The author has asserted his right to be identified as the author of this work in accordance with the Copyright, Designs and Patents Act 1988.

First published 2003 by

PALGRAVE MACMILLAN

Houndmills, Basingstoke, Hampshire RG21 6XS and
175 Fifth Avenue, New York, N.Y. 10010

Companies and representatives throughout the world

PALGRAVE MACMILLAN is the global academic imprint of the Palgrave Macmillan division of St. Martin's Press, LLC and of Palgrave Macmillan Ltd. Macmillan® is a registered trademark in the United States, United Kingdom and other countries. Palgrave is a registered trademark in the European Union and other countries.

ISBN 1-4039-0287-9 hardback

This book is printed on paper suitable for recycling and made from fully managed and sustained forest sources.

A catalogue record for this book is available from the British Library.

Library of Congress Cataloging-in-Publication Data
Knowles, Jonathan, 1967-

Norms, naturalism and epistemology: the case for science without norms./Jonathan Knowles.

p. cm.

Includes bibliographical references.

ISBN 1-4039-0287-9

1. Knowledge. Theory of. 2. Naturalism. 3. Norm (Philosophy) I. Title.

BD161.K596 2003

146—dc21

2003046941

10 9 8 7 6 5 4 3 2 1
12 11 10 09 08 07 06 05 04 03

Printed and bound in Great Britain by
Antony Rowe Ltd, Chippenham and Eastbourne

Contents

Preface

vii

Introduction

1

1 Traditional Epistemology and Naturalism

6

Introduction

6

1.1 Norms and epistemic norms

7

1.2 Traditional epistemology

9

1.3 Critique of TE

14

1.3.1 The epistemological problems

15

1.3.2 The metaphysical problems

18

1.3.3 The methodological problems

21

1.4 Naturalism

23

2 Naturalised Epistemology

34

Introduction

34

2.1 Quine, psychology and anti-foundationalism

34

2.2 Goodman and the method of reflective equilibrium

39

3 Anti-psychologism

48

Introduction

48

3.1 In what way is AP a genuinely distinct position?

50

3.2 The dominant research strategy in the philosophy of science (DRSPS)

53

3.3 The fundamental problem for AP

55

3.4 Kuhn on exemplars and tacit knowledge

60

3.5 An alternative conception of AP

64

4 Anti-foundationalism

66

Introduction

66

4.1 Mill's inductivism

69

4.2	Evolutionary epistemology	79
4.3	MREAF	81
4.4	Laudan's normative naturalism	83
4.5	Epistemic norms as hypothetical imperatives	99
4.6	Why does it appear as if scientists use norms if they do not need to?	105
5	Psychologistic Naturalised Epistemology	114
	Introduction	114
5.1	Ballpark psychologism and 'mixed' varieties of NE	119
5.2	Are we naturally rational?	124
5.3	Can norms be principles of our reasoning competence?	129
5.4	Normativist psychologistic NE	135
	Conclusion: Science without Norms	143
	Notes	150
	Bibliography	176
	Index	181

Preface

This book represents the fruits of my main research in philosophy over the past five years. It began life as a project description for which I gained a post-doctoral fellowship at the Philosophy Department, University of Oslo in 1998, and has been completed after I started working, in 2000, at the Norwegian University of Science and Technology (NTNU) in Trondheim. Further back in time, its original inspiration was my teaching on a course at the University of Oslo entitled 'Vitenskap og rasjonalitet' (*Science and rationality*), for which I tried to write a textbook (cf. Knowles 2000b). This teaching and writing gave rise to a lot of ideas; you hold the distillation of what I see as the most coherent and original of these.

The main ideas of the book and some of the material have been presented at various different seminars and conferences; two portions of it are modified versions of published journal articles (or parts thereof). Chapters 2 and 5 derive from notes for a graduate seminar held in the Spring semester of 1999 in Oslo; the Laudan material of Chapter 4 from another graduate seminar, at the same institution, in the following Spring, held together with Professor Nils Røll-Hansen. An overview of prototype versions of Chapters 1, 2 and 3 formed the basis for a paper presented at the 23rd International Wittgenstein Symposium, *Rationality and Irrationality*, in Kirchberg, Austria, in August 2000 (published: cf. Knowles 2000a). Portions of Chapter 4 were presented as a paper entitled 'Do we need norms in a naturalised epistemology?' at the *Vitenskapsteoretisk Forum*, NTNU, in February 2001, and in the UK at the philosophy societies of the Universities of Birmingham and Lampeter and of King's College London during March 2001. An overview of the book's overall argumentative thrust, based around the concept of *wide reflective equilibrium*, was presented at the *4th European Congress of Analytical Philosophy* in Lund, Sweden, June 2002 (published: cf. Knowles 2002a). I thank all those who gave me feedback on these occasions.

The portions of the book that derive from previously published work are, first, Sections 4.3 and 4.4, which constitute a modified version of 'What's really wrong with Laudan's normative naturalism'

in *International Studies in the Philosophy of Science* volume 16:2, 2002, pp. 171–186; and, second, the latter half of Section 5.4, which is a modified (and abridged) version of Section 2 of 'Is folk psychology different?' in *Erkenntnis* volume 57:2, 2002, 199–230. These appear by kind permission of, respectively, Routledge (Taylor & Francis Group, <http://www.tandf.co.uk>) and Kluwer Academic Publishers.

I acknowledge personally: Jo Sivertsen for having conceived the aforementioned course *Vitenskap og rasjonalitet*, and written the original textbook (cf. Sivertsen 1996), without which I would almost certainly not have written this book; Nils Roll-Hansen and Alfred Fidjestøl for discussion of and comments on the Laudan material; Harvey Siegel for putting me by way of, and making me take seriously, the idea of instrumentalistic epistemology; Øystein Linnebo and Mikael Janvid for very useful discussion of naturalised epistemology; and David Henderson, both for listening patiently to my ideas over a beer on a balmy evening in Kirkeberg, and later, for more sober, written comments on a draft version of the whole book.

On the more practical side of things, I would like to thank Halvor Eifting for his not inconsiderable efforts in gaining permission from Taiwan for use of the cover photograph (thanks also to Arthur Sand for reproducing this); Gary Barham from Macmillan, for suggesting the project and advice in the initial stages; and my colleagues at the Philosophy Department in Trondheim for their understanding attitude towards my personal circumstances, something which has definitely facilitated the completion of this book.

Finally, thanks to Janne, Edvard, William, Gina, Janet and Lewis – for being there.

Introduction

Naturalism is definitely the flavour of the month in analytical philosophy – indeed, of the last three decades or so. From W.V.O. Quine to Sir Peter Strawson, Jerry Fodor to John McDowell, Jaegwon Kim to Donald Davidson – everyone wants to be in on the naturalism act. This book is not an exception to that trend, but it is only partly concerned with an articulation and defence of a naturalistic philosophical position. The main arguments of the book ride on the naturalistic wave and seek to draw some significant and perhaps unexpected consequences of doing so.

Its focus, as the title suggests, is epistemology (the theory of knowledge) and specifically the issue of norms in epistemology – *epistemic norms*, as they are known (as opposed to moral norms, social norms or rational norms in a broader sense). Insofar, it is a book about what has, after Quine,¹ come to be termed *naturalised epistemology* (NE). However, it is not about NE in what I would call 'the narrow sense'. My philosophical background does not lie first and foremost in epistemology, and the ideas for the book sprung originally from my thinking in the philosophy of psychology and the philosophy of science. This is reflected in the end-product, which, though I hope clear, cogent and provocative in its argumentation, will perhaps be felt by many main-stream epistemologists to pass over much of what makes epistemology exciting – questions about the proper analysis of knowledge, justification, belief, and so on, plus how to respond to scepticism. This, I need barely add, is not something I see as a fault in my book. Rather, my somewhat cavalier march into, for me, foreign territory will, I hope, give active

practitioners in epistemology, and especially naturalised epistemology, pause for thought about the credentials of arguably their central project: that of deriving and justifying epistemic norms.² In so doing, I also hope to point up a general philosophical lesson that should be of interest to a wider audience.

What I aim to make at least plausible is that there is, along this central dimension, no work these practitioners need to perform. In particular, I will be arguing for the following overarching claim:

Any theory that makes use of just naturalistic resources to specify epistemic norms is either defunct or superfluous to requirements.

In embracing naturalism in the theory of knowledge – as I will argue we absolutely should do, in Chapter 1 – we cannot, *contra* the presuppositions of many, continue to hold out the hope of delivering on the substantive goals of the traditional epistemological project, as conceived by a thinker like Descartes. We cannot simply do away with the bath water and retain – perhaps in some envisaged collaboration with psychologists and natural scientists – the baby of the activity which that project's *apriorism* gave a point to. It strikes me that many philosophers of today see the epistemological and metaphysical perspectives of former generations as intellectual vices barring the path to real philosophical insight and progress. This piece falls into the ranks of those which rather see the traditional ideas as noble but forlorn fictions, which, once perceived for what they are, leave behind only a sense of how lithe the philosophy and philosophers can actually produce by way of concrete results. This should not make us leery of philosophy itself, of course, but only of a certain conception of its remit (an attitude which is winning increasing popularity amongst thinkers today).³

Perhaps this prologue will suffice to indicate that this is not an introductory text, insofar as it assumes familiarity with basic philosophical ideas and concepts, as well as the rudiments of the history of philosophy. On the other hand, it is not a very technical book, and it should be readily accessible to anyone interested in contemporary analytical philosophy. In the rest of this introduction, I will provide a more systematic introduction to my theme and sketch briefly the structure of the book.

My overarching thesis has its point of departure in a perennial area of concern in philosophy: that of giving and justifying a theory of sound epistemic inference, of identifying the kinds of belief- and hypothesis-forming practices that will yield, or are most likely to yield the kinds of beliefs and hypotheses we cognitively value – traditionally, those that are most likely to be true, or approximately true (whatever that amounts to). Since the birth of Western philosophy, this project has been seen both as vitally important insofar as it should provide a *foundation* for knowledge, and as primarily amenable to the *a priori* methods of reflection traditionally employed in philosophical thinking. The envisaged product has been seen a set of correct epistemic norms which we may use to steer and evaluate our practices of belief-formation, including our systematic research into the hidden well-springs of the universe, in order to gain knowledge.⁴

In more recent years, however, there has been a move towards reassessment of this traditional project. One of the main claims of Quine's seminal article referred to above is that attempts to specify epistemic norms by scientists and philosophers over the ages have conspicuously failed: all the proposals have been susceptible to devastating counterexamples. Much recent philosophy of science has also been aimed at undermining various different proposals concerning the overarching epistemic and methodological norms of science, indeed, the very idea that such norms exist or should be sought.⁵ For Quine, these failures and insights should lead us to give up doing epistemology as a substantive, normative discipline, and to turn, instead, to descriptive empirical psychology – to charting the course of events whereby we *actually* acquire knowledge, rather than laying down prescriptions for how we *should* do this. The line to be defended in this book concurs in many ways in this line – though, as I shall indicate, it is important to be clear about exactly what it does and does not imply. Many, however, have seen Quine's response to the failures of traditional epistemology (TE) as unwarrantedly pessimistic.⁶ For such philosophers, though we certainly need to get out of the armchair and look to empirical findings, it is still possible to deliver something like the epistemic norms dreamed of by the traditional thinkers. This is the project of what I call in this book *NE*.

My main thesis is thus that *NE* (in this sense), like its traditional ancestor, is misguided. In particular, I will be arguing that we cannot

make sense of *naturalistic epistemic norms*, or at least of the *need* for such norms for optimally rational belief-formation – that is, for norms proposed from within a *naturalistic epistemological theory*.⁷ In a word: *naturalised epistemology is either defunct or superfluous*.

Neither Quine (around 1953–1969) nor myself are thus naturalised epistemologists on the foregoing understanding of the expression – even though we both are what I call *naturalists* and reject traditional epistemology.⁸ To show that NE (in the sense specified) is either undoable or not worth doing, though obviously a negative project, is clearly of utmost significance, for at least two reasons. First, there are many philosophers today who are seeking to give such theories, or who argue that they can and should be given; what I have to say should, as said above, give them pause for thought. Secondly, if their theories fail for the reasons I give, and if TE fails, then we need to face up to the idea of rational belief-formation without norms – to *science without norms*, as I put it.⁹ Though perhaps unpalatable in prospect, I believe such an idea is fully tenable, and articulates well the fundamental insight of Quine's holistic idea of the web of belief as a naturalistic metaphor for science, purged of 'First Philosophy'. In both of these ways, my arguments will also serve to reinforce the 'deflationist' conception of philosophy's remit alluded to above.

The rest of the book paves the way for this conclusion, and is organised as follows.

Chapter 1 outlines the motivation for NE by sketching the project of traditional, *a priori* epistemology, the fundamental weaknesses of this, and the motivation these supply to *naturalism*. I defend a variety of naturalism that meshes with Quine's holistic conception of science, and suggest how it might be seen as constituting a basis for the more ambitious and substantive theories that would class as what I term NE.

Chapter 2 motivates a taxonomy of different varieties of naturalised epistemology, in which I distinguish three broad classes: *anti-psychologism*, *anti-foundationalism* and *psychologistic naturalised epistemology*, the outlines of which I sketch. The discussion is not only pursued to a large extent through the lens of Quine's NE, but also draws on Nelson Goodman's conception of the justification of inferences rules as using what has come to be termed the *method of reflective equilibrium*.

Chapter 3 offers a critique of the first of these three classes of NE, *anti-psychologism*, in its various different forms. The basic idea behind anti-psychologism is the use of the method of reflective equilibrium (of a 'narrow' or 'wide' variety) to justify epistemic norms, understood in a naturalistic manner, something which can be seen as operative in relation to both scientific- and common sense belief-formation. Chapter 4 offers a parallel critique of the second class, *anti-foundationalism*. Its two main *foci* are John Stuart Mill's conception of inductive rules in science, and the recently popular *instrumentalistic* conception of epistemic norms as defended by, amongst others, Larry Laudan (whose particular views I discuss at some length).

In relation to both anti-psychologism and anti-foundationalism, I argue that, though the norms they propose perhaps can be coherently derived and justified, they are not necessary for optimal rational belief-formation. Chapter 4 also offers an account of how this can be so consistent with the apparent widespread reliance on norms in actual science.

Chapter 5 concerns first and foremost the final class of NE, *psychologistic NE*. This comes in both a *descriptivist* and *normativist* variety, but the idea behind both of them is that norms can somehow be identified with psychological principles that naturally govern our reasoning and inference practices. I argue that the very idea of such principles being norms is hard to credit, and at the very least unmotivated, given naturalism. The chapter also includes a general discussion of the significance of psychologism in epistemology, plus critical discussion of what I call 'mixed' varieties of NE, which combine different sources of information or different theories of norm in their overall theory.

The conclusion – *Science without norms* – considers the implications of these critiques for our epistemic practice and for philosophy. It also indicates some caveats to my overarching thesis (though not ones that compromise its significance!). It closes with a brief discussion of the relationship between the line I espouse and the views about the role of norms or rules in reasoning and rational activity associated with Wittgenstein and Lewis Carroll.

1 Traditional Epistemology and Naturalism

Introduction

My overarching thesis – as I will henceforth refer to it – is that, to the extent we can make sense of naturalistic epistemic norms, there is no need for such norms in the process of rational belief-formation. To appreciate the significance of this claim, one of course needs to know something about NE (and how I understand this), and to do that, one also needs – since epistemology was originally conceived and pursued as a discipline that could only be described as anything but naturalistic – to look first at how epistemology is conceived on a more traditional model. In this chapter, I will present the outlines of what I call TE before showing how this, arguably, breaks down, paving the way for naturalism and NE.

The aims in this chapter are largely expositional, and the substance lays no great claims to originality (apart perhaps from the exact characterisation and defence of naturalism). The picture of TE I offer is simplified and to an extent, perhaps, a caricature; the discussion, at least, rides roughshod over a number of philosophical concepts and distinctions. A more thorough presentation would, however, bog us down in issues that are not relevant to my dialectical tack. The basic idea behind my characterisation of TE is both historically real and, at a certain basic level, intuitive, whilst the objections I raise to it are by any account strong ones, as well as being widely accepted, in some form or other, by most naturalistically inclined philosophers today.

In this way, the discussion serves its chief purpose of locating and motivating the ideas of naturalism and NE. It should be borne in mind that this book is centrally about the failings of naturalised, not TE (we are riding on the wave of naturalism!). Thus, if this chapter is at times written in the manner of an obituary, and a committed traditionalist is alarmed to see her favourite theory being buried (as she sees it) alive, she can console herself with the thought that this is not really the book for her.¹

To begin with, however, I should say a few words about the central concept which both traditional and NE (as I am understanding them) lay stress on, and which I will ultimately be arguing need have no application in our belief-forming practices.

1.1 Norms and epistemic norms

This book revolves around discussion of different conceptions of epistemic norms, so it is appropriate to begin by precisifying how I understand the basic concept *epistemic norm*. I am not going to go into great detail about this or discuss relevant literature, because I think the concept is clear enough at the outset, and that the relevant complexities will emerge in later discussions of the book. But a little to set the scene will be of help.

Firstly, then, I see a norm – in general, whether epistemic, more broadly rational or even moral – as a proposition, principle or statement with an essentially *general normative* content. A normative content is one which (in English) is canonically expressed by the use of the imperative mood, or by modal expressions like ‘should’, ‘must’, ‘it is permitted to’ and so on. By stressing the idea of generality, I mean to exclude, say, a command to do such and such on a particular occasion; I do not mean that norms must concern everyone (*viz.*: *Everyone must help to alleviate the present famine in Ethiopia* would not be a norm by this definition, whereas *People in the West should help alleviate famine wherever and whenever it occurs* would be). Thus, if the canonical forms of the norms we are concerned with are viewed as expressible by sentences like ‘Do X!’ or ‘Everyone shall do X!’ or ‘If situation Y obtains, do X!’, X (and Y) must always be understood as referring to a *type* of action or inferential move.²

In relation to matters epistemic, I thus understand a norm to be a rule or prescription³ behoving certain general reasoning and belief-forming practices (possibly under certain specified circumstances), not the acceptance of a particular proposition, hypothesis or whatever (at a particular time, or always). I take this to be reasonably in accord with ordinary usage, and in any case to capture the notion that is significant from a philosophical perspective. The whole point of epistemic norms has been to systematise, regulate and justify belief-formation; it has never been supposed that 'particularistic' normative statements – e.g., 'Believe what you just saw' – could suffice as (part of) a substantive epistemological theory.

Here are three further points of clarification concerning epistemic norms (or just 'norms', as I will for the most part say) which might be useful at this stage in the discussion:

Firstly, as I have suggested and as will become further apparent, I do not suppose norms to be generally tenseless, that is, valid for all times, or valid for all people (if they are valid). Some norms will be, others not (see Section 1.2 on *universality*).

Secondly, I will not be distinguishing systematically between norms in respect of their different types and grades of modal force (*viz.*: 'Do X!' versus 'Do not do X!', 'It is permitted to do X', 'Be strongly inclined to do X', 'Consider doing X', etc.). Nothing I have to say should depend on the existence of such differences, whilst as far as I can see the different theories of norms I consider naturally open for them. It is also usual to distinguish between what are generally called *categorical* norms (or, after Kant, *imperatives*) of the form 'Do X!', and *hypothetical* or *instrumental* norms or imperatives, of the form 'If you want Y, do X!' or 'Since you value Y, do X!'. In contrast to the others, this is a distinction that I will be discussing explicitly as it relates to more recent work in naturalised epistemology and philosophy of science (cf. Sections 4.4 and 4.5).⁴

Thirdly, and finally, I allow that norms can, in principle, be based on or derived from non-norms, that is descriptive propositions – for example, that norms for deductive logical reasoning might be based on descriptive logical truths. As far as TE is concerned, it strikes me that the distinction between norm and any truth it might be based on is rather blurry: whereas in NE, the nature of this relation has, of course, been a large issue (we shall return to it later in this chapter).

1.2 Traditional epistemology

The central idea behind what I will be calling traditional epistemology is, first, that we have available epistemic norms that we must use to steer the formation of and to justify our beliefs about the world and, moreover, that these derive from some source, some body of knowledge, that is essentially *prior* to and *independent* of the beliefs whose formation they steer and justify. The norms are, as it is put, *a priori*⁵ in character – given in advance of and thus capable of serving as the basis of the claims that needs grounding. The classical statement of this conception of epistemology and epistemic norms is arguably to be found in the philosophy of Descartes. It is of course Descartes who effected the epistemological turn that indelibly influenced all subsequent philosophy, conferring upon the theory of knowledge the status of *prima philosophia*, a foundation for all other cognitive disciplines. A more explicit connection to traditional epistemology, as we understand it, can be appreciated through Descartes' stress on uncovering *the method* for the pursuit of scientific work of all kinds – something that was for him precisely a set of *apriori* norms for the attainment of knowledge of the world.

But TE is not just something we find in Descartes' philosophy, nor do I mean this to be definitive of it. A similar commitment is also evident in the thinking of one of Descartes' more empirically inclined contemporaries, Francis Bacon, who set up elaborate canons for inductive reasoning to replace Aristotle's syllogisms as the appropriate rules for the new science. In general, TE forms the backdrop, albeit sometimes implicitly, for the highlights of epistemological thinking from the seventeenth century onwards (many of which broach naturalistic reactions to traditionalism): that of Descartes himself on the foundations of the method; of Hume on the problem of induction; of Kant and his conception of *transcendental idealism* as the appropriate backdrop for understanding the *apriori*; of Mill, Whewell and other nineteenth century figures concerned with the nature and credentials of scientific method; and not least of the twentieth century's philosophy of science, as inaugurated by logical positivism and pursued further by Popper, Quine, Kuhn and others. There are also many, working in philosophy today, who I think could fairly be classified as traditionalist epistemologists, and who tussle with contemporary naturalists.⁶

Turning to the other end of the historical spectrum, I would also maintain that there is much in ancient Greek philosophy that can be seen as cleaving to the fundamental tenets of TE. The very idea of an axiomatic system, classically exemplified in Euclid's *Elements*, is in many ways the formal manifestation of the fundamental epistemological intuition behind TE. Many today are inclined to downplay the epistemological significance of axiomatic systems. Nevertheless, it is clear that axiomatic method does historically have such significance, not least insofar as it owes its first explicit exposition to Aristotle's conception of scientific method in the *Posterior Analytics* – a conception which it can seem natural to view, as Bacon did, as a specific suggestion concerning the epistemic norms for substantive scientific work.

What then, in a little more detail, is the conception of epistemology underlying TE? The intellectual problem we face is that of justifying our beliefs, or – perhaps more pertinently, at least at the time of the scientific revolution – to decide what it is we should, rationally, believe, and thus can be said to know, about external reality. Feeding into this process we might imagine some initial data in the form of sensory impressions or intellectual intuitions, but moving from these to conclusions about what to believe requires normative guidance about correct reasoning: furthermore, the veracity of the initial data must be grounded. The whole process thus rests on the availability of a set of norms or rules which will tell us what we can trust of our initial data, plus how to transform these into a coherent conception of reality. According to TE, it is part and parcel of this idea that these norms – the epistemic norms – be apriori, in the sense that they must be known or knowable independent of the substantive beliefs whose formation they are meant to dictate and justify. If they were not thus knowable, we would, it is alleged, be left without a rational basis for our beliefs; we would face a circularity in our reasoning which would undermine our goal of discovering the nature of the reality, or at least that of ever justifying the belief that we have discovered it.

At this point, one might delve into the details of a selection of the historical examples of this project I adduced above, but I will instead make some remarks of a more systematic character to further clarify TE. Some of these will also help to understand the various naturalistic reactions to it to be discussed later.

First, a few words to clarify the notion of the apriori that is relevant to TE. It is of course traditional to contrast apriori knowledge with that derived from (sense) experience. However, apriori – as most would now accept, and at least as far as we are concerned – should not simply mean that same as *innate*. It might be that one learns many mathematical truths through experience (by counting oranges, say), but this does not, *ipso facto*, render them non-apriori. Conversely, recent developments in cognitive science would suggest that the idea of innate knowledge should be taken seriously, but that would seem to have no implications for the idea of apriori norms for enquiry in the sense assumed by TE (though it might, as we shall see, have implications for naturalised epistemology). A related point is that, on my understanding of TE, classical empiricism will be just as beholden to apriori norms as classical rationalism – precisely in order to underwrite its reliance on observation.⁷

These are relatively humdrum observations in view of recent discussions of the apriori. My aim in making them is to stress that the notion of apriori assumed by TE is a very strong one, indeed, is meant to be that of *pure* aprioricity. I will not be offering a detailed discussion of this notion, since my overarching aim is to motivate the naturalism that rejects it. Nevertheless, I think it can be captured roughly and in a way that does not prejudice subsequent evaluation, by the idea that apriori truths or norms differ from *a posteriori* ones in that the latter are *knowable* – with separate stress on *know-* and *-able* – at least partly through experience. Thus, I think it would be reasonable to deny that one could really come to *know* a mathematical truth like $2 + 2 = 4$ (a paradigmatic apriori truth) through, or with the help of, experience – by counting oranges or being told; whereas it seems obvious that a piece of innate knowledge about physical bodies 'hard-wired' in us in the course of evolution⁸ could be so acquired, and therefore would not be classed as apriori.⁹

The point of norms being (thus) apriori for TE is, as already indicated, that they should not presuppose substantive world-knowledge; that they should instead provide a *foundation* for this knowledge. This means that the set of correct epistemic norms (or our knowledge of these) is *justificationally* prior to our substantive knowledge of the world: the latter depends in this way on the former but not (ever) *vice versa*.¹⁰ Purely temporally, our substantive knowledge could come first, in that we might entertain, or even believe,

certain propositions without bringing to bear, or having access to, the norms. However, these beliefs about the world would only become knowledge, be justified, when warranted in the light of the norms.

(This last sentence might be felt to be ambiguous: is it necessary for our beliefs to be [or become] knowledge that the possessor of these beliefs *herself* applies the norms to justify them – that she self-consciously thus comes to see the beliefs as justified – or is it sufficient that she merely holds the beliefs, and that they are in fact justified in relation to the appropriate norms (in a non-accidental manner)? In recent epistemological theorising, this issue has been characterised in terms of a distinction between the positions of *internalism* and *externalism*, holding respectively to the former and the latter disjuncts in the question just posed. It seems fair to say that TE in most of its manifestations has held to an internalist view on this matter, but in fact this does not seem essential to its central tenets. The debate between internalism and externalism, along with many other issues in contemporary epistemology [such as the tightly interrelated issue of exactly what knowledge is], will not occupy centre stage in this book. Nevertheless, I would like to take the opportunity now to point out that my overarching thesis [that naturalistic norms, insofar as they are genuine, are not necessary for optimal belief-formation] is not merely a form of externalism – the idea that one can have fully rational beliefs without applying the relevant norms oneself. After all, externalists do not, as far as I am aware, claim that deriving epistemic norms is *quite generally* superfluous to requirements.)

It is implicit in TE that the norms hold *universally*. In saying this, however, we should immediately distinguish between two different ideas: on the one hand, that the norms are universal in that they always and necessarily yield true beliefs, in something like the way logical rules are meant to give true premises; on the other hand, that the norms are universal in being the same for all rational agents at all times. We can dub the first idea *universal validity*, the second *universality*. As I understand TE, it is only the latter that is, strictly speaking, presupposed; in other words, it does not seem an essential part of TE that the norms necessarily generate true beliefs, or justify with 100 per cent certainty. If statistics can be viewed as apriori in the relevant sense, there will no doubt be apriori norms that may enjoin belief in propositions which are false. This brings out one way in which TE, as I have so far characterised it, is compatible in

principle with *fallibilism*: the idea that knowledge claims need not be 100 per cent certain to be justified. In fact, there seems some room to doubt whether TE really is compatible with fallibilism, given that its norms are meant to be apriori. For could we know, strictly apriori, that a certain rule is by and large, but not universally valid? (What would be the relevant feedback-telling us this, and how would we evaluate it?) That issue, however, is a portal to a large debate which absolutely need not detain us here. For present purposes, we shall allow that TE might operate with fallibilistic and/or statistical norms.¹¹

What *does* seem axiomatic for TE is just that norms are universal, that is, valid (universally or otherwise) for all times and for all rational agents, for this would seem to follow simply in virtue of their being purely apriori and foundational for all enquiry. The Canadian philosopher Harold Brown has argued that universality lies close to the heart of our common sense concept of rational belief-formation: given a certain body of information and the task of drawing a conclusion from it, it seems there cannot be several different, conflicting conclusions that are all rationally best.¹² Perhaps this intuition can in turn be explained by reference to the idea that the rational aim of belief-formation is truth which, being unitary, provides a universal yardstick for all beliefs (though I do not mean to imply that TE is committed to truth as the goal of enquiry).

Universality can be seen as entailing a tacit disavowal of relativism, the idea that either knowledge or truth (or both) are irreducibly relative to groups of individuals, intellectual epochs, cultures or whatever. In this book, I will be assuming that at least when it comes to epistemic norms – traditionally or naturalistically construed – they are valid if they are valid in an absolute, non-relativistic sense. A relativism about norms would undermine itself insofar as norms that were supposed to be only relativistically valid would lose their normative status – they would not genuinely state what one *ought* to believe, for others in epistemically identical situations would not be under the same obligation.¹³ However, being absolutely, that is, non-relativistically valid does not, at least obviously, strictly imply being valid for all times and all rational agents. The latter certainly seems to be an implicit presupposition of TE. But it is not, as we shall see, a presupposition of all varieties of NE.

Finally, it should be stressed that it is not a presumption of TE that the correct set of epistemic norms will constitute anything like an algorithm for the formation of belief – a mechanistic procedure that can be applied unthinkingly and automatically to yield a unique output, given a certain input (I leave open whether there are historical examples of thinkers who have viewed epistemic norms thus).¹⁴ There has been a deal of discussion of this idea in recent philosophy of science, particularly after Kuhn stressed that his scepticism about rules in science really only amounts to a scepticism about the idea of a methodological algorithm. I shall discuss Kuhn's idea later in the book (Chapters 2 and 4), but for now the point is merely that TE is not committed to this surely implausible view of epistemic norms.

1.3 Critique of TE

TE may in many ways seem to enunciate an attractive and indeed necessary project for the pursuit of rational cognitive enquiry, even if not quite in the manner conceived of by Descartes; there are also many philosophers today who would concur in such an assessment.¹⁵ The view of most who would call themselves naturalists, however, is that it suffers from irredeemable problems of diverse kinds, and that apriori norms do not and cannot constitute, contrary to first impression, a necessary foundation for our substantive world-knowledge.¹⁶ In this section, I shall review some relatively well-known objections to TE, alluding along the way to naturalism's alternative conception of our epistemological condition.

I see the chief problems with TE as falling into three main classes:

- (a) There seems to be no viable way of justifying norms purely apriori, nor do there seem to be plausible candidates for such norms (*the epistemological problems*).
- (b) We cannot make sense of the nature of the truths that would constitute the objects of the assumed, purely apriori apprehension, nor of the mental act that would be involved in grasping them (*the metaphysical problems*).
- (c) It is unclear that we can even make sense of the need for epistemic norms, or motivate particular proposals about what these norms should be, without assuming a good deal of substantive scientific knowledge (*the methodological problems*).

Of these three classes, (b) is probably the most contentious standing alone, whilst, in a broad sense of 'methodological', (c) is perhaps the most fundamental. Their exact interconnections and presuppositions will become apparent through a detailed presentation of each.

1.3.1 The epistemological problems

TE needs to provide an apriori justification of its apriori norms – so much, surely, is obvious. As we have seen, if one has to appeal to substantive knowledge of the world to justify them, then the norms will no longer function as foundations for substantive knowledge and the traditional conception will fall apart. And if one cannot give a justification for the norms at all, then surely one has no reason to regard them as genuine norms – that is, as genuinely saying what one *ought* to do to gain knowledge – rather than arbitrary stipulations which one could take or leave at will.

However, the test of time suggests strongly that such apriori justification is not generally forthcoming. It would seem first that the idea of an apriori justification for a set of epistemic norms threatens a regress of norms: if the first set is required to justify world-knowledge, and this must in turn be justified, presumably by higher-level norms, then surely we need to ask, of the set of norms at any given level, how that set is to be justified. (Some might see the so-called *analytic* [or, in a narrow sense of the word, *conceptual*] account of the apriori associated with logical positivism and ultimately Hume as an alternative, non-regressive way of grounding apriori norms; this founders for reasons that go somewhat beyond the theme of this essay, and thus I offer only a brief discussion of it in the following note.)¹⁷ It seems the only cogent response to the regress problem is to maintain that the norms at some fundamental level must be somehow self-evident, indubitable, self-demonstrating or something of that ilk.¹⁸ But to this suggestion, there immediately arise at least three good critical questions: What exactly are the criteria for self-evidency, indubitably (or whatever) – and how are these themselves to be non-regressively justified? Have any such proposed criteria ever actually been fulfilled by a concrete proposed norm? Has any proposed norm that might be seen as meeting them ever stood the test of time in relation to what we later have taken to be new scientific discoveries?

Perhaps some would want to insist that logical or even certain mathematical principles satisfy the relevant criteria for self-evidency (or something similar), or at least immediately follow from norms that do so. Take for example *modus ponens*, which says that one should infer a proposition of the form q from two others of the form p and *if* p , *then* q . If one denies q given such premises, then, given the standard truth-table definition of the meaning of 'if... then...' it can be shown that one is led into contradiction – either one must accept p and *not*- p or q and *not*- q . This breaches the so-called *law of non-contradiction*, which says that one should not simultaneously hold a proposition and its negation. This in turn can be viewed as something which cannot be rejected on pain of lapsing into incoherency: if I genuinely and consciously believe any proposition, p , I cannot simultaneously believe, consciously, *not*- p . Given I can entertain beliefs at all, it seems I must abide by the parameters set by the law of contradiction, and hence the other laws of logic.

This justification is, however, far from water-tight or uncontroversial.¹⁹ Moreover, even if one sees no flaw in it, it is a fact that several logical and mathematical principles that have been viewed in the past as utterly indubitable have been brought into question with the advent of quantum mechanics and relativity theory (in particular, the law of non-contradiction and Euclid's geometrical postulates).

Having said all that, it is not to be contested that the epistemological status of logic and mathematics is a deep and vexed philosophical problem – one that I should hardly hope to pronounce definitively on here. Though, for a range of reasons (mostly contained within this and following section of the present chapter), I find attractive the Quinean view that mathematics and logic do not fall into a distinct epistemological category relative to ordinary scientific statements, nothing of what I have to say here would, I am sure, convince a dissenter. However, even if logic and mathematics are in some way special – and here we mean not just ontologically, but epistemologically – this does not make my project to any great extent less interesting, nor undermine the significance of my overarching thesis. To begin with, nearly all epistemologists have felt a need to go beyond logic and mathematics in specifying epistemic rules for enquiry, so my objections to the idea of apriori norms, and the alternative naturalistic construal of norms, can be seen as applying first and foremost to non-logical and non-mathematical ones

(insofar as these can be separated). Secondly, and more importantly, my main remit is to deflate the hopes of NE; assuming that norms quite generally are meant to be naturalistically grounded, what I have to say can then be seen as drawing out consequences of that assumption. In line with this, for most of the rest of this book I will write as if logic and mathematics do not furnish apriori but at best naturalistic norms (in one or other of the senses I will later be outlining).²⁰

Logic, then (along with – by assumption, and something I will no longer flag explicitly – mathematics), may or may not be a special case. However, as already noted, one needs in any case to go beyond logic to effect a workable set of epistemic norms (logical laws do not, after all, tell you anything about what you should believe *per se*; they say rather what you should believe, given you believe certain other things). At the very least, they do not seem enough to answer the kinds of questions that science is typically concerned with, relating to the nature of objects and properties beyond what is given to us in sense perception (or intuitive thought). Accordingly, philosophers of science and epistemologists have sought to specify epistemic norms that go beyond logic. Yet when they have done this, giving a purely apriori justification for them has been anything but straightforward. There would seem to be no such norm whose rejection would, for example, cast one's beliefs into incoherence (as perhaps is the case with logical norms). And to judge by the history of philosophy and science, which bears testimony to a thoroughgoing and almost endemic lack of consensus concerning what the further norms should be, or by what more fundamental standards they should be justified (beyond logic), it would seem there is no norm that deserves the epithet of self-evident, indubitable or self-demonstrating, whatever these exactly might involve.

Most traditional suggestions seem, moreover, to be clearly inadequate in the light of what we now take to be substantive scientific knowledge. Straightforward numerical induction – Hume's famous critique aside – is open to *de facto* countereexamples. If you have seen many things of kind F in many different circumstances and they have all had property G, one might be tempted to believe that all or most Fs are G; but all too often, as in the case of the discovery of black swans in Australia, experience has shown that such inferences break down. Of course, more sophisticated inductive

rules than this have been proposed. However, not only are many of these also inadequate, but, perhaps more significantly, the process of justifying them has come to be understood as a process of refinement in relation to empirical and/or pragmatic factors, rather than as apriori.²¹

Again, Descartes' requirement that one should be able to make clear, intuitive sense of the basic ideas one accepts has also been breached insofar as the theory of gravity – not to mention relativity and quantum theory – seems to defy intuitive intelligibility. The logical positivists' demand that the content of scientific hypotheses be reducible to basic sensational elements renders most of science meaningless, whilst Popper's enjoinder that scientists expose their favoured theories to continual and rigorous testing in the hope of falsifying them consigns the majority of successful science to the category of the epistemically irrational.²²

It will not have escaped attention that the latter objections to proposed apriori norms are empirical in character: the putative norm is evaluated in relation to what we experience, or what we (think we) know from or about science. This might suggest that they have no force against norm-proposals that are meant to be steered by apriori considerations. However, we must bear in mind the overall dialectical situation. No consensus concerning apriori epistemic norms has emerged in the history of philosophy; in addition, most of the concrete suggestions that have been made seem inadequate by any reasonable lights. Therefore, it seems unlikely that apriori justifications for epistemic norms will be forthcoming; indeed, it seems the most tenable means of evaluation of a norm will be by reference to empirical considerations (exactly how this is meant to happen is, of course, something we will have much more to say about).

1.3.2 The metaphysical problems

Assume now, for the sake of argument, that one *had* reached some kind of consensus as to the correct set of apriori norms, and that they therefore could be reasonably classed as self-evident, or whatever (even if we had not quite understood what the relevant notion involved). The supporter of TE still (arguably) needs to tell us something about the nature of the truths that would be seen as constituting (the basis for) these norms. A metaphysical problem could then be posed by the following question: can we countenance the kinds of

truths that would be necessary to play the epistemological role of norms assumed by TE?²³

What is the truth – the proposition or state of affairs – that underlies or constitutes a valid apriori epistemic norm? In a suitably general sense of 'conceptual', we might allow that it should be a *conceptual truth*, that is, a proposition or fact that is independent of the truths that constitute the physical world. But this just begs the question as to what one means by 'a conceptual truth' – and also how we are meant to grasp these. In the history of philosophy prior to the twentieth century we have two dominant models that might be put into service here: the Platonist model, and the theist-mentalistic model of Descartes, Leibniz and other classical rationalists. (I exclude from discussion the Humean view of conceptual truths as 'relations between ideas', which only became prominent in the twentieth century when expressed in terms of logical positivists' *analytic* view of the apriori. Though this raises few metaphysical problems, it is unsuited to the epistemological purposes of TE: see note 17.) According to the former, Platonist model, conceptual truths exist outside space and time as Forms, and we are possessed of an equally immaterial soul that enables us to perceive these. According to the theist-mentalistic model, conceptual truths are ideas in the mind, itself a non-spatial substance, whose veracity is secured by virtue of the existence of a benevolent deity who has placed them there precisely as a foundation for our scientific enquiries.²⁴ However, surely neither of these models is acceptable in our secular era, either as an account of what such truths consist in or of how we gain access to them. Reality is fundamentally the self-contained universe of space-time as conceived in Einstein's theory of relativity. Human beings are just complex biological organisms on the planet earth that have evolved from simpler life-forms largely through processes of natural selection. The metaphysical assumptions of Plato and the classical rationalists are simply unacceptable given this world-view. The same would seem to be true of any other account of the truth-makers of apriori principles. Hence, the very idea of purely apriori truths or norms, at least in the sense relevant to TE, must be rejected.²⁵

A defender of TE might point out at this juncture that this objection is quite blatantly question-begging in its assumption of a naturalistic world-view: a non-naturalistic epistemology can demand, *mutatis mutandis*, a non-naturalistic metaphysics. To this,

the opponent of TE could point out that a broadly naturalistic world-view is widely accepted today and is independently motivated by the success of science. A supporter of TE might of course simply retort: *in relation to what does one evaluate science as so successful?* – intimating the need for precisely apriori norms, and setting up the dialectical interchange anew. But, even if one accepts that there is a logical gap in the anti-TE argumentation here, it is hard to see that many today would want to dig their heels in at this point in support of TE's metaphysics.

Having said that, the argumentation employed here suggests that the metaphysical problems are in some way subordinate to the epistemological ones (whilst these in turn might be seen as fundamentally methodological ones – see below). In a similar vein, one might also argue that the main problem with TE is not that the two models for understanding the nature of apriori truths canvassed above are – merely, or even primarily – metaphysically extravagant, but rather that they are in some deeper way explanatory otiose. This is not again just because they rely on theories – about the Forms, or God – that themselves stand in need of justification, but rather because they explicitly posit an entity with *the express role* of playing the kind of guarantor for our beliefs that their epistemology requires. Thus they can be seen as *reifying* the mere *idea* of a foundation for our beliefs without having any independent reason for positing, or indeed any real conception of, such a foundation or how it functions. In many ways, this is just what modern science presumably does *not* do – reify the desirable. The two models represent both a kind of dogmatism – an insistence on truths that have no obvious connection with the rest of what we believe – and at the same time a lack of faith in our existing cognitive practices and belief-set. Naturalism, building on science, is, as we shall shortly see, precisely concerned to rectify such an attitude.

Some might feel that there is a further metaphysical problem with TE insofar as it countenances *normative* truths, and that one can make no sense of such truths – possibly at all, but at least not given a naturalistic world-view. However, if we allow that apriori truths could be seen to have some special metaphysical status (as Forms, or whatever), it is not clear that seeing these as normative, or as capable of grounding norms, raises any further, substantive problem. To this, one might reply that the assumption of normative truths is really the

most pressing problem for TE. In this connection, it might be pointed out – something a defender of TE could also have done in response to the basic metaphysical objection presented – that, given there is apriori consensus about the norms, then it is simply capitious to demand some substantive account of what these norms consist in. On this way of understanding the debate between TE and its dissenter, the metaphysical problem boils down to the very idea of whether you can countenance normative truths or not.

However, if we look now a bit ahead, it is arguable that the problem of normative truths is also a problem for naturalised epistemology. NE, as I understand it, still wants to set up certain principles as epistemic norms; one could then ask, what part of the world, exactly, are these normative truths or principles meant to reside in? This question might be seen as raising at least two different kinds of doubt: first, as to whether a purely descriptive, naturalistic world-view can find place for normative truths at all; second, about the existence of appropriate descriptive truths corresponding to particular putative naturalistic norms. We will discuss the first problem further in Section 1.4 and the second in Section 3.1, but for now it suffices to point out that, even though TE might have certain metaphysical problems, NE has its own as well – to an extent overlapping with those of TE, and which may even be more serious insofar as they strike at the very coherence of the naturalistic project. This is the main reason I said above that the metaphysical objections to TE are the most contentious, and probably therefore the weakest part of the case against TE (though not, I think, without any independent weight).

1.3.3 The methodological problems

It would appear that all suggestions as to the priori norms by supporters of TE, at least those that go beyond logic, presuppose substantive views about the nature of the world. But this undercuts TE's aim of providing a foundation for world-knowledge: the project is pragmatically self-defeating. For example, if one is seeking to give a normative account of how to get from sensory data to scientific theory, one assumes that our observational powers are limited in relation to the underlying physical reality we seek knowledge of.²⁶ But that is presumably an empirical fact, not something given

apriori. In that case, any norms adduced (partly) on this basis cannot function as an independent, apriori foundation for our knowledge, as they are, according to TE, meant to do.

Generalising this line of thinking, it seems that any epistemic norm can only make sense in relation to the contingent abilities of the creatures for whom it is proposed; one must take account of the kind of being one is seeking to give prescriptions to if those prescriptions are going to be meaningfully applicable to it. This might be seen as resting on the familiar idea from Kant that anything that a being *should* do must be something it *can*, as a matter of fact, do; and thus, conversely, if it is literally incapable of doing something, it cannot be prescribed rationally that it do it. Whether one invokes such a formal principle or not, however, it seems clear there will always be certain empirical presuppositions underlying any suggestion as to appropriate epistemic norms. In relation to epistemology as usually conceived, namely, as concerning *our* belief-forming practices, there will therefore be empirical presuppositions about human cognitive capacities. Thus are the strongly foundationalist ideas behind TE undermined.²⁷ (Note the argument here is *not* that we have *just* these capacities, and that therefore TE's non-naturalistic metaphysic must be wrong, which was one form of objection in the previous subsection; rather it is that TE always assumes or presupposes knowledge about such capacities and thus cannot deliver on its aims of pure apriority.)

A recently much-discussed example of how this constraint yields substantive and perhaps somewhat surprising results is provided by the problem of checking if one's beliefs are logically consistent.²⁸ It is generally accepted as a minimal requirement on a rational agent that she seeks to make her beliefs logically consistent – if they are not, she is in effect committed to any proposition, in particular to (say) *P* and *not-P*, which is self-contradictory. Given this, one might think that a norm enjoining one to *check* that one's beliefs are consistent would be reasonable, and perhaps could be grounded apriori. However, it seems that such a norm could not be countenanced, since we are – though only as a matter of contingent, empirical fact – incapable of checking our beliefs are consistent when these increase beyond a trivially small number. This is because consistency checking amongst a larger number of beliefs leads to (or at least, in the view of the majority of mathematicians who work in this area, most

likely leads to) a combinatorial explosion, with the number of operations increasing exponentially with the number of sentences one is meant to check. We finite human beings could not hope to check the consistency of even so few as one hundred sentences, which is presumably far less than corresponding to the number of beliefs most people hold. Thus it seems that the proposed norm is not after all acceptable. But that is because we tacitly accept the above constraint that we need to take account of empirical facts about our abilities in giving a theory of epistemic norms, thus breaking with the precepts of TE.

In these and similar ways TE would appear to be fundamentally misconceived as a means of specifying what it aims to specify: how rationally to form beliefs. On the line of thinking currently being investigated, it is not most fundamental that there are substantive epistemological and/or metaphysical obstacles in the way of an aprioristic way of doing this, but rather that apriorism is self-stultifying – hence '*methodological problems*'. One might also class the epistemological problems, and thus the metaphysical ones, as methodological too, albeit in a somewhat broader sense. Thus, in the previous subsections, we have seen that it is essentially TE's failings as a viable epistemological *research programme*, as one might put it,²⁹ that is fundamentally unsatisfactory about it: in the history of the subject, it has simply failed to deliver the goods, suggesting we should try something new. The objections of the present subsection underline this conclusion in a more principled manner.

1.4 Naturalism

The above critique of TE alludes, at various points and various ways, to ideas associated with the philosophical view that has come to be known as *naturalism*. This is perhaps most evident in the subsection on the metaphysical objections which made use of the idea of a naturalistic world-view. But another dimension of it, which I believe is more fundamental, has played a pervasive role: the idea that we can only reasonably start our enquiries *in situ* – assuming some ordinary, empirical knowledge. In this final section of the present chapter, I want to outline and defend the basic tenets of my understanding of naturalism, and explain how it connects to NE.

I understand naturalism as committed to roughly the following two ideas:³⁰

N1: In anything we might call a justification, or justified formation, of a belief about the world, we cannot ultimately dig deeper than to other beliefs about the world and to basic cognitive practices.

Justification is always in this way situated.

N2: The fundamental cognitive practice for us today is that of modern science, which operates with a *naturalistic picture of the universe*: one that does not trade in immaterial, magical or supernatural entities or explanations.

This understanding of naturalism owes a lot to the work of Quine. N1 draws on his holistic conception of science as an interconnected and ever-evolving *web of belief*, as first presented in his classic paper 'Two dogmas of empiricism' from 1953; N2 draws on his predilection for the ontology of natural science (though I should indicate immediately that this does not for me, as it does for Quine, indicate a commitment to any strong form of physicalism).³¹ On the other hand, it is important to note that N1 and N2 do not, as they stand, incorporate the main thesis of Quine's 1969 article, 'Epistemology naturalised' – that normative epistemology should be ditched in favour of a descriptive account of how knowledge is constructed on the basis of sensory evidence. If we accept Quine's empiricism as a further premise, I think this latter view naturally meshes with the web of belief picture, though only on a certain understanding of 'normative epistemology'. In the following, I will explain in more detail what N1 and N2 commit one to, spelling out some of the connections to Quine, before turning to the relation between naturalism and NE.

Consider first the idea in N1 of justification of belief as 'situated' *vis à vis* our actual beliefs and cognitive practices – as presupposing at least some of these. As I see it, this essentially simply builds on the failure of the idea of the apriori charted in the previous section: there is no ultimate court of appeal for assessing the veracity or, more broadly, rationality of our beliefs (indeed, in attempting to find it, one invariably assumes what one is supposed to be grounding). In asking whether one should believe a certain proposition, one is faced, in the first instance, simply with more beliefs that have

ultimately the same epistemological status as it, together with the basic cognitive practices that we bring to bear in manipulating beliefs (more on these very shortly). If the belief seems to follow from (or be ruled out by) beliefs one already holds dear when there are strong empirical reasons for rejecting (or accepting) it, there is no principled restriction against rejecting one or more of the beliefs that recommend it in order to allow the empirical reasons to hold sway. Anything might in principle be revised in relation to new evidence.³² (The opposite situation is also possible, of course – that empirical reasons can be overridden for the sake of preserving more 'central' or 'theoretical' beliefs – something that should make clear that 'empirical reasons', understood in terms of what one is supposed to be able to see, hear, etc., are themselves also anchored in the web of belief.)

Given just this idea of epistemic situatedness, one might think that beliefs could be justified by entering into so-called 'local holisms': sets of beliefs that mutually support each other but do not encompass the whole system of belief. However, though we seem able to make clear intuitive sense of such local holisms, the web-of-belief model takes it that it is some kind of ideal that beliefs should ultimately cohere with *all* other beliefs. In the final analysis, it is the one, natural world we are seeking to understand, whose different parts and aspects presumably hang together somehow (note again that I do not see this idea as committing one to a radical physicalism – see below). It is also the case, as stressed by the greatest influence on Quine's holistic model, Pierre Duhem, that in processes of scientific reasoning (indeed, we might add, in reasoning generally), though one might initially be concerned with some local problem, there are always potential ramifications to an ever wider set of belief – as Jerry Fodor has put it, '[i]n principle our botany constrains our astronomy'.³³ And ultimately, once one has given up the hope of finding apriori foundations, the only meaningful sense in which one might ground a belief would seem to be to bring it into a coherent epistemic relation with an ever-greater proportion of one's total set of beliefs. For all these reasons, the idea of an *all-encompassing* web of belief is appropriate as a metaphor for naturalism. (This should also make it clear that we are not talking, first and foremost, about the belief-system of any individual, but rather the amassed knowledge of the scientific community – *contra* what Fodor tends to assume in his discussions of scientific reasoning.)

So much is essentially just Quine; by contrast, the notion of *basic cognitive practices* in N1 is not something that derives from Quine's philosophy. By stressing it, I mean to indicate that, in addition to beliefs, we must presuppose, albeit perhaps only in the very broadest sense, *ways* we have of manipulating these beliefs, of relating them to one another – ways which are not (in the first instance) encoded as explicit beliefs or rules, but simply exist as certain reasoning dispositions that we must take to be fundamentally valid (though not of course infallible, in the sense of *always* leading us to the right beliefs). Reasoning is, after all, a process, and if, as naturalists, we must take certain beliefs we reason from for granted, then we must *mutatis mutandis* take for granted the means by which we reason to and from them. Naturalism in itself leaves it open whether these means or ways embody any kind of *system* beyond that which might supervene on the contents of the beliefs they relate. Further, talk of basic cognitive practices does not necessarily imply that we are talking about *psychologically* or even *individually* construed competences – any more than talk of the holistically interconnected web-of-belief should be understood primarily in relation to the belief-systems of individuals. The idea of basic cognitive practices, under its various different guises, will be pivotal in understanding the different varieties of naturalised epistemology to be discussed in this book.

In the last-paragraph-but-one, I made use of the idea of specifically scientific reasoning and of a scientific (or naturalistic) world-view, which connects with the second part of naturalism as I understand it: N2. Now, I believe there are sufficient resources to motivate a web-of-belief model, of the kind described so far, without substantive assumptions of this kind; however, the resulting position would arguably not be one worthy of the epithet 'naturalism'. The basic idea that we must start within our web of the belief and with our actual cognitive practices is reminiscent of some aspects of Wittgenstein's later philosophy, in particular his so-called *quietism*, to the effect that there is no coherent sense in which we can ground our basic beliefs, nor a perspective from which we can meaningfully criticise them. I think this quietist notion is central to naturalism, but it is not enough. The scientific belief-system is, arguably, not the only system of belief; others revolving around ideas about superhuman gods, transcendent virgins, star signs or whatever might be in some

Traditional Epistemology and Naturalism 27
 Put the above way for acceptance of naturalism
 sense capable of elaboration into a distinct and coherent network.³⁴

The web-of-belief model, to be fully naturalistic, needs to be supplemented with a commitment to a scientific world-view. As it happens, it would seem that the scientific world-view is the product of an enterprise which, as noted above, seems to use a holistic form of reasoning; at the same time, this kind of reasoning would seem requisite to understanding the causally interconnected universe. So, I see the anti-foundational and metaphysical sides of my naturalism as intimately interwoven. Nevertheless, neither is sufficient on its own to characterise naturalism (remember that a scientific world-view is also shared by people who would reject the web of belief model, such as logical positivists).

It is important not to overestimate or misunderstand the use I am making of the scientific world-view in characterising naturalism. Here, there are three points to make. To begin with, I do not mean to be committed to the doctrine of *physicalism*, the idea that all properties (phenomena, entities) are, should or can ultimately be, reduced to physical ones. For many, and particularly for the majority of contemporary philosophers of mind, 'naturalism' denotes a view of the world as containing only properties that can be cashed out in terms of the concepts of fundamental physics.³⁵ I join the ranks of those who find such a goal inherently vague,³⁶ and at least as nowhere expressed and hardly ever realised in actual science.³⁷ It is thus unmotivated from a naturalistic point of view – indeed, it stands in danger of becoming precisely a non-naturalistic, apriori dogma. Contrary to what many of its supporters seem to suppose, it also lacks motivation by way of general metaphysical considerations concerning the nature of intentional psychology and other 'higher-level' explanatory schemes.³⁸ Though nothing of great significance hinges on these assumptions in the present work, it is at least important to establish a clear demarcation between my own view and Quine's predilection for behaviouristic and neurophysiological explanations in psychology – and his concomitant rejection of notions like *tacit knowledge*, which I will be making free use of later on.

The second point is that my commitment to the scientific world-view is not really meant to be substantive in the sense of something that, at least potentially, might stand in need of further backing. In deferring to science, and thus to a naturalistic world-view, one builds on the beliefs and cognitive practices that *de facto* prevail in our

modern, secular society, and that (I assume) most reasonable people, if they thought long and hard enough about it, would feel compelled to endorse. This attitude might be backed up with the idea that science has facilitated our happiness and survival in a way no other belief-system has or seems ever capable of doing. However, I do not think it is necessary to see this connection as necessary or even sufficient to explain our faith in science (perhaps one might be sceptical of the connection between science and progress; perhaps one might doubt the inference from progress to truth). Nor need one who forewent such an explanation appeal to some kind of religious conviction as a basis of their scientific faith (in the manner famously associated with Einstein). Ultimately, it is part of naturalism, as I understand it, that we are to an extent simply beholden to whatever system of beliefs prevail (at a given time), which in our case is that which science offers. Connectedly, I do not offer any definition of science, except, again, to stress that the naturalism of its world-view is intimately related to the epistemic stance which stresses coherence with the belief-system as a whole, and thus abjures the dogmatism that is traditionally associated with apriori projects (see the discussion of the metaphysical problems in the previous section).

Finally, I should also say a word about the relationship between the world of everyday belief and reasoning, on the one hand, and science, on the other, since this will play a significant role in later discussions. My underlying conception of the relationship between common sense and science is that the latter builds on, refines, supplements and to an extent corrects the former, rather than completely overhauling it. (This applies to both at the level of world-view and of implicit cognitive practices.) This conception fits, as I see it, with the idea of the web of belief as an evolving, holistic network: common sense gives us only an unsystematic and underspecified realm of cognitive activity which what we call scientific thought offers our (surely) best development of; it is not a self-standing realm which science competes with. I think this a reasonable view, and will assume it in the following, though I do not have space to argue for it here.³⁹ (On the other hand, I do not think that anything essential in the argument of this book would be undermined by an alternative view of the relationship between common sense and science.)

This, then, is naturalism as I understand it (and is what I mean to refer to whenever I unqualifiedly talk of 'naturalism' in the following).

It is important, I stress again, to see it as essentially a reaction to the failures of TE. It does not see itself as a competing first philosophy, an alternative to the various metaphysical and epistemological theories put forward by philosophers throughout the ages in the spirit of TE. It builds rather on what is a reasonable scepticism to any such theory in view of the dismal track record of trying to find one, as well as on the project's internal tensions. It also relies to an extent on the secular world-view and methods of contemporary science that are widely accepted in the world today. These three ideas, though not totally decisive against TE, lend each other mutual support and make what I am calling naturalism the obvious starting point today for any further epistemological project. In itself, however, naturalism represents essentially a position of neutrality – a form of quietism, though one that accepts and records the fact that we live in a scientific age.

If that is naturalism, what is NE? For Quine, it would seem to be just the combination of naturalism with his empiricist view that knowledge is something acquired through the use of the senses and the subsequent elaboration of the evidence thereby delivered. If we allow that this idea and all the complex working out of what it involves are facts, and if we also, in accord with naturalism, view these facts as part of the web of belief, then Quine's idea of NE would seem to be a natural fall out: we replace the traditional, normative, apriori project of, as he puts it, 'deducing science from sense data'⁴⁰ with the purely descriptive, scientific project of detailing the course of knowledge construction from sensation 'upwards'.

Of course, many would balk at the idea that it is the senses that give us our basic cognitive access to the world (or that they should play such a privileged role in this connection). Indeed, this aspect of Quine's view would seem to amount to little more than an apriori empiricist prejudice on his part, and thus to fit in badly with the naturalism that one otherwise sees him as advocating.⁴¹ A related but more significant point is that 'deducing science from sense data' is hardly the most obvious model for understanding the kind of apriori project that TE saw itself as engaged in, let alone the only such model. This suggests that, even when once one has renounced TE and its apriorism – thus embracing naturalism in my sense – there may be many other possibilities for motivating a normative epistemology that go beyond what Quine is concerned to attack in 'Epistemology

naturalised'. Indeed, Quine himself has later precisified that his aim was not to renounce the normative.⁴²

In this book, I will be precisely considering different conceptions of a theory of genuine⁴³ epistemic norms that takes as its backdrop the naturalism I have been describing in this section. It is (first and foremost) such theories that I see as propounding the idea of *NE*, and will talk of as different forms of *NE*, or simply different *NEs*.⁴⁴ *NE* as I understand it is thus, in contradistinction to Quine's original usage, essentially *normative* in nature.⁴⁵ The idea that naturalism should open for such possibilities may or may not be immediately evident; that it does do this, or at least appears to do so, will be the point of Chapter 2 (in large part) to demonstrate – notwithstanding that I will ultimately be arguing that none of the possibilities adumbrated are capable of delivering genuine norms that are also *necessary* for optimal belief-formation.

At this point, some readers may nevertheless feel that the naturalism I have outlined – and that arguably sounds the death-knell of at least one brand of normative epistemology – must be too radical a position to lend even hope to constructing *any* normative epistemology on its basis. As we shall see in this book, this reaction is refuted by the substantive attempts that have been made to motivate normative epistemologies on austere naturalistic grounds (including one from Quine himself); in addition, there are principled reasons why one might hope to offer a normative epistemology even once naturalism is embraced. Even so, it is arguable that many who propose normative epistemologies and call themselves naturalists (and thus naturalised epistemologists) cleave to a less committing kind of naturalism than that I have outlined here. Thus, a very common strategy in recent epistemology has been to view empirical, scientific knowledge, particularly concerning the ways our cognitive system functions – deriving from studies of perception, memory, reasoning and so on – as a necessary *constraint* on a theory of epistemic norms, that is, of how we *should* reason, but not to see this knowledge as, even potentially, *exhausting* such a theory. The guiding thought here appears to be that even if a meaningful epistemological theory must take its starting point in how we in fact are cognitively built, it must also go beyond these facts. Thus, there is still a special role for the work of philosophers in epistemology. This is often seen

as connecting to the fact that epistemology is by its very nature a normative discipline, and that purely descriptive information will thus always be insufficient as its basis. We may need to naturalise epistemology, but to go the whole hog with the Quine of 'Epistemology naturalised' is epistemological suicide; as Kim puts it: 'For epistemology to go out of the business of justification is for it to go out of business.'⁴⁶

I agree with Kim and others that an epistemology in business is in normative business. However, I disagree that my naturalism obviously rules out normativity, as I will shortly explain. Moreover, I think there has been a deal of uncertainty from supporters of naturalised epistemology who stress the role of philosophy as to exactly *what* this role is meant to amount to. Perhaps all that is meant is that we need philosophers to give theories of epistemic norms – in which case the theories would presumably be fully naturalistic in my sense and susceptible to the criticisms to be presented in this book. In some cases, however, it seems clear that the reference to philosophy is meant to be more substantive in that the theory makes play with the kind of aprioristic ideas I have been attacking in this chapter.⁴⁷

I have not completely ruled out that some logico-mathematical principles might be justifiable apriori, even though they would not furnish a complete epistemological theory. This might suggest there is scope for a more complete theory of norms to be secured through a rapprochement between apriori and naturalistic considerations. Now, this book is concerned first and foremost with *NE* in a strong sense, that is a conception based on the Quinean naturalism I have outlined above, which eschews the notion of the apriori entirely. There are also, as I said, several extant varieties of *NE* conceived along these lines, as well as principled reasons for thinking such varieties should exist. In view of this, a legitimate strategy for me in response to the possibility of 'weaker' or more moderate naturalistic epistemologies would be simply to disregard them. And indeed, I will not devote extensive space to an examination of the different ways they might be developed and defended. Nevertheless, in Section 5.1, I will consider a few of varieties of position which would see norms as a joint product of apriori and empirical sources of information, which I also believe are representative of a wider range of similar views. I also believe, as I will seek to show, that they only establish a need for norms that

correspond to the apriori sources. With the exception of this section, I will assume a sense of NE as based on (my) naturalism, and not any weaker versions thereof, for the rest of the book.

This chapter is almost at an end. In closing, I will consider the two most famous objections that have been levelled against NE mainly for the purpose of putting them behind us.

The first is that NE is circular. Traditionalists are accused of landing in a regress of justification, but when their opponents embrace naturalism, the tables are turned: NE assumes scientific knowledge and the validity of existing cognitive practices as the basis for any further normative evaluation of science and cognitive practices. In this way, one is simply assuming what one's project – that of deriving valid norms for science – is meant to serve to ground.

Though there is much to be said about this kind of objection, it is widely accepted that it is – at least when pitched at such a general level – wide of the mark. A proposal as to a certain norm should not, of course, assume precisely what it is meant to ground, but it can presumably assume other items of knowledge. If we accept naturalism, this latter kind of circularity cannot be seen as vicious, for naturalism is precisely the idea that our actual beliefs and practices are our only possible starting point. In the light of the failures of TE, we are left with no option.⁴⁸

The second objection was alluded to above in discussing weaker versions of naturalism and NE; I shall call it 'the problem of normativity'. It originates in Hume's attack on cognitivism in ethics, and is also connected to what G.E. Moore dubbed 'the naturalistic fallacy'. In its crudest form, it points out that one cannot derive an 'ought' from an 'is', that is, that one cannot derive valuational or normative statements from purely factual or descriptive ones. Hence NE, based on the radical kind of scientific naturalism I am propounding, might seem bankrupt from the start. The naturalistic world-view has simply no place for 'normative facts', such as that a belief might be *justified* in the light of some other beliefs.

I think a supporter of NE should, as with the objection from circularity, simply demur at giving a principled answer to the problem of normativity. All naturalists, in rejecting TE, reject the idea apriori, foundational norms. But they do not thereby reject the idea

of *normativity* – the idea that there can be, if only in particular circumstances, *reason* to believe one thing rather than another; that in these circumstances, one *ought* to believe the one thing and *ought not* to believe the other. Naturalism, as I understand it, involves the idea that we must assume our basic cognitive practices are fundamentally in order. This means that we must *assume* we reason (for the most part) as we ought to, not that we *actually* reason as we ought to; but at least the idea of *correct reasoning* is not something that is abnegated at the outset. In that sense, the problem for NE is perhaps not so much that it fallaciously *derives* the normative from the descriptive (at least in general); the problem is that there may not exist anything like real normativity at all. However, since to give up this idea appears to invite intellectual suicide (how could one recommend doing precisely that, for example?), one presumably should not – if only on methodological grounds – see rejecting normativity as an automatic consequence of embracing naturalism. As we shall see in subsequent chapters, it is also intuitively plausible that certain descriptive facts can suggest that certain specifiable norms should be accepted. All in all, this should be sufficient for NE to get off the ground.⁴⁹ (Note that this general answer to the normativity problem does not automatically discharge one of having to say what naturalistic state of affairs *particular* norms consist in – something I will suggest *anti-psychologistic* versions of NE precisely have problems doing [see Section 3.1]. Nor does it automatically enable one to regard principles of a natural reasoning competence as rational norms, as *psychologistic* versions of NE seek to do – see Section 5.3.)

With these objections behind us, the challenge that now remains for NE is that of showing what form it should take, and then of validating the coherence and necessity of the epistemic norms it proposes. In the following chapter, I will outline what I see as the space of possibilities for NE in the light of the work of two highly influential thinkers in contemporary epistemology. The rest of the book is devoted to criticising these various possibilities in accord with my overarching thesis.

2 Naturalised Epistemology

Introduction

The aim of this chapter is to motivate a taxonomy of different varieties of NE and outline their essential characteristics. My intention is that this taxonomy be, in all essential respects, exhaustive of the possibilities for NE. The different varieties – which I see as dividing again into several different sub-varieties – will then be critically discussed in the subsequent chapters in support of my overarching thesis.

As a way into the distinctions relevant to this taxonomy, I will present the ideas of two philosophers who, one directly and the other more indirectly, have had a large influence on debates about naturalism in epistemology. The first of these is someone we have already said a good deal about: Quine. The other is Nelson Goodman, and in particular his work on the justification of rules of inference.

2.1 Quine, psychologism and anti-foundationalism

We have already discussed the main thesis of Quine's article 'Epistemology naturalised', but it will be instructive to review the arguments of it, and its relation to other aspects of Quine's philosophy, in a little more detail as a point of entry to my presentation of NE. Much of the article is concerned with charting two distinct phases in the demise of what Quine calls 'traditional epistemology', corresponding relatively closely to our usage of that expression in Chapter 1. The first is Hume's attack on the possibility of sound

inductive inference on the basis of observation; the second is Quine's own attack on Carnap's semantical reduction of theoretical terms to observational terms, which perhaps turns most fundamentally on the idea that scientific statements 'face the tribunal of experience not individually but only as a corporate body'.¹ (These points thus overlap with those adumbrated in Chapter 1.) Quine concludes that there are no sound normative principles on the basis of which we may understand the relationship between observation and theory, and that we should content ourselves with understanding the actual, descriptive course of events that leads from sensation to belief. He dubs the discipline that results from this NE, though it is unclear whether it really deserves the epithet 'epistemology' at all. As we have already noted, what Quine really seems to be advocating is that we *ditch* epistemology – arguably something essentially normative – and *replace* it with a particular form of descriptive psychology.

As we also noted in Chapter 1, Quine has since insisted that his view is not nihilistic in relation to matters normative; moreover, it seems this commitment is evident even in 'Epistemology naturalised'. As a metaphor for science, we find here Quine extolling the virtues of Otto Neurath's boat, afloat at sea and far from any port² – a metaphor Quine is particularly fond of, and which lies close to his own notion of science as a web of belief. The idea behind the boat metaphor is that as a mariner must repair planks of the boat while standing on other planks, unsure as to whether they also stand in need of repair, so must we assess and correct our scientific beliefs *in situ*, each being in principle as fallible as any other. The present point is just that this account seems to open for, if not to advocate, significant normative evaluation of our beliefs – at some basic level, we continue to *assess* and *correct* them in relation to one another. It is also relevant here that the web of belief is not conceived, first and foremost, as the belief-system of an individual, but rather the amassed knowledge of the scientific community as a whole, such that individuals operating within this may be more or less apt at respecting and exploiting its holistic interconnections. In view of all this, it would seem that normativity is firmly ensconced in Quine's metaphors for science. On the other hand, it is significant that these do not suggest endorsement of a *substantive* epistemological theory in the sense of something delivering norms (i.e. what I would term, an NE.)³

In relation, however, to the task immediately at hand (i.e. motivating a taxonomy of NE), the point of direct importance is that though Quine himself appears only to see a seamless development of a single, core intuition relating his various ideas, his stress on psychology and his stress on the metaphor of the boat/web of belief may be seen as two quite distinct ways of reacting to TE. The former emphasises the centrality of an *a posteriori* examination of the way our sensory systems, brains and/or cognitive systems are actually organised in any viable descendent of TE – something that has since been taken up by many who do not accept Quine's strongly empiricist and physicalist views on the nature and genesis of knowledge. The stress on the boat/web-of-belief metaphor enunciates by contrast just a move away from TE's foundationalism – from the idea that we can ground our beliefs about the world by reference to something that is justificationaly prior to them. As noted in Section 1.4, Quine's idea of a scientific replacement of traditional empiricist epistemology flows naturally from the web-of-belief model together with his background empiricist assumptions. But as we have also noted, Quine's empiricism is far from mandatory. Moreover, starting with just the web-of-belief model of science, there would seem to be scope for other, substantively epistemological (i.e. normative) projects; as we shall see later on in this chapter, the same is true given just the emphasis on the centrality of scientific psychology, at least when psychology is understood in a suitably wide sense.

These are, then, two distinct ways of retracting from TE that can be recovered from Quine's writings; as we shortly shall see, they play a role in characterising my taxonomy of positions within NE. To the same end, we should now also note that we may distinguish two distinct senses in which one could conceive of the anti-foundationalist metaphor of repairing the boat at sea. Given my definition of naturalism, it follows that all putative varieties of NE must give up on the strong foundationalism of TE – give up the aprioristic idea that, as we may put it, we can and should repair our boat by docking at some port and standing on completely firm, dry land. However, there are two ways one might retract from this, one a more radical departure from TE than the other. The more radical retraction might be captured by the metaphor of continually moving around to many different parts of the boat to repair others, without ever finding some particular plank or planks that are definitively the most sturdy to

stand upon; just as science evolves indefinitely, progressing in some sense, but without arriving at any end-point. This retraction might be dubbed *strong* anti-foundationalism, but, for reasons that will soon become apparent, I will for the most part call it simply *anti-foundationalism*.

The less extreme reaction may be envisaged as eventually deciding on some particular part or parts of the boat as especially suited to standing on whilst repairing all the other parts; one could not have found this out apriori, that is, without first trying out many other parts, but one could, presumably, come justifiably to believe that certain areas of the boat are particularly suited to standing on in order to repair all the other parts – even though they will not be 'infallible', and we can revise our opinion about what they are. We might call this idea *weak* anti-foundationalism – or, as I shall say, *anti-apriorism*, since, whilst denying that we can do epistemology apriori, it holds that we can still make sense of some of our beliefs as being justificationaly prior to others (and thus even – in some suitably weak and diluted sense – foundational for them). (Exactly what I have in mind by anti-apriorism will become clearer in Section 2.2 and Chapter 3.)

Now we are in the business of putting labels on things, let *psychologism* be the view that we should replace TE with a study of our psychological belief-forming mechanisms, leaving open how we conceive of these mechanisms and whether or not such a study might have normative implications.⁴ Quine's writings on naturalism and NE open for all these ideas. As I read his earlier papers, he would seem actually to espouse both (strong) anti-foundationalism and a certain form of (non-normative) psychologism. At the same time, he seems to have no interest in propounding a substantive epistemological theory delivering norms, that is, an NE.⁵

This position, modified in the direction of a less behaviouristic/empiricist psychology, lies very close to my own view. What we need to do right now, though, is to consider how supporters of NE could exploit the ideas we have uncovered in Quine to motivate their different epistemological theories. To start with, NE might embrace (strong) *anti-foundationalism*, and see the job of positing norms as a project which interacts with ongoing scientific practice in the most general sense, and as having no more (clearly defined) end-point than that practice. On this line, we can propose certain norms *pro tem*, but there is no prospect of ever definitively deciding on some

suitable subclass of the totality of our beliefs as *the* norms for belief-formation. I call this *anti-foundational NE* (or *AF*).

Alternatively, NE could adopt *anti-apriorism*: some parts of the boat can, after some period of empirical enquiry, be identified (albeit fallibly) as the logical (i.e. not merely temporal) basis for our epistemic activities. Epistemology turns out, if you like, to have some specific object of study – alongside that of physics, chemistry, biology and so on – which it is the job of the epistemologist to uncover; it is just that for the naturalised epistemologist this object is not conceived of as some realm of apriori truths. Within this anti-aprioristic camp one can place a *psychologistic NE* (or *P*, as I shall refer to it), which seeks to give a theory of norms through detailing the structure of our actual belief-forming mechanisms (where these could perhaps be viewed as tacitly known principles in a natural reasoning competence – more on which below). On the other hand, there might also be those who, while wishing to preserve the idea of a unique and, in some (weak) sense, foundational set of naturalistic norms, would reject the idea of seeing this as concerning the human cognitive system. We might call such a view *anti-psychologistic NE*, or just *anti-psychologism* (henceforth *AP* – a position I will say more about in Section 2.2).

Noting that *AF* is *ipso facto* an anti-psychologistic form of *NE* (since *P* is a *weakly* anti-foundational form of epistemology), we can classify these three positions in relation to the distinctions we have already drawn as follows:

Varieties of NE			
	Anti-apriori		Anti-foundational
Psychologistic	P		–
Anti-psychologistic	AP		AF

As we shall see, each of these broad categories accommodates several different sub-varieties. It is also possible to propose some kind of 'mixed' view, which combines elements of two or more of the three main types (as well, perhaps, as philosophical or apriori elements); we shall consider some views of this kind in Section 5.1.

2.2 Goodman and the method of reflective equilibrium

To help begin to put some flesh on these different theories, I believe it is instructive to examine the line on the justification of epistemic norms defended by Nelson Goodman, something that has had a deep influence on naturalising moves in epistemology. In Goodman's famous discussion of induction in *Fact, Fiction and Forecast*,⁶ the issue is not, as for Quine, whether normative inductive principles are any good. We are reliant on some or other form of induction, so we must simply assume that some or maybe several different versions of the principle are valid. The burning question is how we *justify* such principles, in a way that can quell Humean sceptical doubts. A tacit assumption of *TE* is that we can have some fairly direct, intuitive contact with the appropriate epistemic norms, or can deduce them from truths that we have such contact with. Goodman is sceptical to this line, both in relation to inductive and deductive principles. Instead, epistemic norms (generally, as I will henceforth assume) should be arrived at and justified by reference to judgements and intuitions about what is a valid inference in concrete cases. If a proposed norm comports well with our judgements about concrete cases, then it is insofar justified. However, it is not the case that a norm must be immediately rejected simply because it conflicts with one or more judgements about particular cases. These latter judgements do not form a totally incorrigible and infallible bedrock for epistemological reflection. We may need to revise one or more of them, especially if a proposed principle accords well with many other such judgements – in somewhat the same way as the existence of anomalies for an otherwise highly explanatory and well-confirmed theory in science may reasonably lead to rejecting or reconceptualising these in ways conducive to the theory. The aim of the enterprise is to arrive, by gradual mutual adjustment of norms to judgements, at a set of coherent epistemic norms which accord with the judgements we are disposed to make about particular cases – again, at something like a *theory* of these judgements, where the hypotheses constituting it make claims about, or simply yield, the correct norms, which may then be viewed as justified insofar as they are arrived at in the manner described.

This method has also been advocated by John Rawls in relation to the justification of ethical principles.⁷ Rawls termed it the *method of*

reflective equilibrium, since it is precisely a *reflective, rational* accord between rules and judgements we aim to arrive at, though this need never be achieved to perfection. Though we will only be concerned with *epistemic* norms here, we shall (following recent practice) refer to Goodman's method for finding and justifying epistemic norms as the method of reflective equilibrium, or *MRE*.

As far as I am aware, Goodman himself did not see MRE as relating to NE, at least as such (his original discussion predates Quine's 'Epistemology naturalised' by over a decade). That something like MRE would function as a least central component in justifying epistemic norms is a widely held view amongst many calling themselves naturalists today (both those who cleave to the strong variety of naturalism I am assuming and others). Nevertheless, these naturalists usually see the process that would eventuate in a *naturalised* epistemology as issuing from a process of so-called *wide* reflective equilibrium, as opposed to the 'narrow' variety proposed by Goodman. The expression 'wide reflective equilibrium' also stems from Rawls, and was used by him to denote a process whereby one sought to bring intuitions about concrete cases and rules in accord with wider philosophical theories, for example from ethics or metaphysics.⁸ In co-opting the idea of wide reflective equilibrium, naturalists stress the need to take account of *empirical* evidence concerning the structure of our reasoning capacities and other psychological data in the reflective process of deriving valid epistemic norms (together with – for a *moderate* naturalist – the kind of philosophical theories Rawls originally stressed; I discuss an extant example of such a position in Section 5.1.)⁹

Notwithstanding all this, I wish to maintain that MRE, even in narrow, unmodified form, already enunciates a naturalistic perspective on epistemology – or at least does so when interpreted in a manner that makes it plausible as a way of justifying norms. This interpretation of MRE tends to be obscured, I believe, because the intuitions about rules or principles are themselves often seen as having the role of *data*, along with the judgements about concrete cases, both of which are then fed into and steer the process towards reflective equilibrium. On this understanding, it becomes natural to regard the intuitions about rules as corresponding to philosophical or *a priori* elements, and the whole process as differing from traditional accounts of justification insofar as these are fallible and

susceptible to correction by taking into account judgements about concrete cases as well (also viewed as fallible). The problem with this interpretation is that it makes MRE susceptible to a seemingly cogent objection: if MRE builds on just fallible intuitions, how can a state of reflective equilibrium that has taken account of these intuitions be guaranteed to correspond to the correct set of norms? How can the process reliably bootstrap itself up to correctness from half-correctness?¹⁰ A related problem is that the process has no inherent systematicity; we aim at reflective equilibrium between rules and judgements, but our progress towards this can only, it seems, be desultory.

The right response to these objections, I believe, is to take seriously an analogy I hinted at in first presenting MRE: the idea that the whole process is in fact part of an ongoing, *theoretical* enterprise in which one always proceeds *in situ*, seeking to improve upon the *theory* – about the norms – that one has at any given time by taking account of further *data* – that is, judgements about concrete cases, made in accord with our basic cognitive practices – and modifying the theory on the basis of these, though in a manner, no less scientific, that also permits rejection of the concrete judgements, rather than the theoretical tenets. What makes this not merely a good but also a good *naturalistic* response is precisely that it eschews completely the idea of an *a priori* starting point – and of any external vantage point – for the justification of epistemic norms, fallible or otherwise. In this way, it is consonant with Quine's idea of science as an ever-evolving web of belief (or a boat afloat at sea), which for us lies at the heart of the naturalistic approach to knowledge.

It should also be noted that exactly how such theory-construction proceeds or should proceed is in itself something that would be revealed through MRE.¹¹ This of course introduces a circularity into the justification, but it does not impugn the use of MRE, for, by naturalism, we must in any case assume the soundness of our basic cognitive practices in manipulating the web of belief. In short, MRE can be seen as a way of trying to make these practices explicit by using them – with the aim of systematising and subsequently expediting or streamlining them.

Thus, I see MRE as fundamentally naturalistic. At the same time, in being a theoretical enterprise, the basic procedure is open to modifications in order to take account of 'wider' naturalistic considerations,

that is, empirical data from psychology (see above). Moreover, I believe such a naturalistic understanding of MRE – construed narrowly or widely – can be finessed in relation to all of the three main varieties of NE outlined in the table above.

Though I see Goodman's intentions as perhaps cohering most closely with AP (as we shall see), I have structured the ensuing discussion by starting with the link between MRE and P (i.e. psychological NE). The connection here can best be approached through the vexed question in the philosophy of psychology of whether it is possible to show that humans are irrational by experimental means. This has arisen in the wake of various psychological experiments that seem to suggest that (even) normal, intelligent and well-educated human beings perform rather 'badly' on certain kinds of reasoning task – badly, that is, in relation to certain traditionally accepted standards of reasoning, in particular, logical and statistical rules of inference. For example, they ignore base rates in judging conditional probabilities, in violation of Bayes' theorem, and commit fallacies of deductive reasoning, such as 'affirming the consequent'.¹² Whilst many have taken these results to indicate that humans are naturally irrational, certain philosophers – most notably Jonathan L. Cohen – have argued that such factors must be due to *performance* interference of an underlying body of knowledge, or *cognitive competence* that is fully in accord with the correct rules of inference.¹³ This is because there is a *constitutive* link between the concrete intuitions of ordinary people and what the correct rules of inference are, that is the latter cannot diverge significantly from the former.

To explain and support these claims, Cohen uses the idea of MRE, as well as ideas from cognitive science. Building on Goodman's writings, Cohen claims that normative standards of inference are determined in relation to judgements about concrete cases, that is, by MRE. But it also appears that MRE bears striking resemblances to the method used in various areas of cognitive science to determine the nature of underlying cognitive competences. The paradigm example of such a cognitive competence is our language faculty as understood by Noam Chomsky and his followers,¹⁴ though in recent years, especially under the influence of evolutionary thinking, the notion has been extended to several other domains, such as our rudimentary knowledge of the behaviour of physical objects (so-called *common-sense physics*) and belief-desire psychology.¹⁵

Chomskyan linguistics aims at adducing grammatical principles which may be seen as corresponding to a natural cognitive competence, that is a body of knowledge that is in essential respects the product of an innately specified cognitive structure. This knowledge is drawn on in understanding and producing linguistic material, but the actual utterances and linguistic judgements of a given individual may from time to time involve deviations from the knowledge, or competence – issuing in the so-called *performance errors*. For example, many people instinctively regard sentence (1) below as ungrammatical and would never utter it. However, on reflection it is clear that it is built up from sentence (2), in the same way sentence (2) is built up from (3). Since (2) and (3) are clearly acceptable, this suggests that (1) is also grammatical and endorsed by the rules in our competence.

- (1) The boy the girl the dog bit kissed died.
- (2) The boy the girl kissed died.
- (3) The boy died.

Linguistic performance must of course constitute the chief source of data on the basis of which theories of competence are adduced; it is just that we do not need to reject a grammatical principle simply because concrete judgements or utterances at times fail to accord with it (even when this is not just a matter of interference in the physical manifestation of speech, e.g. an unbidden cough or hiccup).

Cohen's argument for the impossibility of human irrationality is then this: Given that the methodology in linguistics and other areas of cognitive science for uncovering cognitive competence is as described above, it seems that MRE may also be construed as a means of arriving at a theory of an underlying natural competence, in this case what we might call our *reasoning competence*. (The idea that we do, or at least might have such a competence, in line with that we have for grammar, has recently gained popularity insofar as we seem naturally inclined to make certain kinds of inferences and not others.)¹⁶ But, given MRE is also the correct method for determining epistemic norms, it follows that our reasoning competence is 'rational', that is embodies the correct norms.

Now, in itself, this argument does not entail a psychological NE (i.e. P), insofar as Cohen is not committed to seeing the norms for

belief-formation as *themselves* psychological structures in our reasoning competence. His plaint is to show that this competence cannot diverge from the relevant norms, that is, that we cannot be shown, at least by experimental means, to be irrational. Nevertheless, the move to P might seem tempting. P is the view that a theory of epistemic norms literally concerns one or more cognitive structures,¹⁷ and a reasoning competence in Cohen's sense would be the obvious way of understanding what such structures amount to. Assuming such a competence exists, and given that the method we would use to uncover its nature resembles MRE – in the same way the method in linguistics, used to uncover the nature of the principles in our grammatical competence, does – it could be tempting to conclude that a body of epistemic norms derived via MRE simply is a theory of our reasoning competence (just as a grammar, for Chomskysans, just is a theory of our grammatical competence). Such a line might be viewed as a naturalised version of traditional rationalistic philosophy, according to which the well-springs of our knowledge are prior to and independent of our knowledge of the world around us because embodied in the soul. Understanding the soul as the brain – at a suitable level of abstraction – we can view P as suggesting a transposition of this traditional idea to the naturalistic setting furnished by modern cognitive science.

Of course, there are also strong reasons for being sceptical of this way of conceiving of epistemic norms, as we shall see later on. One objection that we should consider straightaway, however, is that a theory of any cognitive competence should and will take account of more data than just intuitions or judgements. For example, it will take account of neurophysiological and behavioural evidence that MRE (narrowly construed) would not appeal to.¹⁸ However, though this may be the case, and though it may undermine Cohen's specific argument for the view that experiments cannot show we are other than naturally rational (since MRE will not by itself give an adequate account of which principles are in our reasoning competence), it does not suggest that MRE is not an important and proper *part* of a more complete empirical method for studying our natural reasoning competence. Alternatively, one might argue that MRE itself needs to be 'widened' to take account of the kinds of empirical data mentioned. Either way, the attraction of P, as a way of understanding the significance of MRE, remains.¹⁹

So much for the link between MRE and P; what about the other varieties of NE? As suggested, P would, I think, be viewed by many as problematic as an account of epistemic norms; however, one might also reject it simply on the basis of one's conception of MRE. For example, it might be held that, unlike the methods for uncovering competence in cognitive science, MRE does not have any obvious end-point; that MRE, though at any point in some sense *regulatively* governed by the aim of reaching a state of reflective equilibrium, will repeatedly destabilise as we learn more and more. The cognitive practices of science, or even of everyday reasoning, need not admit of some final, definitive codifiability; there need be no specifically epistemological object for us to uncover. Yet that would not, at least obviously, undermine (the point of) MRE and the proposing of norms *pro tem* in the ever-unfolding course of scientific investigation. Allied with this conception of MRE, NE takes on an anti-foundationalist guise: AF, as discussed above.²⁰

Yet, even if we view MRE as capable of delivering some best theory of norms, it is not obvious these should be seen as corresponding directly to some psychological capacity. Allied with this conception of MRE, NE takes on its *anti-psychologicalist* guise: AP. AP is distinctively naturalistic insofar as it sees norms as derived in the manner of an empirical theory seeking to systematise concrete judgements of validity, assuming the tenability of our basic cognitive practices. But it demurs at the idea that it is merely a study of a psychological competence.

We can shed some light on this way of understanding AP by bringing it in relation to a theme from philosophy of science. It strikes me that MRE, understood in a non-psychologicalist fashion, can be seen as having affinities with what we might term *the dominant research strategy in philosophy of science*: the appeal to the practice of actual, prominent scientists, today and throughout history, as the basis for providing a set of normative methodological rules for science.²¹ On this line – which is perhaps most famously and self-consciously prosecuted in the work of Imre Lakatos²² – the data for the construction of a theory of scientific rationality (as Lakatos terms a theory of norms) are the judgements and actions of scientists whose science we today judge to have been most successful and thus to constitute the canon of scientific knowledge – the works of Galileo, Newton, Lavoisier, Darwin, Maxwell, Einstein, Bohr *et al.* (together with

perhaps, at least in the latter cases, some less celebrated members of their research teams). At the same time, the norms we arrive at can (or should, ideally)²³ function as a corrective and guide to subsequent scientists, and, insofar as the methodological theory is generally well-confirmed, they will be asked to resist a number of *prima facie* counterexamples. Clearly this method has much in common with MRE, the main difference appearing to lie in the database: in philosophy of science we use scientist's judgements, with Goodman's original MRE those of ordinary people. It would also seem that Lakatos' proposal both (a) *does not* concern an innate reasoning competence possessed by all humans, yet (b) *is* meant to be a set of more-or-less universal rules for science – thus placing it firmly in the AP-camp. At this point, all sorts of questions arise, concerning for example the difference between AP in relation to ordinary people's judgements and AP in relation to those of scientists. I defer discussion of these to the following chapter.

Before closing the present one, I should point out that it is possible to motivate various varieties of NE independently of MRE; further, some of these would not involve the kind of justificatory process involved in MRE at all. That MRE is not absolutely essential to P, for example, should be clear from my presentation of it – even though I think that P gains much of its initial plausibility from the fact that its norms are uncovered through a process that is, in the first instance, specifically epistemological rather than scientific (the idea then being that the former *turns out to be* in fact just a special scientific method for revealing the nature of our reasoning competence). It is also the case that we shall be considering another variety of psychological NE in Chapter 5 which is conceived completely independently of MRE (I dub this *NP*, for reasons that will later become apparent).

As far as AF is concerned, as will appear from our discussion of it in Chapter 4, an MRE-based version thereof is only one of several varieties, most of which do not draw on MRE (indeed, the central intuitions motivating most versions of AF are probably quite removed from those that motivate MRE). It is also possible to conceive of a variety of AP that incorporates MRE but which does not rely on this method alone for deriving norms, as we shall see in the next chapter.

All in all, then, one should not get the impression that MRE is *definitive* of NE *tout court* (though I think it is of certain varieties,

mainly those that fall within AP). I have discussed it somewhat extensively here because it has featured centrally in many recent discussions of NE, and because it can be seen as central to at least two of the three main varieties of NE I identified through the work of Quine.

Finally, and relatedly, a word of reassurance about my overall strategy to those readers who feel that the taxonomy I have offered in this chapter fails to capture the fault-lines of real interest to an understanding of NE. I urge these, firstly, to remember that we are considering first and foremost *purely* naturalistic epistemologies, in accord with my understanding of naturalism in Chapter 1 (and that some less committed naturalistic theories will be presented in Section 5.1); and, secondly, to simply to read on, for the complexities and discussions to be offered in the following chapters will supersede the necessarily coarse-grained taxonomy I have offered here. I believe the latter is appropriate, and allows me to give what I think is an exhaustive assessment of the possibilities for NE; neither is it idiosyncratic.²⁴ However, the substantive arguments against NE have yet to come, and if there are those who would have divided up the terrain differently, I offer in advance my apologies that they may have to do a bit of chapter-hopping in order to assemble the objections relevant to the position they are concerned with.

3 Anti-psychologism

Introduction

In the previous chapter, I discussed some influences on modern retreats from TE, suggesting how we might see Goodman's method for the justification of epistemic norms, MRE, as meshing with the different forms of NE identifiable in the wake of Quine's arguments for naturalism. In this chapter, I will be considering the variety I dubbed *anti-psychologism*, or AP. AP holds that epistemology has some kind of subject-matter, and that MRE is thus capable of yielding a definitive theory of epistemic norms that might function as a kind of foundation for science (in a weak sense), but that these norms do not merely detail the structure of a cognitive reasoning competence.

I will be arguing against the tenability of AP. I will make use of several lines of argument, but I will ultimately identify as the most serious problem that we cannot make sense of epistemic norms in the way AP conceives of these in such a way that they do not amount merely to an inessential *overlay* in relation to our basic cognitive practices. For norms to be not merely such an overlay it should be the case that, at least some time, we are in a position where we *must*, if we are to form our beliefs most rationally, resort to and make use of these norms. I will seek to argue that this requirement is never satisfied in any variety of AP (at least which is not, for other reasons, implausible).

The first two sections of this chapter revolve around clarification of AP, especially in relation to the two other varieties of NE, P and AF,

and to what I called in Chapter 2 *the dominant research strategy in the philosophy of science* (henceforth *DRSPS*). Several problems with AP emerge here, which I believe carry some weight without being decisive.

Readers will soon gather that the problems I present for AP are of a 'meta' nature, in the sense that they do not seek to suggest that the norms gleaned through the methods endorsed by AP (i.e. MRE) are inadequate in relation to some or other 'intuitive' standard of normative correctness (e.g. that these methods would endorse as a logical rule something like *affirming-the-consequent*, which is 'obviously' wrong). Insofar as we are viewing MRE through the lens of our naturalism – as essentially a theory-building process, internal to the web of belief – there is ultimately no such external standard by which one might judge a proposed norm as inadequate (or if there is, one has shifted to a different form of NE). Moreover, I will take it that AP-norms, if derived in the appropriate manner, do give genuine normative guidance – specify what one rationally *ought* to do to gain beliefs – given the assumption of naturalism that we cannot but take our basic cognitive practices as valid (cf. Section 1.4). (It is these that MRE seeks to give a systematic account of.) Having said all that, my presentation does take up most of the issues that have been brought up in connection with recent discussions of MRE, albeit from a slightly different angle – for example, I will be considering the impact of putative *reasoning experts* on my arguments, as well as the possibility of 'wider' psychological factors being taken account of in the overall balancing process (i.e. a process of wide reflective equilibrium, cf. Section 2.2).

For nearly all of this chapter, I will assume that AP, of whatever kind, sees norms as being derived and justified *solely* through a process of MRE (either wide or narrow), which is how I personally have come to conceptualise its scope. However, in the last section I will broach the idea that MRE needs supplementing as a way of justifying norms for AP – in particular, that MRE only functions as a way of establishing the *goals* of our belief-forming activities, not the *means* we should pursue for attaining these goals, which should be established and justified in relation to more straightforward empirical knowledge. Hence AP delivers instrumental rules (as well as justifying the goals these norms presuppose). My arguments against this instrumentalistic conception of norms are not presented in this

chapter, but instead deferred to Section 4.5, in which the instrumentalistic conception is examined in the context of AF.

3.1 In what way is AP a genuinely distinct position?

That many have taken AP to be a genuinely distinct position – especially in relation to DRSPS – seems beyond doubt. Nevertheless, given that is not, in our terminology, anti-foundationalist, and given also that *both* AP and P (psychological NE) build on MRE,¹ there might seem to be some unclarity as to in what way AP really is or can be distinct from P.

To see this, consider the situation in which one has completed the task of providing a theory of epistemic norms in line with MRE, based on judgements of ordinary people, but which one does not want to regard simply as a theory of a cognitive reasoning competence. Saying this is of course consistent with supposing that we have such a competence, which there is some reason to think is the case² – suppose then that we possess such a competence. But then, given that, one might wonder in what way the theory of norms delivered by AP could really be anything other than a theory of this competence. An argument to this effect might go as follows. Assume the norms delivered by MRE are correct, but that they diverge from the principles embodied in our reasoning competence. These norms have been based on our intuitive judgements of correctness. But these latter judgements are presumably steered by the principles in our natural reasoning competence (albeit imperfectly, in view of performance errors); and thus, since these principles are (*ex hypothesi*) at odds with the correct norms, our intuitive judgements will give rise to a faulty theory of norms – contrary to assumption. So, if AP has a chance of delivering the correct norms, it seems these cannot diverge from the principles in our reasoning competence. This might suggest that AP must collapse into P.

A second kind of reason for thinking AP must reduce to P is that it arguably owes an account of what *makes it the case* that a putatively correct norm of inference is in fact correct: what is the supposed *truth-maker* of such norms? What in the world corresponds to the fact (if it is such) that one should do such and such to gain optimal beliefs, understood in an AP-manner? This can be viewed as a version of the second problem for NE which we discussed towards the end of

Chapter 1. In line with that discussion, the objection we are pushing here is not that there is a problem about deriving normativity from something descriptive, but rather that, in relation now to AP in particular, there seems to be no naturalistic facts for the norms to correspond to – unless we appeal to facts about our natural make-up, and say that the relevant rule is part of (or at least can be understood exhaustively in relation to) the mechanisms of our natural reasoning competence. But then it can again seem as if AP will collapse into P. (*A propos* the paranthetical remark in the last sentence but one: note that the point does not have anything to do with the ontological status of rules *per se*, for example, whether they are abstract or concrete objects. If a norm or rule exists as an abstract object, this does not make it a correct rule of inference, any more than if it exists as a concrete one. Rather, the question concerns the *state of affairs of being a correct rule*. Insofar as we are eschewing outmoded metaphysics such as Plato's Forms or the rationalists' innate ideas as accounts of what these might consist in, it seems there is little for AP to say here – again, unless it appeals to our psychology.)

I think these objections against AP being distinct from P present quite a strong case. Nevertheless, a supporter of AP would no doubt see possibilities for resisting them. Against the first line of attack she might, to begin with, insist that we can transcend our natural competence: we can acknowledge rules as correct that do not form part of our natural ability, and make judgements that also flout the principles of that ability. Judgements of correctness, though they *may* have their bases in some natural reasoning competence, need not do so. We may, in particular, learn to be more rational. More radically, a supporter of AP might seek to argue that we do not possess anything worth describing as a natural reasoning competence – an empirical claim which is certainly viable as the evidence now stands.

The idea that we can learn to be more rational is often regarded as conceptually delicate: How are we meant to bootstrap ourselves up from a state of irrationality, or partial irrationality, to something approaching full rationality? In answer to this, a supporter of AP could stress that one must in any event assume that we can learn from experience; so, why not when it comes to rationality? Further, if the case against our possessing a natural reasoning competence could be made good, it might seem that some kind of learning would have to be assumed to explain how we arrive at correct norms. The

price of emphasising all this, however, is that it now becomes unclear how AP is supposed to be distinct from AF: If our judgements of correctness are not uniquely answerable to some natural competence, but may be fashioned in accord with our experience, there would seem to be no clear way of making sense of the norms governing these judgements as any kind of foundation for knowledge, as AP seeks to maintain they are.³

Nevertheless, even if it is not very clear what the boundary between AP and AF is (given we do not see AP as basing its norms on a natural reasoning competence), this does not show that there can be no difference. That there is such a difference might moreover be supported by an appeal to the fact that AP has underlying affinities with DRSPS, a project that seems to be aimed at giving a, in some sense, foundationalist methodology, but not by way of tapping the structure of a psychological competence. Perhaps one could also credit the idea of certain methodological facts subsisting in the social practices of the scientific community as the truth-makers of the proposed norms. (We shall have more to say about the connection to DRSPS in the following section.)

Alternatively, a supporter of AP might insist that even if the epistemic norms gleaned by way of MRE do *correspond* to the principles in our reasoning competence, it would not follow that the norms are *identical* with these principles. Since facts about our reasoning competence might seem to be irredeemably descriptive, whilst norms are precisely meant to be normative, she might, to begin with, take herself to have very good motivation for denying the very idea of such an identification. She could then go on to say that whilst our norms may have a necessary *causal* underpinning in our psychological make-up, these two things – the norms and their causal underpinning – are, necessarily, distinct from one another. Taking this line could, perhaps, also serve as a way of rebutting the objection that we can only give an account of the truth conditions for any AP-norm by embracing P, for AP can now see its norms as *corresponding* to the structure of a reasoning competence without being *identified* with them.

These last three replies in defence of AP might be further buttressed if there turned out to be some irreducible *indeterminacy* in arriving at rules of inference through MRE. This would arguably exclude seeing any set of rules so gleaned as simply delineating the structure of a

reasoning competence, whose nature is presumably fully determinate; it would also cast doubt on the coherence of a demand to say what the norms are made true by. Yet indeterminacy need not, it would seem, render the project of giving a set of rules completely otiose.⁴

For all these different reasons, it seems the case for saying that AP reduces to P – or else AF – is inconclusive (even though the same undoubtedly goes for the counter arguments in defence of AP). In the following section, I will explore in a little more detail the connection between AP and DRSPS, before proceeding to what I see as the fundamental problem for AP.

3.2 The dominant research strategy in the philosophy of science (DRSPS)

As noted in Section 2.2, though AP can be understood purely in relation to MRE as applied to ordinary people's concrete judgements, it can seem more intuitively attractive when seen in relation to DRSPS, which applies something like MRE to the judgements and practices of scientists. If DRSPS is viable as a way of gleaning methodological norms, as Lakatos and others hold, then, as good naturalists, we should both accept these norms as valid, and accept that they are not concerned merely to delineate the structure of a psychological reasoning competence. After all, insofar as we are deferring to science, and insofar as DRSPS is a scientific way of studying science's reasoning practices, who is in a position to say that, contrary to appearances, they are merely doing psychology?

However, insofar as this connection is important for AP, its supporters must contend with the fact that DRSPS itself has been subjected to swingeing criticism in much recent philosophy of science. In the light of both historical and contemporary sociological analyses of science, the idea that there is one overarching method or set of epistemic norms for science seems to be a quixotic ideal. If there are rules or norms for science, it seems that these change over time. Newton relied on careful observation and saw his theories as descriptive laws that could aspire to acceptance as highly probable, if not absolutely certain. Galileo, by contrast, used a method involving the proposal of mathematical laws going beyond anything one could observe, which only subsequently could be verified in relation to carefully controlled experiments. In the nineteenth century, the

rejection of apodictic certainty as a goal for science led to a general acceptance of *hypothetico-deductive* reasoning which earlier thinkers would have disparaged. In more recent years, quantum theory has challenged the centrality of certain cherished logical principles, such as the law of excluded middle, when dealing with the behaviour of the most fundamental constituents of matter. At the same time, the flourishing of non-physical science has brought with it a plurality of methods that it would seem difficult to reconcile with one superordinate set of rules. And studies in the sociology of science have suggested that strictures such as 'avoid ad hoc hypotheses' or 'accept the results of decisive experiments' do not have the rigid, prescriptive force that many philosophers of science have supposed them to have.

This meagre handful of examples only gesture at an enormous field of enquiry in the history and sociology of science, but the conclusion they seek to support – that methods in science are indefinitely mutable in the face of whatever is demanded in the name of furthering our understanding of the natural world – is today becoming ever more widely accepted.⁵ In the context of our discussion, they can seem to enunciate a need to move to an anti-foundational conception of scientific method, that is to AF. For present purposes, however, the important point is that they compromise AP insofar as they compromise DRSPS and its monolithic conception of method.

Of course, these considerations are, especially as I have thus briefly presented them, inconclusive with respect to the principled issue which concerns us here: whether giving epistemic norms in the manner envisaged by AP is both cogent and required. Moreover, though it is most famously relativistically inclined philosophers of science who have drawn on methodological mutability in support of their views, more recently there have been others seeking to see it as reconcilable with a rationalistic perspective on science. Indeed, some of these reconciliations might suggest a way of saving AP. Thus Larry Laudan, though an ardent supporter of the changing methods-view, holds that facts about actual scientific practice must and can still furnish the premises for genuine normative evaluation of science, at some level superordinate to that of specific methodological rules.⁶ A related but distinct view is Chalmers', that scientific methods can be shown to have some very general, essentially common-sensical normative feature in common; he

suggests they may all be seen as conforming to something like the following rule: *Take argument and the available evidence seriously and do not aim for a kind of knowledge or a level of confirmation that is beyond the reach of available methods.*⁷

In the present context, I have nothing against such reconciliatory lines. However, I do not think that the possibility of specifying very general foundational norms for science is the reason to think that science needs specifically *scientific* norms to make optimal rational progression. The kinds of general norm proposed have a tendency to be either inherently vague or mere common sense – or both (Chalmers' dictum would seem to fall into the last category). Thomas Kuhn is also known for his view that most viable, inter-theoretic norms and standards of science – simplicity, elegance, problem-solving capacity and so on – are so vague and imprecise that interpretation will always be required for them to be meaningfully applied.⁸ If one interprets this in a way that is not conducive to relativism (more on Kuhn and relativism in Section 3.4), what this amounts to saying is that vague norms offer extremely little by way of substantive guidance to the scientist. I think it is also plausible to add that, to the extent they do offer any such guidance, their content scarcely goes beyond precepts of common sense. Insofar, a theory of norms in accord with DRSPS would coincide with a theory of norms derived through MRE taking as its input just the judgements of ordinary people. Thus the appeal to science will have been an inessential dog-leg, and a defence of AP would have to resort to the considerations adduced in the previous section.

In sum, as far as AP goes, associations with science have arguably little to offer. Nevertheless, this is not a water-tight conclusion; nor have we assumed or established any necessary dependence of AP on the ideas behind DRSPS. Thus it remains to find a more principled argument against AP.

3.3 The fundamental problem for AP

It seems to me that there is in fact a very simple thought that renders AP hard to sustain – even if we give it the benefit of the doubt on the points of contention so far mentioned. Remember that, according to AP, with its stress on MRE, norms are derived on the basis of inferences made by ordinary people and/or by scientists that, in the first

instance, are not governed by norms or rules – at least, by explicit rules – but merely, so to speak, ‘flow’ from their ability – from what we have called basic cognitive practices. We are allowing that this ability may, but need not be based (at least entirely) on a natural psychological competence, and also that, even if it is, it may in any case be reasonable not to identify the norms gleaned from the principles contained in this competence, partly because it may be indeterminate what these norms are. However, even on this description of the situation, one can surely legitimately raise the question: *Why bother setting up explicit norms at all?* Whether or not they correspond to the principles in a reasoning competence, why should not we simply let ordinary people and scientists *get on with it?* We might indeed be able to uncover certain rules that, in some sense, they could be said to be following – *implicitly* guiding principles (which may, but need not be mirrored in their actual psychological make-up). But why should we go to the trouble when it seems these are not necessary for their rational belief-formation – and given we are not primarily concerned with a psychological study? I see this as *the fundamental problem* for AP, and will refer to it as such in what follows.

A supporter of AP would no doubt deny that the rules would be redundant insofar as they could play a substantive role in *streamlining* and/or *expediting* the *application* of the ability we, or scientists possess (as based on a natural competence, or otherwise). But though this is perhaps true, it is another matter whether setting up the principles underlying this ability explicitly would be the only, let alone the most effective way of doing this. Thus consider an analogy with our linguistic ability. One could no doubt imagine that speaking grammatically correctly could, to an extent, be improved by making people aware of the rules that either they tacitly know, or that otherwise ‘systematically characterise’ their ability.⁹ However, it seems clear that this same goal would be more effectively ensured by advising people to concentrate properly on what they are saying, on whether their utterances really mean what they want them to mean and so on – in short, to apply their ability *competently*. Someone might retort there are important disanalogies between the linguistic case and that of epistemic norms conceived as derived through MRE. However, it is unclear to me what these are meant to be; in particular, since MRE is seen as operating on the products of a basic cognitive ability, it would seem unexceptionable to speak of more or less competent use

of that ability. At the very least, the present defence hardly shows that setting up epistemic norms explicitly is necessary for optimal belief-formation.

There is also no question of people needing the norms in order to be simply *instructed to reason rationally*. Again: AP assumes that reasoning and forming rational beliefs is a basic ability we possess when we think about the world, and that MRE is the method for making explicit the principles underlying this ability. That we engage in belief-formation at all is just a fact about human beings (or not, as the case may be!).¹⁰ What AP needs to establish is that, in situations where we in fact reason, we must take account of the norms it offers. I am arguing it is highly unclear why this should be so.

Does the possibility of using a wide form of MRE – taking into account empirical facts about our computational limitations or our cognitive abilities in general – affect these points in any way? Presumably, the idea would be that knowledge of certain limitations could rule out certain intuitively attractive norms that would have emerged through a narrow process of MRE, for example, exhaustive consistency-checking (I will focus on this example).¹¹ Thus, it might seem that there would be a point to deriving and explicitly setting up the norms delivered through a wide process of MRE – in the present case, to get people to abjure exhaustive consistency-checking. However, if consistency-checking really were countenanced by a process of narrow MRE, it presumably would have *something* going for it. Some degree of consistency-checking would seem, therefore, to be apposite, or at least legitimate. Concerning the question as to what extent, we can imagine two possibilities in principle: that there should be a norm, or several norms – for different, specified situations – regulating the degree of consistency-checking one should go in for; or that it should be irreducibly a matter of individual and context-dependent judgement. In the latter case, a norm proscribing exhaustive consistency-checking would not be required, because we could instead just inform people about their computational limitations, and let them limit their natural disposition to seek consistency as they see fit, according to the circumstances. In the former case, the relevant norm or norms would have to be derived and justified – which, for a supporter of AP, would go *via* a process of MRE seeking to derive a theory of norms by simultaneously taking into account concrete judgements together with knowledge of human computational

limitations. But then it seems my argument against the non-necessity of these norms would go through in much the same way as before: instead of using the thus-derived norms, one could simply apply one's cognitive ability assiduously, taking into account, as one did so, knowledge of human computational limitations.¹²

A *prima facie* more promising answer to the fundamental problem would build on the fact that we can distinguish *different kinds* of ability, existing in different groups of people, as the basis upon which to derive norms of reasoning. In much of the recent literature on the use of MRE in gleaning epistemic norms, the idea of *reasoning experts* has been taken up, where such experts are viewed as groups of people whose reasoning ability is somehow 'enhanced' relative to the rest of the population (they are standardly assumed, amongst other things, to properly understand logic and probability theory).¹³ Thus it is often suggested the correct epistemic norms are those that are, as it is put, in *expert reflective equilibrium* (which might also take into account wider factors about our cognitive limitations), rather than the reflective equilibrium ordinary people would reach. For our purposes, the important point is that whilst experts would perhaps not need to have their principles made explicit for themselves – for the reasons already outlined – nevertheless, doing this would be beneficial for ordinary people, for it would enable them to form their beliefs in a more rational fashion than they could simply by employing their own ability.

It has been objected against the idea of reasoning experts that it is circular to appeal to them in justifying epistemic norms. To do this, it is said, we would need to know who the experts are – but how are we meant to do this when the experts themselves are needed to tell us how to form our beliefs? Maybe the people we think are experts, and who turn out to be experts on their premises, are in fact, to use Stephen Stich's memorable expression, 'quite daffy' – or, worse, intent on exploiting us.¹⁴

I am not convinced that what is problematic about the idea of reasoning experts is that it enunciates a vicious circularity. In a holistic, naturalistic epistemological setting of the kind we are assuming, not knowing how to determine that a statement *S* (in the present case: that it is such and such people who really know how to reason) has reliability without, at some point in the process, relying on *S* itself should not – in and of itself – engender scepticism with respect

to *S*. Nevertheless, there is, I think, something fundamentally odd about the idea of reasoning experts in the absence of some further substantiation.

One such substantiation might be to view them as philosophers, engaged in developing logical and mathematical theories of reasoning. But if this idea is to be grounded, ultimately, on these philosophers having access to purely a priori sources of knowledge, then the norms they delivered could not be classed as naturalistic.¹⁵

The only other feasible answer would seem to be that the experts are really just some kind of scientist, and indeed this suggestion will occupy us for most of the rest of this chapter.¹⁶ That reasoning experts could be anything other than scientists or traditional philosophers, and at the same time *bona fide* experts – and not just some upstart group seeking to exploit or oppress other people – seems highly dubious. Surely rationality, or the potential for it, is the prerogative of all human beings, if anything is. Should it become instead a commodity – perhaps something for sale from one group to another – it seems hard to credit the idea of it really being *rationality* we are talking about. In any case it seems unlikely to be the notion traditionally linked to the epistemic practice of getting at the truth (or however one conceives of science's ultimate goal) that I am concerned with here.

What, then, of scientists? If they are the relevant expert group, one could concede that scientists do not themselves need the norms that might be gleaned from an analysis of their own activity, but simultaneously insist that those of us who are not scientists do – given that science represents the most advanced means of finding out about the world, and that our everyday reasoning ability, even if not based (exclusively) on some natural reasoning competence, does not suffice for advanced scientific activity. Learning the rules of science might seem to be an ideal way to enter into the sphere of the scientific, of expediting this entrance, or at least of in some way enhancing one's epistemic ability.

I think it must be confessed that this reply on behalf of AP is rather far-fetched. It seems implausible to suppose there should be any monolithic set of rules one could learn in order to become, or even to expedite becoming, a scientist – not least in view of the considerations concerning variations in methods and the vagueness and/or common-sensical nature of the general principles adduced in

the previous section. (Note, then, that I am not here saying that learning rules of *any* kind is irrelevant to becoming a scientist, for example, instrumental norms linking certain methods to goals – though, as I will argue in Section 4.5, these kinds of rules either essentially are, or at least can always be replaced by, descriptive statements.) However, though I think this point is highly significant, I would like in the following to essay a more principled rebut to the above reply to the fundamental problem for AP; I will refer to it as ‘last ditch reply’. I will be building on what many have seen as central insights into the nature of scientific knowledge and competence in the theory of these developed by Kuhn in *The Structure of Scientific Revolutions*, especially the postscript of 1970 in which he links his ideas to those of Michael Polanyi on the role of *tacit knowledge* in science.¹⁷ The account and arguments offered below might be seen as an attempt at *explanation* of the (putative) fact that the only rules that can be gleaned from an analysis of scientists’ abilities are common sensical and/or inveterately vague.¹⁸

3.4 Kuhn on exemplars and tacit knowledge

One of Kuhn’s central ideas in his overall philosophy of science is that scientific competence is not acquired through learning explicit theories or methodologies, nor does it consist in such abstract knowledge. Instead, it involves the ability to understand concrete, observable states of affairs – what he calls *exemplars*, or simply, in one particular sense of the word, *paradigms* – in relation to the abstract theoretical formalisms. This ability is then applied creatively to new observable circumstances. For example, to become a competent Newtonian physicist, it is not sufficient to apprehend Newton’s various laws of motion and gravitation. Take for example Newton’s second law, *Force = mass × acceleration*. To really understand this, one has to be able to relate it to other more specific mathematical formulae which are instances of it, and, ultimately, through problem-solving exercises, to particular concrete set-ups involving the planets, pendula, gyroscopes, rolling balls and so on. The things one sees in the experimental set-ups become in this way reconceptualised in accord with the abstract theories. Then – if one is sufficiently brilliant – one applies the theories to new observational phenomena, creates new paradigms and so on. Scientists thus become competent

members of paradigms – understood now in the broad sense of a community-based scientific research programme – in virtue of a skill that we must understand, not in terms of that which may be learned through explicit instruction, but in relation to practical experience with concrete, observable set-ups that one seeks to make sense of through the ‘gestalt’ of the relevant theoretical formalisms.

How does this account serve to rebut the last ditch reply on behalf of AP? The important idea is that in gaining scientific competence, one does not acquire something that can be explicated in terms of *abstractly specifiable rules*, that is rules that could be made sense of and applied independently of the concrete exemplars that scientists understand their theories in relation to. Insofar as they use reasoning abilities that can be explicated in this way, it seems scientists use just those ordinary people possess. We can assume that they draw on these reasoning capacities *inter alia* in order to understand how certain concrete situations can instantiate the properties which their abstract theories suppose them to have. This latter kind of understanding is, for sure, not something which ordinary people possess, but neither is it the case that its acquisition could be secured by explicitly learning certain rules or criteria abstracted from an understanding of the concrete exemplars.¹⁹ It is in this connection that Kuhn sees links to Polanyi’s notion of *tacit knowledge*: a kind of knowing that is inarticulatable because it is inherently bound to the context of its application.

It bears emphasising here how different Polanyi’s notion of tacit knowledge is from that standardly used in cognitive science, where it concerns unconscious knowledge of particular principles and rules that precisely could – in principle – be laid out explicitly.²⁰ On Polanyi’s and Kuhn’s view, what ordinary people lack in relation to scientists is, for sure, knowledge of the natural world, but this is not something that they can simply be told about, nor can they reason themselves to it with the help of some abstract canon of scientific methodology, in the way the last ditch reply assumes. What they need is to be able to apply their everyday reasoning competence to a way of seeing the world in terms of the new theoretical formalism. But if that is the case, then AP seems straightforwardly vulnerable to the fundamental problem – to wit, the charge that the principles underlying this everyday competence do not need to be made explicit for one to reason rationally.

In the last two paragraphs I have been relying on what many see as insights into the nature of scientific competence and knowledge due to Kuhn and Polanyi. Nevertheless, it would be folly to do otherwise than admit that, however influential he may have been, Kuhn's views (in particular) in philosophy of science are controversial. In the light of this, it may seem that my argument is rather flaccid insofar as I have presented no systematic defence of Kuhn. However, I think that assessment would be unfair.

For one thing, part of what I am relying on is Kuhn's historical and sociological knowledge of science, not just his philosophical analysis; I am not in a position to go into depth about these historical and sociological details, but that, presumably, does not rule out using the conclusions of those who have done this. A more significant point is that I am only building on Kuhn's specific ideas concerning scientific competence and its relation to exemplars, and eschewing his more relativistic ideas – what is chiefly responsible for Kuhn having won so many inveterate opponents. This strategy seems reasonable insofar as Kuhn himself denies he is a relativist, allowing that science has made progress through the centuries in the sense of increasing in predictive accuracy and problem-solving ability. He denies it makes progress towards 'the truth' in any substantial sense – but that assessment is something many others would concur in. Connectedly, that scientific theories essentially involve seeing the world in different ways does not obviously exclude rational comparison of theories, as Kuhn himself has stressed²¹ – nor even if, as my own view entails, such comparison must be accomplishable without resort to explicit rules or standards. Further, Kuhn's notion of a paradigm (*qua* research programme) *can* be understood as a fairly loose and labile concept that can be made to fit in with several historically documented and intuitively plausible aspects of scientific activity – aspects which, on the more strict reading that some of Kuhn's writings foster, seem to be foreclosed (something which admittedly makes this more strict reading untenable).²² Thus, in particular, talk of different paradigms does not need to imply that different schools of scientific thought are always radically opposed to one another, rather, they may peacefully coexist (as for example evolutionary biology and organic chemistry now do), or draw on partially overlapping theoretical resources at the same time as specialising on particular problems within the same general domain. Similarly, the

notion of a paradigm is consistent with the idea of a fair degree of continuity of scientific expertise across episodes of what Kuhn dubs scientific revolution (Einsteinian physics did not overhaul *all* of the foregoing mechanics). It seems, in other words, quite consistent to think of paradigms as articulated into some kind of hierarchical structure, both diachronically and synchronically – and of those that are wholly disjoint as not in conflict with one another.

What Kuhn is arguably most concerned to stress – at least, what I have been drawing on here – is the idea that genuine theoretical knowledge is not just abstract structure, but rather is constitutively tied up with a scientific skill or expertise that mediates, on the one hand, theory and, on the other, what we perceive and how we act. We may of course all gain some kind of superficial understanding of a scientific theory we are not proficient in through popular presentations, and this will be important to a general awareness of the subject area, of how it relates to one's own area of interest and to that of others', as well as possibly influencing one's future research interests. Nevertheless, this understanding is superficial compared to that of the experts, even when it is, in some sense, quite extensive. I am suggesting that this view explains the fact that what a scientist acquires when she masters a theory is not pure formal knowledge, nor an inferential capacity which we could characterise with the help of explicit rules, but a context-specific ability that only has genuine application in relation to the relevant exemplars. Thus, the fundamental problem for AP remains: the last ditch reply to this problem assumes (as I see it) a faulty conception of scientific competence.

Of course, Kuhn's theory of scientific competence is still controversial even when 'cleansed' of its relativistic associations in the way I have just sketched. In particular, it downplays quite generally the idea that scientists are dependent on the following of rules or norms (at least, peculiarly scientific ones), and this may strike some as so obviously distortive of fact that the theory must be rejected. Now, it is still open that Kuhn should be wrong about the significance of scientific norms. Moreover, if some account of them other than that in which they are derived from an ability were tenable – for example, an understanding of them as *instrumental* norms, or something else along the lines to be discussed in the following chapter then his general line would have to be rejected. However, I will be suggesting that such accounts do not show the need for genuine, explicit norms in science.

This still leaves us with the supposed 'datum' that scientists are *de facto* thoroughly dependent on explicit norms in their work. (This is something that the totality of what I myself establish in Chapters 3 and 4 – i.e., subtracting the contribution made by Kuhn's conception of competence – is equally in tension with, since I take myself to show that at least *scientists* do not need norms.) I will thus also be offering explanations as to why this 'datum' might seem to be thus were it in fact an illusion (cf. Section 4.6). None of this will prove Kuhn right on the question of scientific competence, but it should remove barriers to a more sympathetic attitude towards his view – and do that and more with respect to my own arguments.

3.5 An alternative conception of AP

In this chapter, I have been assuming an understanding of AP that builds on MRE exclusively (possibly a wide variety thereof). I have suggested that there is a fundamental problem for AP in seeing its norms as necessary for optimal belief-formation because MRE derives its norms on the basis of pre-existing cognitive abilities, seeking to encode and explicate the nature of these abilities – the point then being, why not just seek to employ these abilities to better effect? I have also suggested that there is probably no monolithic set of norms of science that something like MRE might be used to glean, and that this can be explained by Kuhn's idea of scientific competence as based on (Polanyi's sense of) tacit knowledge. Hence we cannot see ordinary people as having something to learn from scientists reflecting on their fundamental reasoning practices.

But is there any other conception of AP which does not build on MRE, at least exclusively? Arguably there is, for one might, as I so far have not done, distinguish between the *goals* or *aims* of epistemic activity, on the one hand, and the *means* or *methods* for attaining these, on the other. As we shall see in Chapter 4, this is also a central distinction in various different versions of AF (such as Laudan's and, as evinced in later works, Quine's). Coupled with AP, the idea would specifically be that MRE could be used to determine the nature of the relevant goals of epistemic activity, but that the means, the methods, for achieving these would be a more straightforward matter of empirical enquiry. Thus MRE might be employed to determine that, say, truth, logical consistency, accuracy, speed or whatever are epistemic

goals for us, whilst empirical enquiry could tell us what ways are most suited to achieving these goals (consonant with our inherent cognitive limitations).

On this line, there could presumably, at some point, emerge some relatively final set of prescriptions concerning how best to achieve our goals – a final set of *instrumental* norms. But since these prescriptions would not be based on MRE,²³ what I identified as the fundamental problem for AP could not arise. Alternatively, or in addition, these prescriptions might be seen as constituting scientific or expert advice about how best to reason (we have of course cast a number aspersions on the idea of a monolithic set of norms for science, but perhaps the present conception is apt to repel or at least assuage these).

I will have much more to say about this instrumentalistic conception of epistemic norms in Chapter 4. However, the basic problems with the present conception of AP can be outlined here. Concerning the goals, in line with what I have argued earlier in this chapter, I would say that it is not necessary to make them explicit for, in being derived on the basis of our cognitive practices, they are things we will aspire to anyway. Further, there would seem to be no overarching set of goals peculiar to science that scientists, by reflection on their own practice, might otherwise have communicated to lay people.²⁴ That the goals do not need to be made explicit does not of course show that instrumental norms recommending certain means for attaining these do not do so; my objections to the idea of instrumental norms are presented in Section 4.5.

4 Anti-foundationalism

Introduction

In this chapter, we shall be considering varieties of NE that fall into the category of *anti-foundationalism* (AF), as outlined in Chapter 2. The central idea behind AF is not merely that a theory of epistemic norms must be derived from other than apriori sources; in addition, the norms are not seen as constituting a unitary, unchanging basis for science or cognitive inquiry generally. According to AF, norms evolve indefinitely in the same way our scientific knowledge does.

Even thus baldly characterised, AF can appear to have certain virtues relative to the alternative forms of NE (i.e. AP and P; the following points allude to discussions in Chapters 2 and 3). To begin with, it is *eo ipso* not psychological, something which to many would be seen as an automatic virtue if the aim is to give normative principles for inference and reasoning.¹ Secondly, in consequence, there is also a clear demarcation between AF and P, as there is not obviously between AP and P (for, since AP assumes that the study of norms is of some stable set of entities that are derived on the basis of people's concrete inferential judgements, it becomes unclear how this should be conceived as concerning anything other than a psychological reasoning competence). Thirdly, and relatedly, AF would not seem to owe an account, as AP arguably does, of what *makes* its norms true or correct insofar as the norms it offers are always essentially *pro tem* – or else (as we shall see) based directly on empirically ascertainable connections between natural phenomena. Fourthly, and again relatedly, there would appear to be no 'fundamental

problem' for AF as there is for AP in understanding how norms can be necessary for cognitive enquiry (for they are not derived on the basis of some stable ability, natural or otherwise, which one might simply hone the use of instead of using the norms).² Finally, AF does not assume that knowledge and science are built on an unchanging, methodological foundation as AP does; on the contrary, any norms are meant to be as mutable as the science they serve – a feature which arguably accords better with the facts of scientific history (cf. Section 3.2).

Different variants of AF also have their own intuitive appeal, as we shall see in the following. However, AF also faces some intuitive problems. Perhaps the most obvious of these stems from the very fact that, unlike the other varieties of NE, it eschews *universality* in its understanding of norms (cf. Section 1.2). Now, we should not confuse this eschewal with an embracement of some form of methodological or epistemic relativism, that is, the idea that norms are only valid relative to some epoch, culture, scientific paradigm or such like. The idea behind AF is not that a norm is only ever correct or (perhaps better) acceptable relative to some stage in the historical progress of science at which it is proposed; rather, it is that the very notion of something being an acceptable norm presupposes a particular historical stage in the progress of science and scientific knowledge. AF holds that a set of norms derived in an appropriate manner are acceptable at the time at which they are put forward; since there is no prospect of ever completing our understanding of the natural world, we will later replace them with different ones, but since there is not either any question of our transcending the situation we are in at any time, there is no perspective from which we can regard them as only relativistically valid. Nor does this preclude the supporter of AF (as we shall see below) from talking of norms as making progress in step with progress in our scientific understanding: this is also something no relativist would accept. So relativism and AF should not be confused.³

Nevertheless, there is a more subtle problem with AF connected to its rejection of the idea of universality. If AF is going to be viable, then its norms are going to be what I have been calling *genuine* – they must specify which steps genuinely *ought*, rationally, to be taken to gain beliefs at a given time. Now consider that it is highly plausible, for this to be satisfied, that the beliefs which the norms are meant to

help form must take part in some kind of rational epistemic progress. In fact, this assumption is, I think, implicitly contained in our characterisation of naturalism in Section 1.4 – specifically, in the idea that our cognitive practices must be assumed to be fundamentally valid. For a supporter of P or AP, this means that genuine epistemic norms can be gleaned through a systematisation of our basic cognitive practices. A supporter of AF rejects the idea that there is any unitary system to be gleaned from these practices,⁴ and instead seeks to see norms as arising out of ongoing scientific activity. What she can, however, still say about our basic cognitive practices in this case is that they must be assumed to be moving our beliefs in fundamentally the right direction – ensuring that they are somehow, in some sense, making progress.⁵ The point in the present context is then that, without this assumption, there is simply no reason to think any normative guidance that might be offered of an anti-foundational kind – in the form of general norms or otherwise – could genuinely specify what one *ought* to believe.

My main line of argument in the following will be to argue that, to the extent we can make sense of AF-norms being genuinely normative in this way – as specifying general guidance that would expedite the rational progress of our beliefs – the norms will appear to be not in fact necessary for us to make optimal such progress. I aim, more specifically, to show, in relation to several different sub-varieties of AF, that the proposed AF-norms, if they are to have genuine normative force, must presuppose either corresponding norms of other kinds – traditionalistic or naturalistic of other than an anti-foundationalist variety – or else corresponding scientific facts; and that from this basis, it can be argued that the AF-norms are superfluous to the progress of science.

Sections 4.1 and 4.2 contain a relatively direct application of this line of argument to different forms of AF that can be constructed out of well-known ideas from the philosophy of science and epistemology. The first, in Section 4.1, builds on John Stuart Mill's conception of induction, where the ideas behind my main argument against AF are set out at greater length; the second, described and critiqued in Section 4.2, builds on (a certain conception of) so-called *evolutionary epistemology*.

Section 4.3 raises to a large extent rather different considerations against a form of AF based on MRE (as I originally motivated it in

Chapter 2). In Sections 4.4 and 4.5, something more like my main line of argument resurfaces, but here in relation to (at least in the first place) a fully worked out proposal from recent philosophy of science: Larry Laudan's *normative naturalism* (NN, as I shall refer to it). NN can be seen as operating with two different levels of epistemic norm: values, goals or aims, on the one hand, and rules or methods on the other (cf. Section 3.5). In Section 4.4, I argue that, in relation to goals, NN must reduce to some form of AP to be tenable. In Section 4.5, I argue that that NN's conception of methods as *instrumental* or *hypothetical* norms does not warrant viewing these as necessary for rational progress in science. In fact, Laudan's is only one of several AF-theories that stress the idea that epistemic norms are to be viewed fundamentally as instrumental rules. My critique in this section is thus not specific to Laudan's theory, but concerns what has become quite a general conception of norms amongst supporters of NE.

Finally, in Section 4.6, I offer a discussion of how it can be that norms seem to abound in science when, according to me, they are not in fact necessary for its rational progress.

4.1 Mill's inductivism

I believe that one of the clearest instantiations of AF is to be found in John Stuart Mill's views on scientific method, which he sees as fundamentally based on induction. According to Mill, enumerative induction is the most basic form of reasoning used by us, and cannot but be assumed to be valid at the outset. However, it is open to adjustment and improvement:

[I]f induction by simple enumeration were an invalid process, no process grounded on it would be valid; just as no reliance could be placed on telescopes if we could not trust our eyes. But though a valid process, it is a fallible one, and fallible in very different degrees: if therefore we can substitute for the more fallible form of the process, an operation grounded on the same process in a less fallible form, we shall have effected a very substantial improvement. And this is what scientific induction does.⁶

This passage encapsulates (in my terminology) both a central naturalistic and a central anti-foundationalist intuition: science is capable of

producing normative guidance for theory-construction from *within* this same ongoing activity of theory-construction, a process which must all along be assumed to be fundamentally in order but which can be expedited by reflection upon, so to speak, 'where we have to go to so far'. Such a view of the source of epistemic norms implies that there is no need for an analysis of our or scientists' reasoning capacities in order to provide substantive normative guidance for science (as P and AP assume).

But how exactly is the process of improvement meant to take place? The simple answer is *via* experience:


[W]e make experience its own test. Experience testifies that among the uniformities which it exhibits or seems to exhibit, some are more to be relied on than others; and uniformity, therefore, may be presumed, from any given number of instances, with a greater degree of assurance, in proportion as the case belongs to a class in which the uniformities have hitherto been found more uniform.⁷

In a more modern and didactic idiom, George Couvalis summarises the Millian line as follows:

[...] Mill thinks that we start our investigations of nature by assuming all enumerative inductions with true premises prove their conclusions. In his view, there is no better place to start than with such assumptions. Any methodology which attempts to test scientific conjectures without using induction will fail to provide any justification for believing particular conjectures. However, Mill holds that our starting assumption needs to be substantially modified. We quickly realise that enumerative inductions over some sorts of cases produce many cases which are falsified [...]. For example, we quickly realise that when we infer from the colour of a few instances of a species that all members of that species have that colour, our premise provides little support for its conclusion. However, in other sorts of cases, enumerative induction seems to produce statements that are not falsified, so we come to understand that enumerative inductions over such cases provide stronger support for their conclusions. [...E.g....] generalisations about broad structural anatomy of all

the members of species [...] are rarely wide of the mark. [...] We can and do decide on the basis of large numbers of trials that enumerative inductive arguments about some sorts of cases are at least much less probative [i.e. cogent] than arguments about other cases.⁸

Thus, the inductive canons in science – our epistemic norms (at least those we need beyond general logic and perhaps other background principles)⁹ – can be justified in pretty much the same way that we justify 'first-order' scientific hypotheses, that is, by proposing general hypotheses on the basis of observation, and refining them in the light of subsequent observation. What Mill is advocating is thus a hierarchy of inductions, with higher level norms being used to steer the creation of lower level hypotheses, and the higher level norms in their turn being modified in relation to the observational success of the lower level hypotheses they generate. There need be no end to this process, even though one can expect greater stability to emerge in our inductive canons after a while, and even though it seems reasonable to suppose we will thus be making progress towards evermore adequate canons. Pictorially, the idea can be represented as follows:



The sloping/vertical arrows indicate directions of justification or support, and the horizontal, bold arrow indicates the progress (assumed to be) being made over time.¹⁰ N1 and so on are inductive canons, or norms, H1 and so on are hypotheses. In the following, I use the term 'hypothesis' to refer to ordinary scientific hypotheses concerning the nature of the world, and 'norm' to refer to inductive canons and norms in general, even though both, according to Mill, have a hypothetical status.¹¹

Is this view of scientific norms and their justification tenable? A general worry someone might have about Mill's approach is that it offers no *guarantee* that induction is at all on the right track – no guarantee, that is, that controlling the formation of hypotheses and epistemic norms in relation to observation has anything to do with obtaining true beliefs, or making rational progress. Understood in one way, this is just a version of the worry about the circularity of naturalistic epistemological theories that we discussed – and set aside – in Section 1.5. However, a naturalist herself might also reject Mill's theory for reasons specifically to do with its reliance on *observation* as the basis for forming hypotheses (an alternative version of AF to be discussed in the following section uses instead pragmatic success as its basis). I will not, however, be concerned here to assess which 'database' is most apposite to the gradual refinement of hypotheses and norms. Instead, I want to concentrate on a failing in Mill's account that I believe points up a fundamental weakness common to it and other structurally similar versions of AF.

Let us therefore suppose that the process of refining hypotheses iteratively in relation to (as it is in Mill's case) observation is fundamentally apt as a way of gaining evermore adequate hypotheses. That the process is thus simply assumed to be fundamentally 'on course' is, in any case, something we are viewing as part of its naturalistic credentials, and thus not as controversial (see the introduction to this chapter). At the same time, as we also noted, this assumption is a necessary condition for any norms suggested being genuine, that is, it is essential to the idea that the norms specify what one *ought* to do (in some non-relativistic sense) in order to gain rational beliefs. Take the assumption anyway, and the whole process becomes nothing more than an intellectual game. Moreover, it is important to bear in mind that it is *just* an assumption; the naturalist has, *ex hypothesi*, no a priori account of what makes induction generally, or any particular inductive rule, correct.

With the assumption in place, it should nevertheless be clear that using information about which kinds of hypotheses turn out to be most reliable (in relation to observation) in order to derive norms about which hypotheses to form in the face of observation can be assumed to be a reliable method for forming evermore adequate (though always fallible) norms. The question I want to

raise, however, is this: why, given that we must assume that shaping one's hypotheses in relation to observation is a reliable process, will it also be necessary to postulate norms in order to steer these hypotheses – from 'above', so to speak? In other words, why will such norms be *necessary* to form rational beliefs, even if they are, as we are allowing, *sufficient* (within the bounds of their fallibility)?

This might seem like an obtuse question. It might seem obvious – under the assumption that observation can steer us towards better hypotheses – that there will be a potential for greater steering by generating higher level norms. But what grounds this feeling of obviousness, and what kind of steering is this really meant to be? Note first that it is not a steering whose epistemic authority differs in *kind* from the steering hypotheses gain from observation. If there is anything fundamentally lacking in the latter with regard to the *kind* of steering it gives, steering by norms as Mill conceives of them will clearly not compensate for it.

The idea must then presumably be that the norms will somehow *assist*, along with observation, in the creation of good hypotheses – that they will lead to *better* hypotheses than observation alone could sanction. However, it is highly unclear how – since the norms are merely derived on the basis of hypotheses so far chosen, just as the latter are derived on the basis of available observational data – they can be in a position to do this. To see this more clearly, consider the following purely illustrative example (an extension of that in the Couvalis quote, above). Suppose we have observed certain portions of the world, and found good evidence for the following inductive hypotheses, and others of their ilk:

- H1: All swans have one head, two wings, two legs, two eyes, two kidneys *etc.*
- H2: All fish have two eyes, gills, several fins *etc.*

For H3 and H4 (below) and related hypotheses, we have, in spite of initially promising support, ultimately collected much disconfirming evidence:

- H3: All swans are white.
- H4: All pigs are pink.

On this basis, we might propose the following norm:

N1: Draw inductions based on structural anatomy, but not colour. N1 is of course fallible, just as H1 and H2 are, and will probably have to be revised or refined in relation to future observation. That may not seem to render it superfluous, given our current state of knowledge, in relation to future hypothesis formation. However, the appearance is deceptive – at least unless we are going to make some further assumptions. H1 and H2 might after all turn out to be faulty, to a greater or lesser extent. All we can say, given our basic naturalistic assumption, is that, through trial and error, gradual convergence on truth (or whatever) will obtain. In other words: to what extent H1 and H2 and similar hypotheses stand the test of time is only something time will tell (to put it in what I hope is an illuminatingly tautological manner!). The same thus goes for N1: it is in no worse, but also no better position than H1 and H2. Since there is, *ex hypothesi*, no a priori way of delimiting the correct set of inductions, but only further observation, the norm can – in the absence of further assumptions – play no substantive role in this delimitation. Trivially, since norms are derived inductively, and we are assuming such a process is reliable in the long run, we must assume the process will converge on correct norms. But this assumption rides piggy-back on the assumption that the process will converge on the correct hypotheses. Therefore the norms are unnecessary – an inessential overlay, a mere summary description of where we have to get to so far.

I said above that the norms can be sufficient, though not necessary for forming good hypotheses. Now I am saying they can play no substantive role. But it might be unclear that these amount to the same thing. Well, we can allow that N1 will, if used, give us the best possible hypotheses we could rationally postulate at that time, so in that sense it is sufficient. However – again in the absence of further assumptions – ‘grue-like’¹² norms such as the following would also be sufficient:

N1': Draw inductions based on structural anatomy up to 31 December 2010, on colour thereafter.

N1'': Draw inductions based on structural anatomy for birds, mammals and fish, but on colour for reptiles and insects.

It cannot on Mill's account count against a norm that it is intuitively invalid (I take it N1' and N1'' are): we have to find this out empirically. At the same time, we will find out, presumably, that hypotheses like H3 and H4 fare no better in 2011 than they have done hitherto, and that those like H1 and H2 continue to hold, and to hold across all different species of animal. Nevertheless, until we make the relevant observations, N1' and N1'' would be as good guides to hypothesis-formation as N1. I think this should make it clear that saying a norm is sufficient but not necessary for forming rational beliefs really does amount to saying the same as that it plays no substantive role, given the above conception of norms in its bare form.

I should indicate, in relation to these ‘grue-like’ norms, that first-order hypotheses might in principle also take such a form. The point, however, is that the indeterminacy in the norms will always go beyond any that exists at the level of the hypotheses. In other words, though we might legitimately balk at grue-like hypotheses given the assumption that induction is basically a reliable guide to reality, we cannot use that same assumption to balk at grue-like norms.

Having said that, I think we are now in a position to understand why we feel strongly inclined to think that steering by higher level norms, of the kind Mill propounds, *must* have some kind of useful role to play in scientific progress. Though N1' might be just as rational as N1 in the absence of further assumptions, there are, in the usual case, precisely further assumptions in play (albeit perhaps at a tacit level) – namely, that N1, and other ‘valid’ norms, correspond to what we take to be deeper, underlying laws or principles of nature. If one assumes – and, let us allow, assumes correctly – that nature is governed by some very general and simple underlying principles, and that such principles are the basis of more specific laws of nature, then one can see how hypotheses about the former would be helpful in formulating hypotheses about the latter. If the world really does partake of deeper levels of organisation, then knowledge, albeit partial and fallible, of what these are will clearly expedite our understanding of more superficial levels. A more specific suggestion along these lines might be that one can justifiably assume that the world is carved up into certain *natural kinds*. If this were the case, then hypotheses about which categories of thing are natural kinds would, if correct, clearly expedite research.¹³ If anti-foundational inductive

norms correspond to or are backed up by correct assumptions about what the relevant natural kinds are, or about other fundamental principles of nature, they will therefore also be in a position to expedite the progress of science. *However* – and this is the rub as far as we are concerned – they will still not be necessary to do this, since the substantive, factual assumptions about the nature of the world would themselves be sufficient to do it insofar as scientists, perforce will take heed of these in their reasoning.

Alternatively, if one could justify the validity of foundational epistemic norms – in the strong sense of TE, or the weak sense of AP and P – prescribing the forming of, say, simple and general hypotheses, one could perhaps also understand how certain kinds of AF-norms could play a role in the rational unfolding of science. This would obviously blur the distinction between AF and other alternative conceptions of epistemic norm. More significantly, given the foundational norms would already be in place, and that the AF-norms would only be justified insofar as they corresponded to these, the former would again be clearly unnecessary for this rational unfolding. In other words, we can only make sense of how AF-norms, as Mill conceives of these, could play a substantive role in the rational progression of science if we assume they are backed up by scientific statements or by (quasi-)foundational norms. But either way, the AF-norms would also thereby be shown to be unnecessary: the need for something – a statement – that is both specifically anti-foundational and normative would drop out.¹⁴

One might perhaps try to avoid this conclusion by suggesting that putting forward norms in the way Mill envisages may in general be conducive to expediting the search for truth (say); that one does not need to assume *particular* kinds of structures underlie our hypotheses in order to facilitate the derivation of true varieties of the latter, but only that there are such structures at all. Hence anti-foundational norms have a life of their own, independent of foundational ones, as well as of substantive scientific hypotheses.

There is something right about the implicit premises in this argument insofar as we would not regard the norms proposed in the course of scientific study as *ipso facto* the right ones, any more than the hypotheses we propose are thus regarded. They are, to reiterate, fallible. However, this does not alter the fact that to the extent such norms are correct (or appropriate), we need some account of what

this correctness consists in – of how they can facilitate progress – and that is what I am suggesting, AF by itself has no resources to give. The problem I am suggesting for AF is, in other words, not simply that it is not foundational – that it does not involve the idea that the standards for enquiry are in some sense given prior to that enquiry. That would of course just be question-begging. What I am asking for is an account of what substantive contribution an AF-norm can make to the progress of science. Again: norms (for Mill) are put forward on the basis of which hypotheses have been most successful, that is, they have an inductive base, and this is the same kind of base as the hypotheses themselves. But then, if these norms do not correspond to deeper principles of nature, or else to foundational norms grounded in some other manner, how can they really make an independent contribution to progress? Even if one can readily make sense of the idea of refining norms in just the same way one refines hypotheses – of finding some inadequate in the face of observation while others go on to live another day – one should not suppose that one has thereby also made sense of the need for such norms; explained why we should not do just as well in sticking to the hypotheses and refining them without regard to any steering but that provided by observation.

One should also note that the foregoing is not intended to show that there is no *some sense or other* in which it might be acceptable to say that anti-foundational norms are needed in rational science. In some sense, high-level or deep scientific hypotheses might reasonably be described as 'norms', insofar as they function as framework assumptions for any more particular theory.¹⁵ Indeed, it is trivial to render such high-level hypotheses as (genuine) normative statements: if *P* is a high-level, well-entrenched principle of science (e.g. the principle of relativity), then *take heed of P in forming one's theories* will be a norm of science, *ipso facto*.¹⁶ However, none of these compromises the underlying significance of my position. The suggested norm is trivial precisely because the fact that *P* is accepted will in and of itself lead to its being taken account of in theorising of the relevant kinds; we do not need the norm in addition to the factual statement. Moreover, that one might regard the statement itself as a norm is nugatory in relation to my project: I am concerned with the prospects for *substantive* varieties of NE, not just the establishment of a generally naturalistic approach to epistemology in which one

eschews TE and adopts the background anti-foundationalism articulated in Quine's metaphor of the ever-evolving web of belief (cf. Sections 1.5 and 2.1). (I will return to the significance of the *apparently* essential role of norms in science towards the end of this chapter.)

Another possible response to my argument might be to make play, as was attempted in Chapter 3, with a distinction between those who derive the norms and those who might have use for them. Thus, one might think that a norm derived in the Millian manner could be of use to a layperson, or to those at a less-advanced level of scientific understanding, without their having to be aware of the factual principle (or the foundational norm) underlying it. I think this is correct, but it does not compromise my point. For of course, given what I have argued above, if those at the forefront of research are going rationally to be able to sanction a norm, then they will do so in relation to a substantive piece of descriptive knowledge, or else some (quasi-)foundational norm. Hence they could just as well communicate this to their less-advanced colleagues as the norm.

Before closing this section, I should perhaps make it clear why it is I say that factual statements can simply replace the norms that they correspond to in making scientific inferences. Though I think this should be obvious given our background assumptions, some might wonder whether formulating explicit norms might not in fact be necessary to get one to draw conclusions at all. For how, they might say, can one draw a conclusion – make a judgement that such and such is *correct*, or *should* be concluded – on the basis of just beliefs to the effect that so and so is the case? In answer to this, we must bear in mind that no supporter of AF (or NF, or perhaps anyone) can assume that the norms derived could be sufficient on their own, without the use of any other abilities or knowledge on the part of those using them, in order to draw epistemically responsible conclusions. AF-norms, whatever else they are, are definitively *not* algorithms, that is, completely specified mechanical procedures for arriving at ever better beliefs. They must be derived and applied with the aid of some kind of ability in conjunction with other items of knowledge, the correct use of which they do not fully specify themselves (*viz.* the indeterminacy problem discussed above – and what would one do if one observed a swan with three wings?). This shows that the activity of using descriptive items of knowledge in order to support or justify other bits of knowledge is in any case something which a follower of

AF must assume; indeed, it is part of the fundamental naturalistic assumption that allows us to see AF-norms as genuinely normative.

Note also that nothing of what I say turns on assuming any kind of infallibility in this ability to use descriptive knowledge. It seems clear that individuals will vary in the extent to which they connect up items of knowledge and draw inferences in a felicitous manner. Indeed, it is plausibly only in the context of a wider community of thinkers that the holistic constraints inherent in the web of belief will be optimally respected and exploited (cf. Sections 1.4 and 2.1). But introducing normative inductive canons will not ameliorate this situation, for they too can be applied more or less felicitously depending on one's grasp of the relevant background knowledge. In the light of this, the question again becomes why one needs them in addition to the factual statements they must be based on (given they are not based on quasi-foundational norms).

4.2 Evolutionary epistemology

In this section, I will briefly consider a version of AF that I call 'EE' – for 'evolutionary epistemology', though understood here in a rather specific sense. Many different things have been meant by the phrase 'evolutionary epistemology' in the past few decades. However, as I will be understanding it, the central theme is that our knowledge is evolving in a way *analogous* to the way biological species develop. In particular, our beliefs vary, and some mechanism ensures that, over time, the ones that best fulfill our purposes – that further our survival and prosperity – are selected and retained. This mechanism is not natural selection, but something analogous that applies to the creation of knowledge, pruning this as a function, in the first instance, of the pragmatic success of the different claims.¹⁷ If this process of knowledge-creation is steered by epistemic norms which are also under the control of epistemological evolution, then we could presumably make sense of a theory of such norms as a version of AF, insofar as it need have no obvious end-point. This is what I am calling EE.

There are many objections that could be made against EE, concerning such matters as the extent and nature of the parallel between biological and epistemological evolution, as well as the question of whether pragmatic success has any intrinsic connection to truth or

rational growth of knowledge. We will be concerned with the last of these topics in our discussion of biological evolution (cf. Section 5.2), a discussion which could also form the basis of an objection to EE. Here, however I will simply assume there is some significant connection between pragmatic success and approaching truth (or however we characterise epistemic progress). For, insofar as EE is a variety of AF, it seems to me it suffers from failings that essentially parallel those that undermine the Millian version thereof.

Let us then suppose that there is some selection mechanism that steers the growth of our knowledge towards ever greater 'fitness', and, thereby, ever nearer the truth. What would be the role of norms in steering this process, in addition to that offered by the pragmatic feedback operating on the knowledge claims themselves? What special contribution would such norms make to our knowledge becoming ever more adequate? If they too are selected in relation to, and gain their normative force from, what are the most fecund effects, it seems unnecessary to take heed of them (cf. the discussion of Mill above). Of course, one might insist that knowledge will grow more rapidly and better with some kind of 'higher-level' steering; but, again, unless one sees this steering as based on higher level or deeper hypotheses about the world, or else some other kind of norm, it seems one will have no way of grounding this claim.

The rest of the discussion of the Millian view would also apply, *mutatis mutandis*, to EE. A rejoinder that might seem to have particular application to EE is that limiting knowledge claims by using overarching norms will be practically necessary to avoid over-proliferation of knowledge claims, that is, to obtain a manageable set of such claims to test and select from. To this I would reply that if such limits are simply imposed, without reference to pragmatic or observational or any other kind of feedback, they will not represent genuine epistemic norms. If the limits are imposed in relation to some kind of rational feedback, they may represent genuine norms, but one would then want to know in what way they are necessary, given that this feedback will also steer the creation of hypotheses, that is, knowledge claims. Some might want to argue that we are – due to our inherently finite and limited capacities – doomed always to focus on just a subset of the hypotheses that might conceivably be relevant to finding the best theory of some part or aspect of the world; and that some kind of general principle must therefore be used for deciding which

to focus on. But, if that is so, it represents a limit on the rationality of science; it does not legitimate any idea of superordinate and genuine epistemic norms.

The core of the problem with both varieties of AF so far presented is essentially just this. In order for the norms proposed to have genuine normative force, as saying how one ought to form beliefs, one has to presuppose that the scientific process in itself is already in good working order. If one then insists the AF-norms can somehow supplement or give extra guidance to this process, one must presuppose the existence of corresponding facts and/or foundational norms to back up the AF-norms. But either way, the need for specific AF-norms drops away.

4.3 MREAF

In Chapter 2, AF was presented as a theory making use of MRE. As suggested there, the link between MRE and AF is probably the weakest of that of the three main forms of NE I distinguished. Indeed, as I hope will emerge from the following, the idea that MRE might serve as a way of deriving norms, given a fundamentally anti-foundational background is not a very promising way of defending AF. On the other hand, the link seems to be there in principle, and might even seem attractive, particularly as a reaction to the failures of the other theories that build on MRE (i.e. AP and P).

The idea behind MREAF (as I will call it), is, as with all other forms of NE that build on MRE, to give a theory of norms in relation to actual concrete judgements, aiming gradually to bring the judgements and the norms in line with one another. With MREAF, however, the process is not seen as having a well-defined end-point, a subject-matter to reveal, as it is when we see MRE as a method in the service of AP or P. Experience continuously affects the kinds of concrete judgement we are disposed to make, so that the project of arriving at a set of norms repeatedly destabilises. But this does not, at least obviously, render the derivation and use of general norms in this manner otiose.

But is it all the same? To begin with, we need to ask what the point of setting up explicit norms in this way is meant to be when there is confessed to be no final set of norms which MRE will converge on, no 'epistemological object' whose structure one is seeking to chart.

For there to be any point here it seems we must assume that there is some kind of never-ending progress in the MRE-process. But how should we make sense of such progress? According to our fundamental naturalistic assumption as it applies to AF, our basic cognitive practices keep our beliefs basically on track – ushers them ever onwards towards the truth (or whatever). However, I cannot see that this can uphold the idea of MRE making progress indefinitely. Indeed, it seems that the only rational reaction to uncovering constantly changing patterns in the norms that seem to be underlying our reasoning, in the manner outlined in the previous paragraph, would be to suppose that there really is no worthwhile pattern to be discovered at all. In other words, a process of MRE that constantly destabilised would not be one that had any rational role to play in science; the fleeting norms it discovered could, and probably should, just as well be ignored.

Perhaps someone might object that we can see MRE as making progress by seeing scientists as possessing several successive stable abilities, old ones being replaced by new ones in the face of experience (abilities which are then seen as existing alongside the basic cognitive practice that keeps our beliefs progressing). Since these underlie ever more successful sets of beliefs (*ex hypothesi*), deriving their structure would have a point.

However, I think it is very unclear how one is meant to make sense of the talk of successive abilities here. We have, in Chapter 3, allowed that a lay ability may be 'upgraded' to a scientific ability, but that is not what we are considering here; rather, it is the idea that the ability in any individual could be improved upon indefinitely through experience. However, if by experience we simply mean learning new facts, it is unclear how this could eventuate in a new *ability*. Moreover, since it is only the former – the learning of new facts – that we are assuming is somehow rationally on track, it is unclear what other kind of 'learning' could allow us to see a development in the norms derived on its basis as making progress.

I think these are trenchant objections, and show that MREAF is fundamentally misguided. However, even if one allowed that the process of constantly refining norms in the manner of MREAF made progress, this would not show that one needed those norms. To begin with, if we are talking about a succession of abilities, what I called the *fundamental problem* for AP (Section 3.3) would presumably

resurface: the scientist would not need the norms she gleaned, for she could just as well apply her ability more assiduously. It might be retorted that the change in ability might be seen as occurring continuously, so that there would never be a point at which the norms were a simple reflection of an ability at a given time. In saying this, it seems the supporter of MREAF would have essentially left behind talk of successive abilities. Whether this is the case or not, it would still be unclear why the norms would play any substantive role, since the progress in abilities would then presumably suffice to make the judgements optimal, without the aid of the norms. Again one might seek to ground the latter in relation to essentially non-normative and/or foundational factors, but then one concedes that specifically AF-norms are not necessary.

At this point, it might occur to someone to apply an argumentative strategy from the previous chapter (also mentioned in Section 4.1, above). Thus, even though scientists may not have any use for the norms that could be gleaned at any stage in the progression of their thinking, might not ordinary people, or less-advanced scientists, have use for these? After all, on the present assumptions, scientists will be gaining an ever-better understanding of the correct procedures for understanding the natural world, even though this will never be perfected. Against this line, I would appeal to the Kuhnian line on scientific competence, which would suggest that it is not a matter of following rules at any level, even implicitly, which constitutes a scientist's ability (cf. Section 3.4).

Some might recall me saying that I did not want to place great weight on Kuhn's conception of science in my overall argumentation. Be that as it may, the really deep problem with MREAF is that it is highly unclear what deriving norms in the way it envisages could, even in principle, contribute to the rational progress of knowledge – given that the use of MRE is not to tap the structure a stable ability, and given that the progress we can take for granted in relation to AF concerns our beliefs about the world.

4.4 Laudan's normative naturalism

Larry Laudan's theory of, as he calls it, NV can be seen as a substantive version of AF, albeit a rather more complex one than the varieties we have so far considered.¹⁸ Though heralding most directly from studies

in the history and philosophy of science, NN aims to be a genuinely epistemological, or, as Laudan prefers to call it, *metalepistemological* theory.¹⁹ I have chosen to dwell on Laudan's theory for two reasons. Firstly, it offers a fully worked-out conception of epistemic norms as *instrumentalities*, or, as Laudan calls them, *hypothetical imperatives* – that is, as prescriptions that presuppose some ulterior goal, purpose or aim – and this is a conception of norm that I would in any case need to discuss. Secondly, it attempts to give a distinctively anti-foundational conception of the *aims* of science that, in and of itself, would furnish a version of AF. Though Laudan sees NN as a unitary theory, I will, as things go on, be splitting up the discussion in relation to what I see as these two distinct components of it.

For didactic reasons, it is simplest to begin a presentation of NN with the idea of norms as instrumentalities. According to Laudan, all methodological rules in science are in essence hypothetical statements of the form 'If you want *A*, do *M*' where *A* stands for a cognitive aim, goal or value, and *M* some kind of method or procedure which to some extent procures *A*.²⁰ A naturalistically inclined epistemologist will then, according to Laudan, be one who proposes rules of this kind on the basis of investigation into their empirical tenability: the extent to which doing some *M* in fact procures some *A*, in relation to other possible methods. Thus, a particular rule *If you want A1, do M1* will be justified insofar as we believe *M1* leads reliably to *A1* – i.e. in a significantly large number of cases – and also to the extent that there are no other methods that lead to *A1* more reliably (as far as we know). Clearly, the details of uncovering and justifying such rules will be complicated, and epistemology will admit of few clear-cut prescriptions. The empirical basis for this inquiry is for Laudan the history of science, insofar as science has gradually made progress towards scientific goals or aims we today revere; this history will also be subject to interpretational controversies. However, all these complications are fine because our aim as naturalists, says Laudan, is not to ground science on some incontrovertible fundament, but to give an epistemological theory that is in as good working order as the science it serves.

An immediate intuitive reaction to this kind of proposal is to ask how the aims and values of science – our *axiology*, as Laudan puts it – are to be divined and justified. Now, whether justification of the aims of science is in fact necessary to understanding hypothetical norms as genuinely normative – in such a way that they might

furnish a form of NE without further ado – is a good question, and one I will take up presently. For the moment, however, it suffices that Laudan in any case takes up the challenge of justifying aims.²¹ According to him, understanding the justification of aims requires a proper appreciation of the history of science. What this history reveals – again according to Laudan, though also many others – is that the aims of science have changed over time, sometimes dramatically. The scientific aims of Descartes, even in general terms, were not the same as those of Newton, and neither are identical with Einstein's – contrary to what one might have supposed.²² However, Laudan maintains that these changes have been made, not arbitrarily or because of social pressures, but *rationally* in response to factual discoveries and the success or failure of particular research programmes. Theories, methods and aims do not generally change in step with one another, as scientists shift, non-rationally, from one paradigm to the next – as Kuhn, according to Laudan, holds. Rather, they *reticulate* in relation to one another, theories and methodological rules conditioning one another, and both conditioning, as well as being conditioned by, our aims in an ever-evolving, rational unfolding.²³

Laudan also proposes certain general *axiological constraints*, that is, constraints on the rational choice of aims – a *metaxiological level*, if you will. In particular, our scientific aims should not be what he calls *utopian*, in the sense of being demonstrably unrealisable, inherently vague or such that there is, and can be, no criterion for thinking one has realised them.²⁴ In addition, avowed, explicit aims should cohere with the aims implicit in acknowledged scientific practices;²⁵ and, finally, the aims we hold today should be such that they do not render acknowledged scientific achievements of the past 'unscientific' – Newton's work should not, for example, turn out not to have advanced these aims.²⁶ (We will have more to say about these constraints and their significance later on.)

It is also important for Laudan that we can say that, despite the variability in scientists' aims, science itself has made progress. We can say this, according to him, because science has made progress towards what we *today* regard as our scientific aims and values.²⁷ Moreover, this is no accident, for any axiological or methodological proposal is answerable to the canon of acknowledged scientific achievements.²⁸ That canon, as noted above, is necessarily the database for anything that would call itself a theory of science.

If NN were tenable, genuine, naturalistic epistemic norms would thus be vindicated at two distinct levels. First, we would have to countenance norms in the sense of aims or values that are realised through the progress of science, but which cannot be justified in relation to, nor even derived from, the judgements of past scientists – as the dominant research strategy in philosophy of science (DRSPS) would have it they are (cf. Chapter 3). Hence, I would not be able to run an argument of the kind I did against AP, to the effect that norms derived and justified in this way are not necessary to make optimal epistemic progress. Second, we would have to countenance the idea of hypothetical rules linking specified methods to the (justified) aims. However, I will seek to show that NN fails as a substantive theory of methodological rules and of scientific aims, and that therefore the threats it poses to my overarching thesis pall. The chief aim of the present section is to argue against the tenability of Laudan's conception of the aims of science as a form of AF (though some considerations concerning methodological rules will also be mentioned). In the following section, I will then consider the idea of norms as instrumentalities or hypothetical imperatives, considering Laudan's ideas to start with, but gradually opening the debate to encompass wider considerations.²⁹

To set off the critical discussion, I will consider first what for many has seemed the main problem with NN: that it somehow smuggles in, or at least must presuppose, epistemological constraints that are foreign to its naturalistic aspirations. Further, it has been alleged, the normative force attaching to rules and aims within NN must presuppose a form of evaluation that is not naturalistic. There are several versions of this kind of argument; I will concentrate on one which has it that the axiological constraints Laudan suggests for any viable set of aims do not have a naturalistic basis. Thus, according to Laudan, a goal should not be utopian in any of the senses outlined above. The objection is then simply that these seem like *a priori* and conceptual, not empirical constraints on one's axiology.³⁰

Laudan protests against this accusation that it is no goal of NN to deprive epistemology of conceptual resources. Epistemology should be as empirical as science is, but since science also has recourse to conceptual considerations, so may epistemology. Moreover, constraints against utopian aims are not wholly non-empirical insofar as it can be an empirical discovery that an aim is unattainable.³¹

Though these remarks are suggestive, they do not in themselves constitute a satisfactory answer to the charge that NN presupposes non-naturalistic resources. That science uses what might loosely be termed 'conceptual resources' – that it is not a rigidly positivistic enterprise – is well-taken (and in no way conflicts with naturalism in our sense); however, this does not in itself legitimate the *general* use of conceptual resources in a would-be naturalistic epistemological theory. There should at least, for a naturalist, be some *connection* between the conceptual aspects of the latter and those of science more generally. Further, as Harvey Siegel has pointed out, whilst a goal being unrealisable may be a naturalistic, empirical fact, that a goal therefore *ought not* to be pursued seems to be a further, non-naturalistic commitment.³²

Recently, Karyn Freedman has presented a persuasive case that Laudan can turn aside these criticisms with impunity.³³ Referring to various parts of Laudan's writings, she motivates an interpretation of his philosophy on which his more narrowly construed *empirical* naturalism, which underlies his specifically *metamethodological* views, must be understood in relation to a more broadly construed *scientific* naturalism, according to which values, indeed, standards generally, are naturalistically acceptable to the extent they are prevalent in 'the history and practice of science'.³⁴ She supports this interpretation largely through citations from Laudan's *Science and Values*, but I think an even clearer statement of a commitment to scientific naturalism is the following:

The naturalist [...] holds that the same mechanisms that guide the change of aims among scientists can guide the epistemologist's selection of epistemic virtues. There are strong constraints on the aims of science that a scientist (*and thus a naturalist*) can permit.³⁵

This passage strongly suggests that the constraints epistemologists use in evaluating aims are justified in relation to the constraints scientists themselves use. Hence, given this is a form of naturalistic justification, NN remains a thoroughly naturalistic theory.³⁶

Though this might seem an intuitively attractive way for Laudan to defend himself against his critics, it is nevertheless at odds with other aspects of NN. In the above passage, and elsewhere, Laudan

gives the impression that there is a relatively fixed set of constraints for adjudicating and selecting aims amongst scientists, even though the aims themselves will change with changes in theoretical and methodological knowledge. But this fits in badly with his otherwise fully 'reticulated' conception of the relationship between facts, methods and aims, all mutually influencing each other and all susceptible to change over time. In particular, his conception of the criteria for aim-evaluation fits in badly with his insistence that a normative methodology and axiology *cannot* be gleaned from an examination of the aims and methods of particular, historical figures in science, since these aims and methods change over time and need not be those we would espouse today.

I think these anomalies can be amplified to constitute the basis of a fundamental criticism of Laudan's conception of scientific aims and, to a lesser extent, methods. In what follows, I will argue that, given Laudan's conception of scientific progress, it is very hard to see how he can consistently avoid saying, as he wishes to, that a core set of scientific standards – especially at the axiological level – has been held, implicitly or explicitly, by leading scientists over the past few centuries. Moreover, it seems this fact must, for Laudan, play a *justificatory* role in relation to the values we hold today; the price of resisting this is relinquishing the idea that values and aims have genuine normative significance. In other words, if Laudan is to uphold a rationalistic view of scientific norms, he must fall back on the central ideas of the version of AP in which DRSPS is used to derive and justify scientific norms. (This is not to say that DRSPS is apt or its assumption of methodological and axiological unity across history correct, but only that Laudan's account of progress in science must assume something like it not to be relativistic (or something close to this). If the standards of science have thoroughly changed throughout history, this implies that, at least insofar as we want to uphold a picture of science as a rational enterprise, Laudan's account of progress must be jettisoned – presumably in favour of one in which we are far less specific about the sense in which science progresses.)

First, however, we need to go into just a little more detail of Laudan's own terminology. It will not have escaped notice that DRSPS bears a close resemblance to Freedman's notion of scientific naturalism. Though Laudan does not appeal to either idea directly,

he does discuss something that both in turn also have clear affinities with – the *metamethodology thesis* (MMT):

The Metamethodology Thesis: a methodology of science is to be evaluated in terms of its ability to replicate the choices of past scientists as rational.³⁷

According to Laudan, this assumption underlies Lakatos' rationalist methodology of research programmes and also the anti-rationalist critique of method offered by Paul Feyerabend, as well as many other views.³⁸ The basic idea behind MMT is that a proposed methodology – a theory of norms for science – must be answerable to the actions and choices of great scientists of the past. For Feyerabend, MMT entails that there are no methodological rules apart from 'anything goes', since any more substantive rule has been broken at some point or other by some famous scientist or other. For Lakatos, whilst MMT shows that certain, perhaps, intuitively attractive methodological rules – those of e.g. Popper – are untenable, it nevertheless allows us to formulate a more nuanced methodology that can have normative force. Laudan has sympathies with Feyerabend and other relativistically inclined thinkers insofar as he agrees that examining the standards of past scientists offers no grounds for suggesting a monolithic set of methodological norms of the kind Lakatos claims to find. But he does not draw relativistic conclusions from this, rather, he rejects MMT as an adequate metamethodology. Indeed, whether scientists have or have not been rational in the past is, as he sees it, largely irrelevant to methodology; for even if they have been, their cognitive aims are different from ours, as are their background beliefs.³⁹ This means that their actions will be steered by rules that are very different from ours today: whilst our rules will be aimed at putting us in contact with *our* aims in the light of *our* beliefs, theirs will be aimed at putting them in contact with *their* aims in the light of *their* beliefs. Hence use of MMT would be hopeless as a way of finding out what the methodological rules of science should be.

Following this discussion, Laudan introduces the idea of methodological rules as hypothetical imperatives, and the other aspects of NN are sketched out (see above). In the final part of the essay presently under consideration, Laudan takes up the question of the role of the history of science in his conception of methodology, and

of the nature of scientific progress. It appears that, though MMT is eschewed, the history of science still has a central role in epistemology since it is this that provides us with our data for determining the tenability of proposed methodological rules. If one wonders, say, whether we should avoid *ad hoc*-hypotheses in order to obtain reliable theories, one needs to get out of one's armchair and find out, by examining the history of science, whether avoiding *ad hoc*-hypotheses generally has, as a matter of fact, led to reliable theories, and whether other methods have been more or less successful in this regard.⁴⁰ Finally, though aims and methods have changed throughout the history of science, we can say that science has made progress, because it has made progress towards ends that we today value:

We take science seriously precisely because it has promoted ends which we find cognitively important. More than that, it has become progressively more successful as time goes by. [...] Science in our time is better (by our lights of course) than it was 100 years ago, and the science of that time represented progress (again by our lights) compared with its state a century earlier.⁴¹

This is so even though the explicit aims of scientists, and therefore many of their methods, have changed throughout history. For, as Laudan puts it, 'unlike rationality, progress need not be an agent-specific notion'.⁴² The course of science has involved many episodes with unintended consequences that only later have been seen as valuable. But because we now do value them, we can use the history of science as evidence for our methodological rules, and at the same time see this history as making progress towards our values and aims.

An immediate question that might arise here is *why* our aims today should be such that the science has made progress towards them. Though I think this is ultimately a good question (as we shall see below), Laudan would not be perturbed by it insofar as he holds that the history of science also must serve as a '*certifier or decertifier for new proposals about the aims of science*'.⁴³ A scientific axiology cannot diverge radically from the aims achieved by exemplary work in the history of science, such as Newton's *Principia*.⁴⁴ The history of science carries this weight because it is in some sense definitive of what science is,⁴⁵ even though the varying aims and methods of scientists are not.

We are now ready to see why Laudan's overall position is inherently unstable and untenable. Central to his move of rejecting MMT at the same time as viewing science as making progress is the idea that individual scientists' aims and methods can diverge from those of science itself. We are told that though *science* has made progress towards aims we hold dear, and also used methods appropriate for the achievement of these aims, the aims and methods used by individual *scientists* are of little interest to epistemology. It would thus seem that Laudan holds that throughout the history of science, there has been significant divergence of scientists' aims and methods from those of science itself. But really how plausible is the idea that there might have been such a divergence? We can certainly allow that an early scientific episode might have been undertaken using methods and aims that are quite removed from those we take to be best now, and that it nevertheless, without any intention to do so on the part of those undertaking it, made progress towards these latter aims. The example of Cartesian optics, cited by Laudan,⁴⁶ is perhaps an example. But could such a pattern have been maintained *throughout the whole of the history of modern empirical science*? The force of this rhetorical question comes through most clearly in relation to aims; we will turn to its bearing on methods briefly towards the end of this section.

Remember that, for Laudan, the history of science has been *progressive*; that is, it has moved *gradually, step-wise* towards aims we today value.⁴⁷ He wants to understand this progress as having occurred *without* important and influential scientists within this historical progression having embraced the aims that we hold today; *without* them having designed their research programmes such that these aims would (to some extent) be achieved. But does not such a systematic movement towards something relatively stable (our values today) demand explanation in terms of some *vehicle* for the change? And if this vehicle is not the scientists themselves, what else should it be? Of course, it is not *impossible* that progress towards certain aims could have taken place in the history of science without any deliberate steering 'from within', so to speak; but it would nevertheless be a coincidence of cosmic order – a miracle that in, say, the history of physics stretching from Galileo, Descartes and Newton through to Maxwell, Einstein and Bohr there should have been a progression towards fundamental goals that these protagonists did

not share with one another, at least tacitly or implicitly. To put it another way: if science has made step-wise progress, as Laudan maintains it has, towards certain goals, then surely these goals would also have been avowed by scientists 10, 50, 100 and 200 years ago, as well as by us today. It is thus not just 'by our lights', as Laudan holds, that science 50 years ago was better than science 100 years ago, and that science 100 years ago better than science 200 years ago, but *also by the lights of those who conducted research at those times* – and at any other time at which progress had been made, relative to some earlier stage, towards the goals we hold dear today. If this were not the case – if, say, Maxwell need not be viewed as having held (at least implicitly) central aims of physics shared by his predecessors and by us today – then the progress Maxwell himself effected in physics must have been a complete coincidence. Moreover, there must have been many such coincidences in the history of science, at each stage at which progress *de facto* was made without scientists actually conducting research with the aim of achieving the goals they all, to some extent, contributed to achieving. Such a string of complete coincidences could only be described as miraculous. But to see the progress of science towards our aims today as *miraculous* is absurd. Hence Laudan's position must be rejected.

If this line of thought is correct, then MMT, which bases its theory of scientific standards on the actions of individual scientists of significance in the past, must be far more plausible than Laudan suggests it is, given his conception of progress – at least as far as axiology goes. And if Laudan must concede that the distance between his own and other extant positions diminishes dramatically, Laudan would no doubt have many things to say in response to this objection. Most obviously, he would point out that major scientists have, *de facto*, espoused quite different aims from one another. In response to this, a supporter of MMT could point out that her principle does not assume either (a) that all or most important scientists have possessed *all* the aims that science has made progress towards, or (b) that *all* aims of great or important scientists are of interest to the epistemologist of science. *Apropos* point (a), we may note that insofar as there is a plurality of scientific disciplines and sub-disciplines, there may well be a plurality of scientific goals; in addition, within a given discipline, there may be a diversity of aims amongst its practitioners insofar as some will be more interested in

theory, others in observation and experimentation, and others again in mathematical aspects. We can perhaps concede that really great scientists can be required to have some overarching conception of how the different aspects of their discipline relate to one another. But in any case, insofar as it is possible to subdivide science and scientific activity in the ways mentioned, a supporter of MMT will clearly also be able to embrace any diversity in aims this entails in her theory of scientific rationality.

Apropos point (b), one should bear in mind that not all the theoretical endeavours of the great scientists have been successful, and thus neither will the aims that these projects sought to achieve be significant today. This is compatible with anything a supporter of MMT would want to hold: she is not interested in the values of a scientist that steered research activity that did not lead to good science – any more than she is interested in those steering the scientist's non-scientific activity. A related point, due to Laudan himself (see the introductory presentation of his position), is that a supporter of MMT is not (just) interested in explicit pronouncements from scientists as to their aims, insofar as there may be a mismatch between such explicit aims and the *implicit* aims that actually guide fruitful work. Note also that we can allow that some scientists may care more about professional prestige – or their wallet – than they do about the nature of the universe, even though they make substantial contributions to understanding the latter. For it seems clear that most cannot be like that, and that those who are only make worthwhile contributions insofar as *others* explicitly ensure that the procurement of non-scientific aims coincides with the furthering of scientific ones. All these sources of possible divergence in aims between scientists may be acknowledged – and legitimately idealised away from in formulating a database for an axiological theory. What Laudan wants to say, by contrast, is that we should be no more interested in the aims of scientists that *did* lead to good science than in those which did not. That view presupposes that science in the past might generally have been undertaken by people with aims quite other than those in relation to which we regard it as having made progress. I have argued that this latter idea makes little sense.

Laudan would probably insist that even if we restrict ourselves to good science, we still find divergences from the aims we today hold to be important for science. But the protest would not be convincing.

Ad hominem, it can be pointed out that many of Laudan's examples of aim-divergence are taken from the early stages of modern science, when we can allow there was some unclarity as to what science was, and what it was seeking to achieve; as science becomes more and more established, examples of divergence in aims become less common and less obvious.⁴⁸ Moreover, even if examples of good scientists espousing very different goals could be found, the very fact of scientific progression, as Laudan conceives of it, forces us, on pain of having to posit cosmic coincidences, into finding *some* mechanism in the scientific community, and not just science *qua* body of theoretical knowledge, for steering theories towards the aims we value today – or else to drop the conception of progress Laudan assumes altogether. In sum: if we see the aims of science as we have understood them today as things towards which science has made progress over the centuries (as Laudan does), we are forced to accept that these aims will also have been held by the great scientists throughout history. Insofar, MMT must be admitted to be a viable method for doing scientific axiology, that is, for yielding a theory of what the aims of science are.

Could Laudan ultimately accept this, yet claim it to be insignificant? Perhaps he might confess that past scientists have held aims similar to those we hold today, but deny that MMT has any role to play in *justifying* aims – ours today, or any others – in the way many of its supporters have taken it to be capable of. Consider the following passage:

In my view, the chief aim of the methodological enterprise is to discover the most effective strategies for investigating the natural world. That search may or may not involve us in articulating the criteria of evaluation used by past scientists. But the latter task is, at best, a means to an end rather than an end in itself.⁴⁹

Here we see a backing off from the more extreme statements about past scientists having aims and methods very different from ours today. But moreover, it also suggests, as a rebuff to my foregoing criticisms, that when it comes to what is *constitutive* of being 'criteria of evaluation', who has held them is beside the point.

Concerning methods or rules, this point has some force, as we shall shortly see. However, when it comes to aims, it is very unclear how one can understand the aims we hold today as *justified* if one

does not see them as being derived on the basis of the aims of important scientists of the distant and not-so-distant past. In this connection, recall first that Laudan seems to need to resort to what Freedman called *scientific naturalism* to defend the idea that his *constraints* on axiology are naturalistic. This idea is precisely that these constraints are justified insofar as they are prevalent in science, that is prevalent in the judgements made by important scientists, now and throughout history.⁵⁰ Now, it does not *follow* from this that the aims we today hold will be justified in a similar manner, if only because any given axiology will be underdetermined relative to these constraints. Nevertheless, one can wonder why, given one's axiological constraints will be justified in this way, one's actual axiology should or could not be justified similarly (given this axiology has persisted over time, as we have showed it must). Scientific naturalism seems apt to apply not just to axiological *constraints*, but to the aims and values themselves. Indeed – and moreover – in this latter role, scientific naturalism becomes indistinguishable from MMT, assuming the latter has justificational aspirations (something which we can in turn identify with DRSPS).

As also noted in the discussion of Freedman, the idea of scientific naturalism she offers to rebut Laudan's critics runs counter to the spirit of his reticulated model of scientific change, in which no elements are fixed. Of course, the situation becomes no better for Laudan if he rejects Freedman's helping hand in relation to axiological constraints, for the charges that NN is not purely naturalistic would then be unanswered.⁵¹ The more fundamental point, however, is that whether he accepts this helping hand or not, when it comes to justifying values themselves, Laudan faces a dilemma. On the one hand, he can embrace scientific naturalism in order to justify values and aims; but if he does this, any residual opposition to MMT as a means of justifying these must be given up. (On this line his position collapses into a form of AP, which I have already criticised.) On the other hand, he can insist that, in the final analysis, aims and values are one thing, who holds them quite another; but then it seems he must give up the idea that these aims and values can be naturalistically justified, and the interest of his views to our concerns – as an AF-version of NE – pales.⁵²

Someone might protest at this point that I have mis-characterised the second horn of this alleged dilemma. Laudan's remit, as he sees

it is to give a naturalistic account of science as a rational, progressive enterprise. Given his axiological *constraints* are kosher from such a perspective, there is room for rational criticism of aims, as there is for justification and criticism of methods, seen as hypothetical rules for the attainment of different aims. Science has made progress towards certain aims, not others. MMT no doubt *could* be used, given my arguments above, to gain information about science's 'criteria of evaluation', but it is inessential to this project. Is not this, without saying any more, a coherent and indeed attractive position – and one, moreover, that involves, and necessarily involves, naturalistic epistemic norms?

This question marks a watershed in the discussion in that in answering it we must distinguish clearly between aims, on the one hand, and methods or rules, on the other. As far as aims are concerned, I do not think the above position can allow us to see them as expressing genuinely normative constraints on scientific activity. Laudan's view involves a kind of pluralism about aims: as he himself says: '[o]ne may plausibly propose a new aim for science, even one that has never been explicitly espoused or deliberately sought'.⁵³ Indeed, it seems Laudan must allow this kind of proliferation of postulation of goals if he wants to keep his distance from the fully committed supporter of MMT, that is, one who holds that aims are justified in the manner of scientific naturalism/DRSPS. Though, as I have argued, Laudan must concede that scientists have made progress towards certain *shared* goals over the centuries, this should nevertheless – on the present understanding of his position – be *just* a fact. That past scientists have certain shared aims is not meant to contribute to a justification of them; so it had better also be possible for one to espouse goals other than these.

Laudan goes on to defend this kind of pluralism against a charge of relativism by construing the latter as the idea, not (just) that there is a plurality of standards, but rather that scientific theories, methods and aims can be given no rational grounding or be subjected to rational criticism.⁵⁴ We do not need to go into definitional issues, however. For even if relativism does essentially involve a form of nihilism about rationality (as many, by my lights rightly, would hold), this does not show that the kind of pluralism Laudan espouses does not also have, if not nihilistic, then at least quite devastating consequences for the idea of rational inquiry. And indeed I think it

does. For if the only constraints on an axiology are that the aims be anti-utopian, that they cohere with implicit practice and that they should have been realised at some point in the history of science, we surely open for a free proliferation of axiologies that in the end could only be seen as non-rational. Or to put it another way: if these are the only rational constraints on our aims and values, then any *further* specification of them will be quite arbitrary. This is not quite the same as 'anything goes', but it does mean that it is only to the extent that we find a resonance in the judgements of past scientists – reflected, *ex hypothesi*, in the axiological constraints – that our values will have any normative significance. The rest will be up for grabs. We will thus be unable to give any meaning to the idea that some or other of these aims are things we *should* pursue; whilst to the extent we can say that some or other aim should *not* be pursued, because it breaches one or other constraint, we do so in relation to something that is inherent in the practice of scientists and thus does not need to be made explicit as a norm.⁵⁵

This still leaves the question of whether NN, or at least something like it, would be viable given this quasi-relativistic (or even a fully relativistic) view of aims, coupled with a *rationalistic* view of rules, that is of the means for attaining those aims. Though I am open to suggestions to the contrary, as far as I can see, given our background naturalistic assumptions, I cannot see that anything I have said so far would rule this out. Given just our background naturalistic assumptions, it seems we should be open to the idea that, though particular goals conceived of in the course of scientific activity may not themselves have normative significance, nevertheless, specifying means for obtaining these might do.

At this juncture, we should also bear in mind versions of AP (discussed in Section 3.5) which would see our aims as justified through DRSPS – or even just MRE operating on ordinary people's judgements – but which see rules or methods as justified in the manner assumed by Laudan (i.e. by reference to empirical evidence for the relevant connection). Though, by my arguments of Chapter 3, we would not need to set up these aims as explicit epistemic norms, there would still seem to be a point to making explicit the instrumental rules for attaining these. The latter would also be the case for a *mixed* kind of position which saw the aims as justified apriori, but the means for attaining these as establishable empirically (cf. Section 5.1).

The following section is devoted to an evaluation of the idea of norms as instrumentalities or hypothetical imperatives. As a warm-up to that (and for the sake of completeness of the discussion of Laudan), I will close the present section by considering the implications of my arguments above for Laudan's view of MMT as it applies to methodological rules (so far we have concentrated mainly on MMT in relation to aims).

As noted above, Laudan criticises MMT on the grounds that the methods used by past scientists will be suited to furthering *their* aims in the light of *their* background beliefs. However, since aims must, in the light of my arguments, be seen as largely shared by scientists across historical epochs, it seems Laudan must concede that the methods of past scientists have some degree of viability. Indeed, whether or not they have shared our aims, it seems their methods will, at least to an extent, have such viability. For according to Laudan, regardless of whether past scientists shared our aims, their science has made *de facto* progress towards these aims. But if it has done this, the methods they espoused that led to these were good ones – there being nothing more for a method being good to consist in (for Laudan) than it leading to some desirable end.⁵⁶

On the other, Laudan could plausibly claim that it is not the fact that rules were held by important scientists that *justifies* them. Indeed, it seems perfectly consistent with his overall position to suppose that a metatheoretical empirical enquiry into which means are best suited to achieving certain ends could go beyond the history of science and be pursued as an extant empirical enquiry today (even though Laudan himself does not suggest this).⁵⁷ Such an enquiry might reveal that whereas the methods used by scientists in the past were to an extent apt to promote certain ends, they were not the best methods in this respect. Differences in background beliefs could presumably also play a role in explaining why past scientists' methods were not necessarily optimal.

In this section, I have argued that Laudan's idea that scientists' epistemic values have changed radically throughout the history of science is highly implausible, given his conception of progress. Further, his meta-axiology faces a fatal dilemma: either it must conform fully to the ideas behind MMT – i.e., for us, a version of AP – which is something Laudan is precisely concerned to present NN as an alternative to; or be classed as anti-rationalistic. As for his conception of

methodological rules, on the basis of what Laudan actually says, it seems there will again be no great mismatch between the methods used by great scientists of the past and those that are actually conducive to the attainment of the aims we cherish today. However, this match does not obviously serve to justify these methods, and it seems that one could consistently hold that the methods of past scientists are not necessarily the best ones. This might suggest that, as far as methods go, NN still has something going for it.

4.5 Epistemic norms as hypothetical imperatives

Laudan holds that methodological rules are hypothetical imperatives or instrumentalities. They do not say 'Do X!', but rather 'If you want Y, do X!'. This conception of norms might take it that Y is universally and inherently valuable – as the relevant version of AP (see Section 3.5) – but need not, as in Laudan's NN. In any case, the norms are instrumental or hypothetical in that they presuppose an ulterior goal, on whose valuing the normative guidance they offer is conditional. Note, then, that they are not just what one might call *conditional* norms. Examples of merely conditional norms might be *If you believe p&q, you should also believe p*, or *If you are viewing a very weak light source at a great distance, use peripheral rather than foveal vision*. However, these are not instrumental or hypothetical norms, for they do not express an ulterior goal and a means for achieving this, but rather specify epistemic circumstances under which certain inferences or cognitive practices are prescribed or licensed. (Having said that, the second, at least, is probably best understood as an elided hypothetical norm, conditional on a goal of something like *gaining maximum information*.)⁵⁸

The instrumental conception of epistemic norms is quite widespread amongst naturalists today, at least at a tacit level, whether they are seen as falling into the camp of AP, AF or some less strongly naturalistic kind of epistemology than I am chiefly concerned with in this book.⁵⁹ It is arguable that the conception fits most naturally with AF, since the modern idea of instrumental rationality underlying it is one on which goals are not seen as fixed or given, but as essentially mutable. However, though in a way I am assuming this connection in the present section (it being part of a chapter on AF), what I have to say is meant to apply to other theories that assume fixed goals but also espouse an instrumental conception of norms (cf. Sections 3.5 and 5.1).

As noted in previous chapters, Quine's later writings also promote an instrumental notion of rationality as part of naturalised epistemology. Here is the relevant and oft-cited passage:

Naturalisation of epistemology does not jettison the normative and settle for indiscriminate description of ongoing practices. For me normative epistemology is a branch of engineering. It is the technology of truth-seeking or, in a more cautiously epistemological term, prediction. Like any technology, it makes free use of whatever scientific findings may suit its purpose. There is no question here of ultimate value, as in morals; it is a matter of efficacy for an ulterior end [...].⁶⁰

Here Quine would seem to espousing some kind of Af, although the assumption of truth or prediction as the obvious end for naturalistic engineering might seem to suggest a whiff of something less anti-foundational. Since Quine's attitude to truth tends towards minimalism, the issue is complex, and I shall not pursue it further here.⁶¹ The point for us is that there seems to be a conception of naturalistic epistemic norm that does not depend on any of the ideas hitherto discussed, is not psychological, and seems to offer genuine normative guidance. I will be suggesting that, nevertheless, such norms are not necessary for optimal belief-formation.

Since we are already well-acquainted with Laudan's views, and since he is, as far as I can see, as good a representative of the instrumentalist camp as any, the critique will take as point of departure a further objection to NN, made by Alexander Rosenberg. Rosenberg objects to NN to the effect that, on Laudan's conception of the relationship between methodological rules and factual theories, it would seem – contrary to Laudan's idea of mutual 'reticulation' between rules and facts – that the latter are in fact more fundamental, since they explain why the former work or fail to work, but not *vice versa*.⁶² Laudan has replied to this that, intuitively, rules can, if they are justified, explain why theories work, and that otherwise the question of which comes first, rules or theories, is like the proverbial chicken-and-egg problem.⁶³

I think Laudan is right that rules, as he conceives of them, can explain why theories are chosen. However, this does not mean there is no meaningful issue as to which is prior. Moreover, a related but simpler point than Rosenberg's critique is just this: since hypothetical

rules are justified on the basis of empirical connections between means and ends, why not just stick with claims about these empirical connections and forget the rules? Knowledge of such connections will be at least just as useful as knowledge of the rules in guiding action towards the attainment of one's goals. Indeed, rational action in the instrumental sense in question here is not generally conceived of as involving rules – at least beyond the background rules of decision theory (if such exist).⁶⁴ I desire certain things, have beliefs about what will most likely lead to these things, and I act accordingly; I do not need to formulate a rule (additional to those of decision theory) to the effect *If you want such and such, do so and so in order to achieve this*. So the question is, given hypothetical rules will be justified by reference to ordinary, fact-stating empirical beliefs, why bother with the former? If one comes to the conclusion that they indeed are not needed, then, taking this together with the critique of Laudan's axiological proposals in Section 4.4, very little of substance would remain of NN. Moreover, we will have shown – for yet another distinctive class of NE – that the envisaged epistemic norms are not in fact necessary to form optimally rational beliefs.⁶⁵

How might Laudan – or instrumentalists generally – react to this criticism? Perhaps by saying that a hypothetical imperative could be in force in a given scientific community without all the members having access to the empirical reasons behind it being in force – rather as ordinary citizens are required to obey society's laws even though they are partially ignorant as to why they should obey them (assuming they should). But what precisely does it mean to say scientists would not have 'access to the empirical reasons behind it being in force'? This could mean either (i) they are not acquainted with the empirical claim or claims on which the rule is based; or (ii) they are not acquainted with the detailed empirical evidence used to justify the empirical claim or claims on which the rule is based. The second possibility here seems reasonable. Given the cooperative and communal nature of science, it is clear that scientists will often have to accept empirical claims when they themselves have not gathered the relevant empirical evidence for them, or even are not aware – in detail or outline – of its nature. However, it seems far less clear that possibility (i) would be acceptable in a scientific community. Practitioners are told *If you want such and such (A), do so and so (M)*. Surely, in a scientific community it would only be acceptable to

conform to such a rule if one also were given reasons for it, such as: *M leads to A in the large majority of cases and nothing other than M has been found, so far, to lead to A in as many cases.* At the very least, it seems these purely factual statements would be *assumed*, possibly tacitly, by those who embraced the hypothetical imperative and used it, in a rational manner, to steer their belief-formation. In either case, the rule itself would then be essentially redundant.

But what, someone might say, about non-scientific communities? Presumably an instrumental norm could be accepted by ordinary people without them having to be acquainted, even tacitly, not just with the *grounds* for the empirical claim underlying it, but with the empirical claim itself. Moreover, it might seem that science could precisely play a role here in supplying lay people with epistemic norms for the achievement of their goals – so that we get a situation analogous to the kind envisaged in Chapter 3 where, though perhaps scientists will not need the norms themselves, these will be essential for the optimal epistemic practice of ordinary people.

However, this line of argument is also easy to rebut. The bottom line is again very simple: why not just inform people of the relevant empirical connection, rather than the norm? Given what I argued above about the structure of instrumental reasoning, information to the effect that a certain desired consequence is reliably brought about by a certain kind of action and (let us say) only by this action would be just as effective in expediting optimal epistemic practice as a conditional norm. Perhaps one might say that the normative formulation is more streamlined, and better gets across the essential point. However, there is surely nothing in the way of introducing some standard non-normative formulation that is equally streamlined and succinct, for example of the form 'X is the most reliable way (known) of securing Y'. Moreover, it is not even clear that avoiding the complications of the precise nature of the empirical connection, at least in the total way a hypothetical norm does, is actually conducive to maximally rational behaviour. If the only viable, known way of attaining a given goal is in fact far from being a *very* reliable means to that end, this presumably ought to be taken account of by those concerned to achieve it. But even if one disagreed about this, the idea that norms should be necessary, or have some kind of advantage relative to non-normative formulations in steering optimal belief-formation, seems implausible.⁶⁶

Another objection to my line of reasoning might run as follows. If we allow that the goal of a hypothetical (or instrumental) imperative is somehow justified – as AP at least sees it as being – then surely the best means to attaining this goal can stand alone as a norm for belief-formation. There may be people who do not know that Y is a justified goal of cognitive enquiry, or perhaps do know but do not care. Nevertheless, if Y is inherently valuable, surely these people should still do X if X, and only X, leads reliably to Y; moreover, they may need to be specifically told to do so. Relatedly, there might seem to be people within a scientific community who need to be told explicitly what to do – technicians who follow procedures and instructions, or even 'higher-level' scientists mostly concerned about money and prestige – without concern for what scientific ends following these leads to.⁶⁷

In response to this, I think I can only concede that science, indeed, knowledge-gaining activity generally, perhaps does need some kind of rules or norms if it is going to function optimally. If cognitive inquiry requires, as a matter of practical fact, the cooperation of people who are just not interested in truth or rational belief (or whatever general characterisation we choose to give of specifically epistemic activity) then, just as there would seem to be a need for societal rules, so there might also be a need for rules to govern the belief-forming practices of this kind of, let us call them, *philistine* (no negative connotations intended!). However there are several points we should bear in mind in admitting this. First, it is important that these people's lack of concern for the goals of enquiry is not the upshot of limitations on their capacities – remember that the relevant goals can always be specified in quite general and succinct terms. Secondly, the point would seem to be quite a general one in relation to NE, not having anything particularly to do with an instrumentalistic conception of norms. For example, if certain people are needed to reason well when they are not interested in using their ability to reason well (of the kind most versions of AP assume we possess), we might need to set up explicit norms dictating the way they should reason – norms which, for those willing to use their ability competently, would not be required. Thirdly, it is important to bear in mind that the philistines would be performing a job *in the service of enquiry* into the nature of reality, not actually *conducting* that enquiry. In the light of all this, I think the core of my overarching

thesis remains intact, for as far as those who are concerned with cognitive enquiry go, it would remain the case that norms would not be required for optimal epistemic progress. Indeed, it seems ultimately doubtful whether the norms the philistines would be required to follow, though perhaps *coinciding* with norms that non-philistines might countenance for cognitive enquiry, could really be said to be playing an *epistemic* role for the former. They would, typically, be upheld through incentives for abiding by them (i.e. salaries), plus the threat of withdrawal of such in the case of breach. In that sense, the norms they followed could, I think, reasonably be said not to be epistemic but merely practical ones.⁶⁸

If what I have argued above is correct, both Laudan's specific theory of NN and the more general idea of norms as hypothetical imperatives fail to yield varieties of NE whose norms are required for optimal belief-formation. To the extent we might need people to engage in epistemic practices whose goals they are not interested in, we might need to set up explicit rules or norms for them to follow. However, it seems reasonable to deny that, in such a role, these norms really should count as epistemic ones.

At this juncture, a supporter of instrumentalistic NE might say that she can, at least for the sake of argument, go along with my conclusion here, but that it really has no impact on what practitioners of instrumentalistic NE aim to achieve. Whether or not we need normative formulations of the empirical connections between means and cognitive ends, it remains a fact that this knowledge is required for optimal epistemic progress. Hence we need instrumentalistic NE, even if we do not need norms, because we need to be told what the best means are for achieving our cognitive goals.⁶⁹

However, this description of what cognitive agents need to know is tendentious. If what I have argued is correct, talk of what is 'best' only makes an essential entrance in relation to *individuals'* consideration of how to achieve their cognitive aims in particular circumstances. Practitioners of NE do not offer information about what is best or optimal or anything of that kind, or at least doing so is unnecessary (and, possibly, as I have indicated, misleading); their contribution can be restricted to empirical information which individuals may then make use of in their own decision-making. These practitioners might still retort that they are needed for such optimal decision-making – for optimal cognitive progress, on the

present conception of what this involves. However, that empirical information is needed for such progress is not, of course, something I have ever denied – indeed, it is a cornerstone of naturalism. In saying this, I am in a way admitting that what some people who call themselves naturalistic epistemologists do might be necessary for optimal epistemic progress; and thus, if NE is by definition what these people do, then NE is not superfluous. However, it strikes me as highly significant to point out that this discipline and those who practice it are *nothing more* than scientists investigating certain empirical connections of the natural world. That this will be of significance in the progress of knowledge is beyond doubt; but then the same can be said of all other scientific activities. I take it that naturalistically inclined epistemologists have seen their role as something beyond that of ordinary scientists, and that at least those who read them have taken it that they see their role as something beyond this. It strikes me, therefore, as highly significant to point out that it does not go beyond this, or at least does not have to.

4.6 Why does it appear as if scientists use norms if they do not need to?

I have been arguing that norms are not necessary for optimal rational belief-formation if one cleaves to naturalism in the sense outlined in Chapter 1. As yet, I have not considered a psychological version of NE (P), which would have it that norms can literally be identified with principles in some kind of natural reasoning competence, a view I will also be arguing against (cf. Chapter 5). However, even if P were tenable, the norms in question would not be explicit norms of the kind one traditionally associates with epistemological theories (even though I am allowing at the outset that they could nevertheless class as norms). If what I argued so far is correct, and my purview of varieties of AP and AF is exhaustive, we can already conclude that no *explicit* norms are both defensible as genuinely rational and necessary for optimally rational epistemic practice.

This might strike some as a highly doubtful conclusion. For is it not simply obvious that, if not ordinary people, then at least scientists actually use methodological norms to guide their activities and are dependent on so doing? And does not this very datum suggest that something, somewhere in the course of my arguments or

taxological activities in this or the preceding chapters must have gone wrong? Well, if so, I would like to be told what! But in any case, I believe it is relatively simple to explain away the appearance of widespread dependency on norms in science, in a way that will reinforce the plausibility of what I have already argued.

In my view, the norms that science is supposed to rely on will, in general, fall into one of five (not necessarily mutually exclusive) categories whose existence is compatible with the arguments I have presented:

- (i) ideologies;
- (ii) invalid and/or inadequate norms;
- (iii) inveterately vague norms;
- (iv) common sense dictums;
- (v) factual statements about the world masquerading as norms.

In what follows, I shall consider a variety of intuitive examples of norms in science and investigate to what extent they can be placed somewhere within this taxonomy (in so doing it will also become clearer what the different categories amount to). I am not seeking to prove that absolutely all such norms can be so placed. My dialectical strategy is rather that of seeking to dislodge qualms about my *argued* conclusion that norms are not necessary in science or for rational belief-formation, given naturalism. In the course of the discussion, I will also have recourse to comment on the role of epistemic norms in everyday or common sense reasoning.

As a first observation, I think we should all at least be open to the idea that the mere *appeal* to a methodological standard by scientists does not show that the standard is actually used to *guide* their research. There is now empirical evidence from sociological studies that putative epistemic norms are often appealed to by scientists in a way that is rationally idle. I am referring here to cases where scientists put forward hypotheses or reach conclusions and appeal to norms in so doing, but where the norms seem only to give some kind of *ideological* backing to what they have done – that is, the norms fall into category (i) above, *ideologies*. The sociological case-studies in question relate how controversies in science are brought to closure. In this process, appeal is often made to traditional norms of scientific practice – such as utilising so-called *decisive experiments* or

avoiding *ad hoc*-defences of falsified theories – but in fact it seems these norms have little to do with the way in which closure is actually effected.⁷⁰ Now, we would not wish to suggest – and nor do most of the sociologists – that the resolution of scientific controversies undermines the idea of rationality or rational progress in science. But precisely because of this, the norms invoked at times of conflicts would appear inessential to this rationality.

One might perhaps wonder whether these case studies have any bearing on science outwith conflict situations. In response to this, it should be noted that many of the most well-known norms associated specifically with science concern precisely the resolution of conflicts and controversies. Moreover, to the extent that 'normal' science involves the continual resolution of 'small' controversies, as is not totally implausible to hold, the fact that norms function only ideologically in large controversies suggests that something similar might also be the case in smaller ones.

Another problem one might have with viewing norms as ideologies is that it might seem unclear why norms should be thus, assuming they were. In one sense, this is an empirical, socio-psychological question that I neither can nor need answer fully here. The fact that humans seek ratification for their actions in overarching norms seems a relatively deep feature of human nature, but it is not thereby a rationally justified one. However, a more determined objector to my line might not be over-impressed by the sociological data in the absence of some further *explanation* as to why norms in general are unlikely to be capable of functioning as anything other than ideologies. She might concede that, on occasions, they could be misused, but insist that the idea that they generally function in this way requires deeper explanation.

I think such explanation can be offered. To begin with, it seems that one central reason why many norms are used ideologically is that they are *invalid* or *inadequate* (i.e. fall into category (ii), *invalid and/or inadequate norms*). Surely we would want to reject the naive falsificationist norm of rejecting theories that entail false observation statements, or the naive inductivist norm of accepting those that conform with the traditional canons of induction. Yet such crude norms are, at least at times, appealed to by scientists in their work. To the extent their work does not suffer, it seems these norms can be playing little more than an ideological role.

Of course, some scientists have finer methodological sensibilities. However, the price paid for greater plausibility in methodological rules is often loss of precision, a point which also casts doubt on the efficacy of norms in scientific reasoning (here we turn to those that fall into category (iii), *invariably vague norms*). This is, as noted in Chapter 3, a prominent topic in the writings of Kuhn. Kuhn suggests – in work addressed to the charge against him of relativism – that scientists rely on shared *epistemic virtues* in choosing and evaluating their theories, such as simplicity, elegance, fecundity, internal consistency, agreement with available data and so on. At the same time, he confesses that many⁷¹ of these criteria are so vague that they will in practice not function as paradigm-independent standards for theory-choice.⁷² Now, one might well see this as an attempt on Kuhn's part to avoid being labelled a relativist without in fact deviating from the spirit of his original account – as taking back with one hand what he seems to be offering with the other.⁷³ However, the vacillation can plausibly be seen as diagnostic of a deeper point that supports my line,⁷⁴ namely that many shared norms are of such inherent or inveterate vagueness that it is in fact implausible to hold that they will have any general and independent efficacy in steering theory-choice (*simplicity* seems to be an excellent example here).

Note this need not mean that these more or less vague kinds of consideration *cannot* rationally steer theory-choice; indeed, it would, I think, be implausible to claim that they never in fact do so. It is sufficient for my purposes that on any particular occasion it is far from clear that they do *in fact* play such a role, something that can therefore render more palatable my overarching, argued claim that it is never *necessary* explicitly to invoke such vague norms. (Insofar as I would also hold that any substantive content attaching to vague norms is part of our common-sensical, everyday reasoning capacity, there is also another route to understanding their superfluity, as we shall see below.)

But surely not all the epistemic virtues that Kuhn mentions, or that scientists might use, are invalid or inherently vague. Laudan points out that one of Kuhn's own suggestions, that of internal consistency, seems to be a crystal clear notion.⁷⁵ Perhaps the meta-axiological criteria Laudan himself is concerned with, and which we discussed above in Section 4.3 – concerning, centrally, constraints of anti-utopianism – are also sufficiently clear to be generally and

meaningfully applied. However, the apparent problems these considerations raise for my view can also be assuaged. Consistency is a general logical criterion that presumably applies across the whole ambit of epistemic activity, whether lay or scientific. As noted in Section 1.3, whether we, or scientists, can seek to achieve it quite generally seems unlikely. But to whatever extent we can, the goal of consistency presumably remains an aspect of *general* epistemic evaluation, not something peculiar to science. In other words, it is a *common sense maxim* (category (iv), above) – along with, arguably, many others that are often cited as general norms of science (such as Chalmers' *Take argument and the available evidence seriously and do not aim for a kind of knowledge or a level of confirmation that is beyond the reach of available methods*, discussed in Section 3.2).

Some readers might be unclear as to how it can help my case to show an apparently important norm of science is in fact just a general norm of all reasoning. The point is that it is far less clear than it is in science that norms are explicitly appealed to in ordinary epistemic activity. Ordinary people form basic beliefs, draw inferences, even theorise perhaps, but it seems it is only when a philosopher (or perhaps a psychologist) comes on to the scene that we begin to have an inkling as to what principles steer this activity. Indeed, much of epistemology and philosophy of science is often billed as making explicit norms that are only implicit in everyday practice. I have argued, given naturalism and the avenues this opens for justifying norms, that we do not need to make these norms explicit. If I am right, the 'folk' can go on being just as wholly reliant on their epistemic ability – when deployed competently – as they always have been. As for scientists, one could perhaps recommend that they admit reliance on this basic underlying ability rather than appeal to the more dubious normative standards they occasionally espouse (even though they will, of course, apply this together with other, technical and observational abilities to far more difficult and *recherché* problems than most ordinary people deal with).

I believe similar observations can also apply to Laudan's anti-utopian meta-axiological constraints. Chalmers' overarching norm, cited above, incorporates elements of Laudan's ideas. In fact, I think it is far from clear that anti-utopianism is both contentful enough to have substantive implications for belief-formation, and anything like remotely plausible.⁷⁶ But in any case, there seems no reason to

suppose such constraints are peculiarly scientific, that is, that they apply within science but not common sense epistemic practice; they thus fail to constitute a convincing counter example to my arguments.

At this point, I can imagine that many readers, particularly those with scientific backgrounds, may be feeling a growing sense of impatience insofar as I seem to be concentrating on traditional philosophical conceptions of norms – and ignoring the vastly more numerous, more quotidian, but still very real normative standards that scientists routinely apply in their work. It is category (v) – *factual statements masquerading as norms* – that I intend to use to explain the apparent ubiquity of this kind of norm. Now, I should state right away that I fully accept that scientists routinely issue general, normative statements – i.e. norms – such as *In situations of type so and so, one should do such and such to gain optimal information* or *All interpretations of data should ideally conform to such and such theoretical assumptions*. Certain scientific ideas and principles will typically be regarded as ‘standards’ that any other must cohere with, and thus constitute norms in an oblique sense. In saying that such standards are all factual statements masquerading as norms, I do not mean that scientists must formulate themselves without the aid of normative vocabulary. Nevertheless, I do believe that one can and should become clear over the fact that such vocabulary is unnecessary, and that all norms can, in principle, be dispensed with in favour of purely factual, descriptive statements (which of course is not to say that normative statements enjoining *particular* epistemic moves can be so replaced).

It is not feasible for me to prove this quite generally for all possible kinds of normative statements that do or might occur in science. Instead I want to consider two kinds of case which I believe can subsume many of the multifarious varieties of norms one finds. One of these we have already discussed in relation to the conception of scientific norms as hypothetical imperatives. We will return to this presently. But first, I want to consider cases of norms which prescribe accordance with a particular theoretical claim, or set thereof.

A simplified but, I hope, illustrative example of this from physics could be: *Construe experimental results such that they accord, as far as possible, with the basic assumptions of quantum mechanics and relativity theory*. Now, I take it that something like this is in some sense a norm of physics (or perhaps a cluster of norms); let us call it QR.⁷⁷ The

question is: in what sense? If one is aiming to work on certain fundamental theoretical problems in physics, familiarity with the relevant aspects of quantum mechanics and/or relativity theory will be a prerequisite. Moreover, one will be aware, not, of course, of all the evidence that supports such theories, but of the fact that they are very, very deeply entrenched in our scientific web of belief. In other words, they will have the status of non-negotiable axioms, or perhaps, in some sufficiently loose sense, of ‘framework assumptions’ that can help organise other knowledge.⁷⁸ Now let us suppose we carry out some research into some hitherto unexplored region of the galaxy. Our aim, very schematically put, is to reach some conclusion or set of conclusions on the basis of the data gathered and the available background theory. Now ask: do we need QR *per se* to make the best possible inference on the basis of our evidence? Surely not. For that inference must in any case take heed of quantum mechanics and/or relativity theory (i.e. those aspects which are relevant to the present investigation). It is an inherent part of the standard cognitive practice of scientists that they draw inferences on the basis of all the relevant information (albeit, as individuals, only imperfectly). Such a practice is part of what naturalism, as characterised in Chapter 1, involves. Norms that tell scientists to do this in relation to specific relevant propositions are therefore unnecessary (which, again, is not to say that normative statements enjoining particular inferential moves are).

As a further example of a theoretical proposition that arguably plays a normative role, we can mention the general idea that the world is divided up into natural kinds. Some might think the principle has an *essentially* normative character.⁷⁹ However, in the spirit of naturalism, it seems just as reasonable to regard it as a very high-level and deeply entrenched fact, and thus one we perforce take account of in our scientific reasoning.⁸⁰

Let us now turn to the other category of norms that I would claim can be seen as masquerading as factual statements. Many sciences contain methodological rules for conducting research, and these may be quite peculiar to the individual science in question. In social science, for example, one has rules for gaining good survey results; in chemistry, rules for gaining accurate results on the basis of reacting substances together in a laboratory; in physics, rules for using Geiger counters, telescopes; and so on and so forth. What is noteworthy about all such examples is that, in accord with Laudan’s view of

methodological rules, they seem expressible – indeed, seem fundamentally to have the form of – hypothetical imperatives: one has a specifiable or specified goal, and one has a specified means of obtaining that goal (perhaps in some specified circumstances). I have argued, in Section 4.5, that such hypothetical imperatives could be replaced by factual statements stating – albeit perhaps only generally and imprecisely – a connection between the goal and the methods in question. Clearly such factual statements must be taken account of in doing research of the relevant kinds, for reasons that essentially parallel those given in relation to norms like QR. But given they *will* be taken account of, it also seems that one does not need an explicit statement of a rule prescribing actions for the attainment of certain aims. Now I also allow that if science has some specifiable, universal goals that are inherently valuable, then there might be people ignorant of or indifferent to these aims that need explicit normative steering to perform their required services to science. However, I also argued that these people's dependence on rules does not show that science or scientists *as such* are so dependent, or that the rules the former follow really deserve to be classed as epistemic ones.

Some readers might be wondering how all of this can amount to anything more than a verbal dispute between me and supporters of NE.⁸¹ After all, I have conceded (here and in Section 4.1) that norms *in some sense or other* are operative in science. However, as I understand the supporter of NE, her aim is to motivate, and motivate the need for, norms in a *substantive* sense – norms that, at least to some minimal extent, parallel those offered by the apriori, traditional project. If norms in science amount to no more than factual statements, then it strikes me that there is at least very little need for what anyone might want to call 'naturalised epistemology' beyond Quine's general observations about science as a variegated and evolving web of belief (plus the rest of what we said about naturalism in Chapter 1). It is clear that many philosophers have seen both the possibility and the need to go beyond this background metaphor in their theories; I am suggesting that this is an illusion. In this last section, I have therefore also seen it necessary to underline the fact that the norms scientists genuinely rely on – i.e. that are not merely ideological, inveterately vague, invalid or part of common sense – are, in any case, not such as to demand construal as norms. (The points in this paragraph are related to those of the final discussion in Section 4.5.)

Chapter 4 is nearing an end. In concluding, it is perhaps worthwhile sketching my overall view at a somewhat coarser level of grain.

Science is often seen as guided, at the most fundamental level, by certain basic normative precepts which, ideally, procure a truth-tracking character for research (at least according to most traditional conceptions of it). My view is that these precepts are often only deployed for ideological purposes, and that they in any case tend to lack substance because, to the extent they are reasonable, their formulation is inherently vague. This does not necessarily apply to all such precepts, but it is also arguable that those that are usefully precise, as well as plausible, apply equally to common sense reasoning. But the latter seems to get along fine without explicit norms – as we would expect given these have their basis in some kind of ability or competence we all possess. At a less foundational level, it is also often supposed that science employs norms of an instrumental or hypothetical character; some would say that ordinary people should also take heed of similar such norms in relation to their everyday belief-forming practices. But, whilst these in some sense are indispensable, there seems reason to hold that they are facts masquerading as norms, or at least that they can be unproblematically replaced by factual statements. With these considerations in place, my more detailed arguments for the superfluity of norms in this and the previous chapter will, I hope, appear a good deal less fabulous than they might otherwise have done.

I have now argued against the two varieties of NE – AP and AF – that would seek to motivate the idea of explicit naturalistic norms. But what of psychological NE, or P as we are terming it – which precisely identifies norms with *implicitly* known principles in a natural reasoning competence? To show the failings in that we need to consider a host of new issues.

5

Psychologistic Naturalised Epistemology

Introduction

Our verdict on NE thus far is that neither a version of AP nor of AF is capable of delivering genuine epistemic norms that are also required for optimal belief-formation. Even if such norms are valid, they are superfluous.

This leaves us with psychologistic NE, or P as we dubbed it in Chapter 2. In fact, this theory, motivated through our discussion of competence and the rationality problem in cognitive science, constitutes only one of the two main sub-varieties of psychologistic NE that I shall be considering. The version discussed in Chapter 2, and which I shall continue to call P, might be termed *descriptivist* psychologistic NE, insofar as what it wants to regard as norms correspond, in the first instance, to purely factual or descriptive states of affairs – on the basis of the nature or situation of which one then hopes to reconstruct, or justify the idea of, something normative. However, there are also those (albeit perhaps a minority) who hold that a naturalist outlook on the world does not preclude finding normativity there in an undiluted, unreconstructed form. In particular, our mental capacities are naturally such as to be imbued with normativity. This kind of position is perhaps not one which is traditionally associated with labels like ‘psychologism’ and ‘naturalised epistemology’; it is motivated independently of Quine’s arguments in ‘Epistemology naturalised’ and Goodman’s MRF, and may not even be fully reconcilable with my notion of naturalism. Nevertheless, it constitutes, or can reasonably be seen as constituting, an

attempt to understand epistemic norms in both a psychologistic *and* (in some sense) naturalistic way, insofar as it draws its motivation from the special kinds of explanation that are meant to be involved in understanding the actions and intentional states of human beings.¹ I will refer to it in the following as NP (for *normativist* psychologistic NE).

This chapter will offer a critique of both P and NP. But first some preliminary remarks, by way of reminder and sign-posting, are in order. For a supporter of P, we possess a natural reasoning competence and it is the principles that this contains – or at least the principles which its structure *indexes*² – that constitute the correct norms of reasoning and belief-formation.³ On this line, the correct rules of reasoning are part of our fundamental psychological nature, and somehow used by us in our reasoning (even though they do not always manifest themselves in behaviour, i.e. our reasoning *performance* can diverge from our reasoning *competence*). The objections to AP and AF that one does not need to formulate the rules explicitly for us to reason optimally has therefore no application to P, since epistemic norms are precisely seen as *implicit* or *tacit rules* that steer our behaviour (or as rules indexed by a psychological structure that steers our behaviour – a qualification I will just assume from now on). Now, the very idea of such tacit rules is to an extent controversial, but the supporter of P sees it as legitimated by the practice of cognitive science (making it naturalistically kosher), most prominently the study of our grammatical competence as pursued by Noam Chomsky and his followers (cf. Chapter 2). In this presentation, the legitimacy of talk of implicit rules will not be brought into question, nor will I analyse what it really amounts to.⁴ Nor will I take up the question of to what extent empirical evidence suggests we do or do not in fact possess something like a natural reasoning competence on a par with grammatical competence, but again simply assume that we do possess such a competence.⁵

According to P, then, a study of our natural reasoning competence is exhaustive of epistemology. But it is also, in contrast to the role of psychological study on Quine’s picture (as we interpreted it in Chapter 2), meant to yield a *substantive, normative epistemology*, that is, a correct set of epistemic norms: a theory of epistemic norms ultimately *just* is a theory of the principles of our reasoning competence. This theory will probably be arrived partly with the aid of something

like MRE, as outlined in Chapter 2.⁶ At the same time, MRE will fall into place as a part of a broader empirical enquiry into our reasoning competence, this in turn being one of a number of other enquiries into other cognitive competences that are familiar from extant work in cognitive science. P departs from TE in denying the apriority of epistemology, but not, in one central sense, its foundationalism, insofar as the epistemic norms it seeks to adduce form a fixed basis for our beliefs about the world. Indeed, one might even say that epistemic norms for P are in some weak sense apriori insofar as they will be *temporally* prior to the science we carry out on the basis of them. The difference between P and TE lies in that we cannot know what the norms are by means of some occult faculty of intuition, but must uncover them through scientific investigation. In this way, P can fairly be viewed, I think, as the naturalistic heir to the rationalistic project which in Chapter 1 we identified as a paradigm of TE.

It is not, perhaps, obvious which particular contemporary thinkers might be seen as cleaving to P. First and foremost, I see P as one, highly significant way of giving expression to the philosophical impact that many see in the coupling of Chomsky's neo-rationalism with the naturalism of Quine. The following is a more explicit statement of a view which at least strongly resembles P by one such thinker, Louise Antony:

A theory of the well-functioning rational agent is an empirical theory of what *constitutes* rationality, just as an empirical theory of the syntactic knowledge of a competent speaker constitutes a theory of grammaticality. The bet is that such theories will legitimate a good number of our pre-theoretic intuitions about rational strategy [...] A naturalised epistemology, in telling us what we *do*, does not thereby abjure telling us what we ought to do.⁷

So much by way of presentation and motivation of P. As already noted, such a construal of epistemic norms does not face the central problem that AP and AF do, namely that of showing that the norms are needed. Nevertheless, P faces other, at least equally severe challenges. One issue concerns the ongoing debate in cognitive science as to whether we are naturally rational, that is, whether our natural reasoning principles are correct for belief-formation. If there are good

reasons for denying this, or if there is some reason for suspecting it and no good reasons for thinking otherwise, P will obviously not stand on strong ground, whatever else might speak in its favour. Thus the so-called *rationality debate* and its myriad of conceptual, methodological and empirical issues is in principle a central concern for us. Fortunately, we are now in the position of at least having had a lot of the ground cleared by Edward Stein in his recent book *Without Good Reason*. Stein, however, reaches no firm conclusions as to whether or to what extent we are naturally rational. In Section 5.2, I will briefly review the terrain, mainly just to indicate how thorny the issues are. Nor shall I reach any firm conclusion, but will nevertheless carry forward the (tentative) assumption that we are naturally rational. For even if we are, the idea that the principles of a reasoning competence can literally constitute norms is beset with difficulties.

These revolve around the problem showing in what way the primarily *descriptive* principles of a natural competence can be regarded as having relevance to the *normative* matters of the kind epistemology is concerned with. This issue is related to our discussion of the problem of normativity in Section 1.5. There I argued that the problem of how an 'ought'-statement can be logically derived from an 'is'-statement should not be viewed as a general obstacle to NE. However, in relation to P the issue is not – merely – that of deriving an 'ought' from an 'is', but moreover of trying to understand how norms literally can be principles that are fundamentally descriptive. Though I will argue (Section 5.3) that the issues surrounding P and normativity are more intricate and less clear-cut than I think many have assumed them to be, I will uphold the idea that seeing psychological reasoning principles as norms is – at least – non-mandatory and unmotivated, given its naturalistic starting points. Hence P is undermined.

In Section 5.4 I turn to NP. The central claims behind NP can be viewed as related to the rationality debate and the problems of normativity discussed in Sections 5.2 and 5.3. In relation to the former, they connect with an idea that many philosophers have sought to use to rebuff empirical evidence supposedly indicative of our irrationality, namely, that delineating human psychology is inextricably linked with placing conceptual, normative constraints on what can be believed and how. Thus empirical evidence cannot show we are irrational. Supporters of NP, as I understand them, relate to our debate – and

distance themselves from TE – in virtue of the fact that they see this normativity-imbued process of description as responsive to the way things stand in reality, that is, *naturally are*. Thus genuine epistemic norms, and not merely descriptive principles, are a part of the natural (human) world. The philosopher whom I think most clearly can be associated with such a view is John McDowell, especially through his idea of *second nature* as outlined in his recent book *Mind and World*.

The kind of view McDowell defends is often dismissed at the outset by naturalists as simply unintelligible or incoherent. I believe it is motivated by *prima facie* genuine concerns, and though the need to give a substantive epistemological theory is not an explicit one for McDowell, it seems this also could be seen as motivating a position related to that McDowell defends on the basis of reflection on (as it is in his case) perceptual belief and intentional explanation. It is such a position that I term NP. Nevertheless, I do not think NP is tenable. The reasons need spelling out carefully, but they revolve around the fact that norms, as NP conceives of them, cannot figure in explanations of belief-formation (or intentional states/events generally), and hence have no naturalistic grounding. NP probably is incoherent after all, and at least lacks motivation.

Before getting underway with all this, however, I want, in the first section of this chapter, to consider variants of NE which would see facts about our psychology as constitutively *relevant* to epistemology but not *exhaustive* of it. This overall line has seemed attractive to many thinkers, and can also be seen in connection with the *weak* or *moderate* naturalism discussed in Section 1.5. Since this is first and foremost a book about NE in a strong sense, I am not going to attempt any exhaustive survey of possible theories. Nevertheless, I think there are fundamental problems with the weaker naturalist approach understood as an attempted rapprochement between traditional normative concerns and empirical matters that can seem relevant to epistemology. I outline the general form of these problems by focusing on a position, defended by Edward Stein, that is, I think, representative of a broad range of similar views. Since we might consider such a view a 'mixed' form of epistemology, I will also consider several other mixed forms in the same section, and show why in their case too the only substantive normative impact they could have would be that provided by any *a priori* elements they countenance (at least, to the extent they abjure psychological NE).

5.1 Ballpark psychologism and 'mixed' varieties of NE

In recent epistemological theorising, many writers have argued that some kind or degree of psychologism must be accepted in epistemology; that we need at least to *take account* of our actual reasoning capacities in giving a theory of epistemic norms. We might, to loan an expression from Hilary Kornblith, call this 'ballpark psychologism',⁸ in order to indicate the very general nature of the commitment at the outset. The question then becomes in precisely *what sense* psychology is relevant to epistemology (not just 'ballpark-wise'). For some, the study of our psychology will play an important but in the end essentially heuristic role in (normative) epistemology, that is, it will expedite epistemology, but the latter will be essentially *a priori*, and doable in principle without recourse to empirical data.⁹ For many others, facts about our psychology will be actually constitutive of normative epistemology, though not exhaustive.¹⁰ For others, it will also be exhaustive, as P and NP hold.¹¹

One of the main reasons offered for one of the stronger lines is that, if one ignores psychological facts, then one may end up prescribing inferences or practices that human beings simply cannot hope to emulate (cf. the discussion of the methodological problem for TE in Section 1.3, which the present paragraph recapitulates elements of). For example, having consistent beliefs is presumably a desirable epistemic goal, and one might therefore think we should, rationally, make moves to ensure that we have consistent beliefs, that is *check*, each time we acquire a new belief, that it is consistent with those we already hold. But it turns out that actually performing a check for consistency in anything beyond a trivially small number of sentences leads to combinatorial explosion, such that it is impossible for a finite being like us to perform. Yet it (arguably) makes no sense to say that it is rational to do something that we *cannot* do – any more than it makes sense to say that one should morally do what one cannot do. Hence, though aiming to ensure consistency in our beliefs might seem like a more than reasonable epistemic norm, it turns out that, given our radically finite nature, it is something we should in fact renounce. The general moral is that we need to take account of actual human cognitive powers in giving a theory of epistemic rationality.¹²

Nevertheless, according to many who make this kind of observation, it is not the case that a theory of epistemic norms will reduce to a theory about our psychological capacities and limitations – our reasoning competence, understood in a broad sense – as P holds. Thus, according to Edward Stein's *naturalised picture of rationality*, a theory of norms will derive from a process of *wide reflective equilibrium* (see Chapter 2) in which both facts about our psychological abilities and philosophical theories, along with intuitions about rules and (considered)¹³ judgements of concrete cases will constrain the overall balancing process.¹⁴ The idea is that the norms produced in this manner will be more acceptable than if they were derived through only some of these sources. And the theory is naturalistic because what we *can* in fact do, in virtue of our natural capacities, constrains what we can be said to be *behaved* to do.

How should we evaluate this suggestion (which I think is representative of a wide range of 'weakly' naturalistic positions)? Stein does not seem to share my basic conception of MRE as a naturalistic, theory-building process, but rather sees intuitions about rules as constituting a substantive source of information to be taken account of in an overall balancing process. In view of my arguments from Section 2.2, I think he should revise his view on this matter. Even once this is done, however, it seems that his talk of philosophical theories will essentially – and can only, as far as I can see – involve *apriori* conceptions of correct reasoning. These will thus have an independent impact on our final theory of norms, in the same way as facts about our psychological make-up are meant to.

The objection I have to such a view is essentially similar to that offered against a version of AP in Section 3.3 that was seen as building on a process of wide MRE. As with that theory, the idea behind Stein's naturalised picture of rationality seems to be that one should derive a theory of norms by going through a process of MRE that takes into account various additional sources of information. In Stein's case, these additional sources are of two kinds: knowledge of cognitive limitations in humans (or whatever the relevant psychological evidence is), plus *apriori* intuitions about correct rules. The problems with the two theories, however, are essentially very similar, revolving around the question of how the relevant bits of information are seen as impacting on our decisions about which norms to accept. With respect to the psychological information, the point is

exactly parallel to that made with respect to the version of AP. On the one hand, we could choose to let the norms derived without their aid stand as such, informing people about the relevant psychological facts and allow them to make up their own minds, on a situation-by-situation basis, as to exactly how these should be taken account of. On the other hand, we could say that one needs to derive norms on the basis of the psychological information, which would then be incorporated into the process of MRE; that is, we would seek to derive norms through a reflective process that operated on our intuitive judgements augmented by knowledge of our psychological make-up. Given this, we can now argue (in the manner of Chapter 3) that the norms thus derived would not be necessary to achieve optimal belief-formation, for one might just as well rely on the ability-plus-knowledge basis it seeks to encodify.

What about the *apriori* information? Again, it seems that one could either choose to incorporate it into the MRE process, or just let it stand as a kind of ideal to be aspired to as best as one is able to in particular situations. In fact, given it is *apriori*, it is hard to see how it could be treated as anything other than an ideal; but if it were not, the only alternative would again be to incorporate it into the process of MRE as part of the basis upon which that method operates, and the argument above that one would not need the norms thus yielded would go through as before. In sum, it seems that Stein's naturalised picture of rationality might end up with *apriori* norms that we need to take account of (in accord with his weak form of naturalism), but that any other kind of norm would not be required.

We might call Stein's naturalistic theory a 'mixed' view, since it seeks to incorporate elements from different sources in delivering a theory of epistemic norms – in Stein's case both *apriori* and naturalistic elements. Other mixed views are also possible. According to one that, like Stein's, could be seen as assuming a weak form of naturalism, the *goals* or *aims* of epistemic practice would be specified by *apriori* reflection, but the *means* to attaining these goals would be informed by empirical enquiry into what most reliably leads to these goals (possibly in combination with knowledge about our psychology – see below for more on this). My impression is that this kind of view is probably quite popular amongst many who would call themselves naturalised epistemologists.¹⁵

Another mixed view would be one which combined Mill's (or EE's) conception of inductive norms (cf. Chapter 4) with a more 'foundationalist' conception of deductive logic (say), which might be built on something like MRE or apriori intuition. A further possible combination is a view which sees norms as instrumental ones for obtaining specified goals (justified apriori or otherwise), but in which the means specified in the norms also take account of our psychological capacities and limitations. A perhaps related position that we find in the literature is that of Harold Brown, who gives an account of our natural rational capacity that he then cements together with some (supposedly) plausible-sounding considerations of a seemingly non-foundational character to arrive at a set of general norms for science, at least as we have it at the moment; he does not specify systematically how these are arrived at, but it seems safe to assume the use of something like (wide) MRE.¹⁶

Mixed views either simply 'bolt' together several differently conceived sets of epistemic norms into one overall theory, or else seek to bring several different sources of information to bear in deriving an overall theory of norms. (In fact, the second fusion of ideas in the preceding paragraph could plausibly be classed as a version of AP or AF if the aims were conceived appropriately; the important point, however, is that I have not yet considered it explicitly!) Given my arguments thus far, however, it should be relatively straightforward to see that – with the possible exception of apriori norms derived from logic/mathematics – none of these strategies will yield norms that are necessary for optimal belief-formation. I have already criticised views that see epistemic norms as instrumentalities (Section 4.5); that the goals might be justified apriori would not seem to affect those arguments in any way. Nor would the argument be materially affected by taking into account psychological abilities in a specification of such norms. A hypothetical rule such as *If you want A, do B, would still be replaceable by descriptive information* – not perhaps by *B leads most reliably to A*, but at least by *B leads adequately reliably to A and is the most feasible means available to humans for obtaining A*.¹⁷ The combination of Mill's view with an alternative view of logic merely inherits the problems of both accounts. As for Brown's view, it is too underspecified to evaluate properly here. What we can say is that, insofar as the norms are based on empirical facts, we might just use these instead of the norms (cf. Sections 4.1 and 4.6); and insofar as a

wide form of MRE is employed, it seems open either to the objection just discussed, or to the objections raised against Stein's account and the related version of AP (or perhaps both sets of objections). Finally, it is important to stress that even my concession to logic and mathematics should not give any succour to a moderately naturalistic epistemologist, for there will be nothing of *naturalistic* origin in her theory that would have to show up in the final theory of norms.

We are left, then – as far as the prospects of NE are concerned – with the prospects for P (and NP). Before turning to these, however, I would like to point out that I think one can respect the spirit of ballpark psychologism at the same time as denying that it yields any kind of NE (or a weakly naturalistic epistemology). To begin with, one could hold that a purely descriptive study of our cognitive systems, though not itself epistemology, is a worthy and natural heir to the traditional normative study (cf. Section 2.1). More substantially, one could also hold that facts about such systems can have normative implications for our reasoning practices and for science generally. In criticising epistemological theories that seek to take account of facts about our psychological make-up – I have not, it will be noted, argued that these are irrelevant to normative decisions about what should be done in particular circumstances, for example, whether one should seek to find an inconsistency in a certain set of beliefs. On the contrary, I have suggested that they would be relevant in this regard, but that if one seeks to summarise their impact in an overall theory of epistemic norms, then it will turn out that one will not need these norms.

Facts about our psychology could also, in principle, have normative implications for our scientific knowledge. In writings subsequent to his 'Epistemology naturalised', Quine has suggested that findings from psychology about how we acquire knowledge of the world could lead to revisions about what we take ourselves to know about the physical world. To some, this has indicated a change of heart on Quine's part in relation to his original idea that we should replace epistemology with descriptive psychology. But it seems to me that Quine is simply making a distinct and quite compatible point (cp. the discussion in Section 2.1). Insofar as facts about our psychology will be part of the same overall web of belief that our knowledge of the physical world fits into, it seems clear that what we decide to believe

6.5
can we
decide
no believe

about the latter – that is, what we *should* believe – may in principle be dependent on facts of the former variety. Thus psychology, along with any other branch of knowledge, is capable of constraining knowledge claims normatively, simply in virtue of the background, anti-foundationalist picture that naturalism embraces (see Chapter 1) – even though it will play no role in giving a theory of norms.¹⁸

5.2 Are we naturally rational?

As noted in the introduction to this chapter, the question as to whether we are naturally rational – whether our reasoning competence (something we are assuming exists) embodies normatively correct principles – would seem, on the face of it, to be one that any supporter of P needs to take very seriously. In this section, I aim to do little more than gesture at the terrain, drawing partly on issues touched on in Chapter 2, and otherwise on Edward Stein's admirable summary of the debate. I will, however, offer some brief critical commentary of Stein's arguments with the aim of tipping the balance somewhat more in favour of the so-called *rationality thesis* – i.e. the idea that we are naturally rational – than he does.

As discussed in Chapter 2, human beings perform in ways that diverge from what are often taken to be normatively correct standards of reasoning. As also noted, these facts about performance, though suggestive, do not automatically show that humans are naturally irrational or sub-optimal in matters epistemic, insofar as we can distinguish between a person's natural reasoning *competence* and her *performance*. This distinction, which derives from Chomsky's cognitivist understanding of linguistics, has been linked explicitly to the rationality debate by J.I. Cohen. Cohen goes on to argue that, given that the norms of reasoning are derived using the same methods as are used for determining the nature of competence, it is impossible to show that we are irrational by experimental means. Even though one might not accept precisely this line of thought, the distinction between performance and underlying competence is a useful one for framing the central question of the rationality debate, namely, whether we are or are not naturally rational.¹⁹ Using the distinction, we can ask to what extent the empirical data support the idea that our reasoning *competence*, and not merely performance, is 'irrational', that is contains principles that diverge from the normatively correct ones.

Now, it might seem obvious that if the rationality debate ends up delivering a positive answer to this question, then P (and NP) is in deep trouble. However, there is something odd about even raising the issue in relation to P. For in doing so, we seem to be presupposing a set of criteria for epistemic rationality that are *not* linked to our reasoning competence – and thus seem to beg the question against P. In other words, it seems we cannot evaluate P by experimental means: once we tentatively accept it as a hypothesis, the idea that we are naturally rational becomes unimpeachable, since it is precisely the principles of our reasoning competence that are meant to constitute the norms of epistemic reasoning.²⁰ To an extent, this merely reflects one of the difficult conceptual problems surrounding the rationality debate, and that many have sought to articulate: the problem that the norms of reasoning, since they are somehow ours, cannot be used to pronounce on *ourselves* that we are irrational, however plausible such a conclusion might seem, given certain experimental evidence. Nevertheless, P forecloses the possibility of experimental demonstrations of irrationality more definitively and automatically than most other views. In view of this, the best way to proceed would seem to be to see the rationality debate as raising a *prima facie* legitimate question, answering which in the positive – concluding that we are naturally rational – should be regarded as a necessary condition for P to be viable (though not of course sufficient). This is not to say that P's automatic foreclosing of other answers might not be justified in the end, but we need at least to think our way through to how we could accept this.

What reasons, then, are there for accepting, or rejecting, the rationality thesis? As far as reflecting it is concerned, the main argument is simply the mass of empirical data to the effect that we would seem to reason sub-optimally in relation to traditional or widespread normative standards. The task for those who would want to uphold the rationality thesis is thus to show that we can reasonably set aside these data as reflecting just performance errors. Stein considers three broad lines of thinking aimed at doing this, based, respectively, on the idea of reflective equilibrium, on the idea of charity and on the idea of evolution. I will consider these briefly in turn.

The first is basically Cohen's argument: since the method of reflective equilibrium (i.e. MRE) is used both to uncover the nature of rational competence and to uncover the correct epistemic norms, we

must be naturally rational. Against such an argument, Stein considers two broad kinds of objection: on the one hand, one might claim that MRE is not the sole or correct method for uncovering epistemic norms; on the other hand, one might claim that it is not the sole or correct method for uncovering competence. Stein is broadly sympathetic to both kinds of objection. To start with, evidence from experimental psychology and, in principle, neurophysiology, will also be relevant to determining the nature of reasoning competence, in addition to intuitions. Secondly, a simple application of MRE is not likely to yield the correct set of epistemic norms. One needs, it seems, in addition to ordinary people's intuitions, to take account of experts' intuitions, empirical findings from psychology, and general philosophical theories.²¹ Ultimately, Stein argues that principles arrived at on the basis of a naturalistically informed, wide process of MRE are likely to be the best we can make of a theory of correct norms for reasoning.²² I think Stein succeeds in dislodging specifically Cohen's argument. However, given the cogency of my arguments against all of Stein's substantive ideas (cf. Sections 2.2, 3.3 and 5.1), coupled with a strongly naturalistic rejection of apriori sources in determining norms, the general idea that some form of MRE, widened to include empirical data from psychology, should be a plausible method both for uncovering competence and arriving at a correct theory of norms survives.

The second defence of the rationality thesis Stein considers derives from the so-called *charity* arguments associated most famously with Quine and (perhaps especially) Davidson, to the effect that finding rational coherence in a person's beliefs is a precondition for attributing them beliefs at all. If someone does not generally or fundamentally accept that, say, beliefs of the form *P* and *if P, Q* oblige them to form a further belief of the form *Q*, it will be difficult, or impossible, to attribute them beliefs at all; conversely, given they can be attributed beliefs, they must be seen as cleaving to principles like *modus ponens* fundamentally, that is, at the level of their rational competence.

There is an awful lot that might be said about these kinds of argument. One issue revolves around the fact that, since the arguments concern first and foremost the constraints on the *attribution* of belief, it is not clear what implications they have for belief-*possession* and for the principles that actually steer belief-formation – which is what the

rationality thesis seems to concern. Stein is relatively unpersuaded by the idea of apriori constraints on *possessing* beliefs, and hence is sceptical towards the idea of apriori constraints on attributing them. For him, therefore, the charity arguments for the rationality thesis fail. I do not wish to take a stand on precisely this issue. For my purposes, it is significant that there are those who claim that normative constraints apply not merely to the attribution of belief, but also to the possessing of them: having a belief is, as a matter of fact – moreover, of naturalistic fact – inextricably bound up with the drawing of correct inferences. There are normative facts out there in the world, inhering in our natural psychological make-up. In Section 5.4, I will consider this idea in connection with our discussion of NP.

The final kind of defence of the rationality thesis that Stein considers is based on the idea that our rational competence is a product of evolution; *hence* selected for on the basis of its adaptive features; *hence* conducive to forming true beliefs; *hence* rational.²³ He offers an extensive discussion of the prospects of such a 'Darwinian' vindication of the rationality thesis, canvassing the general debate concerning the prospects for adaptationist explanations of biological traits in general, and of psychological capacities as particular examples of such traits. We do not need to go into these issues here; nor does Stein view them as bearing decisively on the Darwinian defence of the rationality thesis.²⁴ What he does think bears decisively is the fact the Darwinian argument does not succeed in showing that having substantially true beliefs is necessary or even optimal as a way of securing behavioural and reproductive success. Processes of natural selection select traits in relation to their survival potential, and there seems no reason to think that in certain circumstances false beliefs – with certain other properties – should not have been just as apt as expediting our survival as true ones. Hence there is no reason to think our reasoning principles should not generally lead to false beliefs given just the constraints of evolutionary adaptation; and hence no reason to think our reasoning competence is 'rational'. On the other hand, to simply equate survival with what is rational is intuitively implausible.²⁵ I would only add that at the very least it empties the notion of epistemic rationality of all interesting content.²⁶

However, I think Stein (along with many others)²⁷ fails to say enough about the first horn of this supposed dilemma. A rejoinder could first draw attention to the kind of mental structure which

beliefs are standardly taken to be involved in, namely, the production of successful action in concert with desires – a structure which we may assume in the present context has itself been selected for through evolution. If one reflects on what this structure involves, then the idea that false beliefs might do just as well in securing our goals as true ones will seem far from cogent. If I desire that some state of affairs *S* should obtain, how else I am in general likely to succeed other than by performing some action or other that leads to *S*? And if that action is going to be partly an upshot of a belief, what kind of belief will in general lead me to perform an action appropriate to bringing about *S* but a true belief to the effect that one of these actions leads to *S*? Of course, *S* might obtain anyway, fortuitously for me in some particular case, even if I do not have such a belief. But given that belief has the role in action-production just described, and given that we do generally succeed in our actions, it seems that our beliefs cannot, barring cosmic coincidences, be other than true in the vast majority of cases.

Someone might object that there is an important ambiguity in the notion of 'behavioural success' which the above argument builds on: are we talking about behavioural success as we, the agents, see it, or success in biological terms? Surely, as far as evolution is concerned, it can only be the latter that is of any importance. I may frame a desire for my wife to play the piano, and, as far as I am concerned, succeed in fulfilling this by whistling a melody I like her to play. But why should I not be radically deluded about what it is I am actually achieving by doing this? As far as adaptive success is concerned, so long as my thought processes lead to something I gain from in terms of reproductive potential, it seems irrelevant whether I have actually fulfilled my own personal projects.

However, this argument is not very forceful. It may be true that evolution is only ultimately concerned with survival, but it is also concerned to secure survival at minimum costs. If our belief-desire system had no connection to what is out there, there would surely be no point in a belief-desire system at all. So, given that we have such a system, and that it is adaptive, it seems there are good reasons to think that most of the beliefs that get used by it are true – and hence, perhaps, that our reasoning competence is rational.

Having said all this, I do not think we should place too much weight on the evolutionary argument for the rationality thesis. To

begin with, it turns on extremely controversial assumptions of evolutionary psychology concerning the need for adaptive explanations of complex psychological mechanisms. Secondly, even if one accepts these assumptions, and also my arguments above, it is not clear that they are enough to establish the rationality thesis, for two reasons. First, even if evolution has ensured that our reasoning competence spits out mostly correct beliefs when these are to be co-ordinated with desires, it is plausible that there are many beliefs – in particular, scientific beliefs about fundamental particles and principles of nature – that have no such role. Why should evolution care about the truth of these? If it does not, maybe our natural reasoning propensities are inadequate to the task of uncovering this kind of truth. Second, even if most ordinary beliefs churned out by our reasoning competence are true, it may still be that there are large numbers that are not, and this may be sufficient to impugn the rationality of the competence underlying their production. To use a crude illustration, it may be a principle in our competence to treat as dangerous anything that looks mildly threatening. This will lead to lots of false beliefs, even though in general our belief-forming mechanisms are reliable. But it may still be enough to warrant the conclusion that we are not naturally, at least optimally, epistemically rational.

In relation to our concerns, the two most significant defences of the rationality thesis are those connected to MRE and the principle of charity, both of which we have seen Stein's arguments fail to dislodge. In any case, in the following sections I will simply assume the rationality thesis. For it strikes me that both P and NP are susceptible to other, more fundamental objections considered as potential theories of genuine epistemic norms.

5.3 Can norms be principles of our reasoning competence?

According to P, we should identify the correct set of epistemic norms with the principles in our natural reasoning competence. Now, as suggested several times already, the idea that the norms of reasoning ultimately reside in, or are intimately dependent on, some faculty or faculties of the mind is not original with P. Indeed, it might be seen as the leading idea of the whole of the so-called modern period in philosophy (i.e. from Descartes to Kant). Given that, the idea of a

naturalistic study of the mind as the heir to apriori epistemology is historically motivated and perhaps even intuitively appealing. On the other hand, the very fact that we now eschew the kinds of truths and capacities that the moderns worked with also casts doubt on whether the replacement can really be quite that simple. In this section, we will investigate to what extent the replacement is plausible by considering different kinds of objection to it.

A first objection might be that psychological principles could not be epistemic norms because, although they may be said to guide us, the sense in which they do this is not the right one. A norm should guide by being *consciously* apprehended by an individual who seeks to bring her thought processes or behaviour in accord with it, whereas those of our reasoning competence are meant to guide us at a pre-reflective, unconscious level. However, even if this intuition reflects some aspect of ordinary language usage, it does not seem, in and of itself, to bar us from seeing unconscious principles as issuing in behaviour that is, in virtue of being steered in the appropriate manner by these principles, correct or rational; and thus it is unclear why conscious following should be a prerequisite for being a norm. Consider further that if the principles of our reasoning competence were consciously known, they could presumably be employed in conscious thinking just like 'ordinary' norms (even though, of course, they would not have to be for one to reason optimally) – which suggests that ruling them out as norms just because they are unconscious is arbitrary.²⁸

The objector might nevertheless pursue the issue. She might argue that genuine norms must be such that it is under the conscious *control* of the individual in question whether or not the norm is followed. If norms are constituted by psychological structures that *causally* issue in behaviour, then the individual really has no choice over whether she follows them or not.

To this it should first be replied that, given the separation of rational competence from performance (more on this below), there may be a sense in, and an extent to, which individuals *can* choose to follow their natural principles of reasoning, namely, by making efforts to bring their performance in line with their competence. There may also be room to question the idea that principles of competence really do *cause* behaviour in any straightforward sense. However, it is not even intuitively clear that norm-governed

behaviour must be non-causal in the sense of being under the active control of an individual's free will. Acting spontaneously and non-deliberatively would seem to be fully consistent with acting in accord with rational norms, and not just in the sense that one might fortuitously hit on the right behaviour.

A second kind of objection might run as follows. If norms are simply psychological principles which all normal human beings possess, and which guide their behaviour, how can we genuinely *evaluate* people's inferences? If being rational is just something people are naturally disposed to be, how can we evaluate behaviour that issues from the relevant dispositions any more than we can evaluate people in relation to activities like eating, sleeping or defecating? Yet it would seem that being able to evaluate others' behaviour as rational or irrational is part and parcel of what epistemic normativity is about.

In answering this objection, a defender of P should avail herself of the competence-performance distinction: psychological principles that are viewed as epistemic norms need not, and will not, always be reflected in behaviour. This will make it straightforward, in principle if not in practice, to classify certain behaviours as irrational or incorrect, namely, those in which reasoning competence has somehow failed to be applied – at least properly – to the task in hand. The idea of a competence-performance divide plausibly also allows P to draw a further distinction between irrationality and non-rationality, the former applying to humans (and possibly other creatures) when their reasoning ability lapses, the latter applying to creatures who lack a reasoning competence altogether (as well perhaps as to humans who refuse to reason or are incapable of doing so).

It might now be objected that the defender of P must implicitly employ a normative notion of *correctness* in explaining what a performance error is – that a performance error, being an *incorrect* or *improper* application of competence, requires, essentially, an external evaluation to establish, independently of any 'internal norms' that might be 'followed'. This objection is interesting, but ultimately fails. This is not because a performance error can simply be *defined* as a piece of behaviour at odds with whatever principles constitute competence. To start with, it seems reasonable to deny that behaviour that merely *by chance* accords with the principles can be competent; moreover, the reply does not explain why performance at odds with competence is to be counted as *erroneous* (and competent behaviour

as *correct*). Another response to the objection might be to say that the evaluation as to whether the *use* of competence was correct would be based on a further application of rational competence (which is of course true for a supporter of P). Some might feel that any residual problem with this response would only be 'epistemological' (i.e. a matter of finding out whether competence was correctly applied on any particular occasion); however, this neglects the fact that the notion of *competent* performance is essential to the very articulation of P's understanding of the difference between correct and incorrect reasoning. Having said that, the notion is, plausibly, a *functional* one: we perform competently in a particular domain when we function correctly in that domain; what this amounts to will be a further, empirical question. Hence, from the kind of naturalistic perspective we are assuming, the objection can after all be turned aside, since the idea of function is one which the empirical research programme of cognitive science – in accord with much of the rest of biology – makes use of (and, surely, independently of any commitment to the more controversial adaptationist paradigm).

Accepting all the foregoing, plus our background naturalistic assumptions from Section 1.5, it might suddenly seem as if P is in the clear. In accordance with naturalism, we must assume that our basic cognitive practices are correct; in accord with our response to the problem of normativity, we are assuming that naturalism does not *ex ipso* rule out genuine normativity. In the previous section, we found no overwhelming argument for believing that we are fundamentally irrational by nature. In the previous paragraphs of this section, we have motivated P as the natural heir to the conception of epistemic norms we find in modern philosophy, argued for accepting the idea of unconscious rule-following, and found place for the distinction between correct and incorrect behaviour that seems essential to the evaluative practices of rationality. What more could be said to vindicate the principles of our reasoning competence as norms?

The answer to this question is, I think, nothing. However, that does not show that we *have* to regard the principles in question as epistemic norms, or that we *should*, taking everything into consideration, do so. It is admitted on all hands that these principles, at the most basic level, are simply descriptive principles or states-of-affairs that somehow underlie our reasoning behaviour. They are mental states, of some kind or other, posited by a thoroughlygoingly naturalistic

cognitive science. Given the rest of our assumptions and arguments, we will be able to say that these principles are what most fundamentally *give rise to* (genuinely) correct reasoning; and that the latter is precisely what is obtained when the principles function as they should, that is, issue in behaviour undisturbed by performance limitations. However, I cannot see that this in itself suggests we should see the principles – in and of themselves – as *norms*. Again: we must, by our naturalism, assume that our basic cognitive practices are fundamentally in order – that we reason as we ought to; hence, for supporters of P, we will *naturally* reason as we ought to. But that does not mean that the natural *ground* of these correct reasoning practices must also be seen as *rules specifying what we ought to do*. That the principles are embedded in a functional organisation makes no difference to this, for this organisation does not consist in the principles specifying how they should be applied.

Furthermore, one should recall that in our discussion of the problem of normativity in Section 1.5, we did not conclude that there was in general *no problem* in deriving something normative from something descriptive. The point was that, to the extent naturalism embodies the idea of correct reasoning at the outset, we should, at least *pro tem*, assume that we can make sense of normative notions like correctness. This means that with theories in the categories of AP and AF that see norms as explicit normative statements, we do not balk at the legitimacy of such statements. But clearly this does not mean that there is nothing problematic about regarding what in the first instance are merely descriptive principles as norms – as somehow embodying content to the effect that *such and such should be done*. And on the face of it, nothing would seem to warrant or make sense of such content-attributions to principles of the reasoning competence.

In the spirit of naturalism, I believe one should avoid laying down a priori strictures on anything – including, thus, what can count as constituting a normative state-of-affairs and what cannot. Perhaps, then, it is legitimate for P to help itself to the idea that the principles of reasoning competence are norms, without saying anything very specific as to how this claim is justified or what it means (at least, beyond what we have said above). Nevertheless, I think this kind of stipulation would be nugatory. To begin with, it is in tension with what seems reasonable. Secondly, given our background naturalistic

assumptions, there is absolutely no need to do it. Let us tell the story as follows. We start to form beliefs about the world – on the basis, unbeknownst to us initially, of our natural reasoning principles. Our first beliefs concern a world of physical and chemical phenomena, which we gradually extend to an understanding of biological organisms and their properties. It later turns out (supporters of cognitive science would hold) that this world also has substantive psychological properties: our perception and categorisation of the world, as well as our use of language, are susceptible to systematic empirical study, parallel to that of physical, chemical and biological phenomena. We then identify, amongst the capacities of the mind, one for drawing inferences and forming the beliefs we do, and have already done, by and large correctly. Now, though perhaps we *could* call the principles of this capacity ‘epistemic norms’, and perhaps even say that they specify what we should do – does this really add anything of import to just saying that we reason, fundamentally, in such and such ways, just as the world otherwise behaves in such and such ways? Even though we must also assume that the reasoning they give rise to is correct, it seems to me that the answer here must be negative, for this fact is not something that naturalism itself requires an account of. Naturalism, let us concede, *allows* for P’s stipulative response. But, insofar as this response is strained at the outset, it also *shows* that it is quite unnecessary.

This concludes my discussion of P. In some sense, we can class its failure alongside those of AP and AF, for all three fail to show that we need a theory that specifies genuine epistemic norms. With AP and AF, the need concerns the norms themselves; with P, what we do not need is to see our (psychological) theory as concerning norms (even though the theory itself is, insofar as we actually possess a reasoning competence, required as part of the research programme of cognitive science).

With this conclusion, the main argumentative burden of this book is discharged. The three varieties of NE I have considered are, as I see things, exhaustive of the potential varieties, given my conception of naturalism, and of what it is for something to be natural: being, in some broad sense, describable by science. However, there is a variety of position which sees itself as naturalistic, and which also views our psychology as the basis of our normative practices, but which would be untouched by the above objections to P. This position sees our

natural psychological make-up as itself inherently normative, and hence the kinds of problem we have been considering do not even arise. In the last section of this chapter, I will discuss this *normativist* form of psychologistic NE.

5.4 Normativist psychologistic NE

NP, as we are calling it, is not, I think it is fair to say, a branch of NE that many naturalists or supporters of NE would include in an overall survey of the available options. Behind NP we find the idea that the faults and failings of much contemporary (as well as modern) philosophy lie precisely in identifying what is natural with (roughly) whatever can be described and accounted for by science. According to some thinkers – in particular, John McDowell, whom I view as the chief exponent of NP’s conception of naturalism²⁹ – such an identification leads to crude forms of either reductionism/eliminativism on the one hand, or dualism on the other: those who accept that what is ‘natural’ exhausts the realm of being opting for the first ‘horn’, those who do not opting for the second. The way out of the dilemma is to reject the conception of naturalism common to both reductionism and dualism: to see our way to accepting the presumed problematic aspects of our distinctively human lives as fully part of the fabric of reality, in unreduced and unreconstructed form.

As noted in the introduction to this chapter, McDowell does not motivate his ideas via epistemology; at least in the sense with which we are concerned, but rather through what has come to be termed *the philosophy of content*, that is, the nature of thought about the world in general. According to him, for thought to have objective bearing on the world, it must be connected to it *via* perceptual faculties, but not in such a way that this constraint in any way compromises the essentially normative nature of thought (i.e. the idea that we can evaluate thoughts as true or false, rational or irrational). Perceptual capacities are realised in us in such a way that they automatically engage our rational, conceptual capacities. However, this seems to create a problem insofar as we share our perceptual capacities, to a large degree, with other animals that are not endowed with the kinds of conceptual capacities we possess. So the question becomes how we can be constrained by the nature we share with animals without being ‘swallowed’ by it, that is, reduced to it – or,

alternatively, without *also* existing in some supernatural sphere, as philosophers through the ages have been inclined to say we do. The way out, for McDowell, lies in diagnosing, and then exorcising a non-mandatory conception of the natural that underlies both the recoil to what he calls 'bald naturalism' – the view that nature is co-extensive with the realm of natural law – and to dualism (or 'rampant Platonism', as he calls it). We must see nature as embracing more than what 'bald' nature embraces, and if we do this, we will with equanimity be able to regard ourselves and our capacities as fully natural, even though they are not part of the nomological world that biology and physics describe. McDowell connects this conception of naturalism to Aristotle's discussion of ethics and his notion of 'second nature' from book 2 of the *Nicomachean Ethics*.³⁰

This kind of naturalism also connects to another prominent theme in McDowell's philosophy – which in turn links it to our theme and NP. This is his insistence on the inherently and irreducibly normative nature of psychological explanation, that is, in relation to rational or intentional events and behaviour.³¹ Explaining the rational formation of a belief in the face of evidence, or a rational action on the basis of an evaluation of preferences involves a *sui generis* form of understanding that has no echo in nomological, physical theory. As McDowell (famously) puts it:

[P]ropositional attitudes have their proper home in explanations of a special sort: explanations in which things are made intelligible by being revealed to be, or to approximate to being, as they rationally ought to be. This is to be contrasted with a style of explanation in which one makes things intelligible by representing their coming into being as a particular instance of how things generally tend to happen.³²

The link between this and McDowell's naturalism can perhaps best be appreciated by contrasting his view with that of Davidson, whose views on psychological explanation he endorses to a great extent (indeed, is largely defending in the paper the quote is taken from), but also significantly diverges from.³³ For Davidson, intentional states and events are genuine denizens of the world, but the referents of locutions such as 'the belief that *p*', 'the desire that *q*' and so on are not fundamentally identified with respect to their psychological

properties, but rather their physical ones. In a similar vein, psychological explanation takes its place as part of an overall scheme of understanding that is first and foremost *interpretive*, that is, that involves *attributing* certain states to people in order to *make sense* of them, but is not, at least straightforwardly or directly, concerned with latching onto or describing genuine patterns in nature. Davidson, in other words, seeks to motivate a stance on the mental that essentially respects the naturalism of 'bald naturalism', without thereby, being reductive. McDowell sees, however, tensions in Davidson's theory, and insists that, unless propositional attitude-attributions and the patterns they enunciate are genuinely seen as 'carving nature at its joints', they will be in danger of lapsing into something second-rate with respect to physical reality – or else of having to be consigned to some supernatural realm. His alternative, Aristotelian conception of nature is meant to account for the irreducible reality of our intentional psychology, along with other aspects of our humanity that are inherently normative, such as aesthetics and morals. At the same time, the fact that our everyday life is so full of explanations that are (allegedly) irreducibly normative in character would seem to go some way towards legitimating this kind of naturalism, insofar as such explanations presumably latch onto some or other, and moreover quite 'natural', tract of reality.

What I am calling NP seeks to exploit this understanding of what is natural, and its connection with normative explanation. For these ideas would seem capable of motivating a conception of epistemic norms as, at least in some broad sense, part of our (natural) psychology, and at the same time as possessing a primary and irreducibly normative nature. Insofar as our psychology is not part of 'bald' nature, but nevertheless part of nature *tout court*, we can make sense of the idea of genuine but also naturalistic epistemic norms regulating our belief-forming practices.³⁴ In other words, the normative nature of psychological explanation can *vindicate* the idea of principles that are both naturalistically kosher and genuinely normative. Indeed, given McDowell's naturalism, it would seem that epistemic evaluation or prescription, on the one hand, and certain kinds of psychological explanation, on the other, pick out a unitary fundamental phenomenon: explaining why someone forms a belief and justifying it by reference to certain norms become simply two different ways of 'latching onto' the same, irreducibly natural but also normative process.

In fact, I think more needs to be said than I am going to say here fully to justify identifying the referent of these two activities with one another. As things stand, McDowell's conception of psychological explanation does not seem to commit him to any *sui generis* general principles for rational explanation (i.e. something that would play a role parallel to that played by natural laws in scientific explanations), and thus it is not clear he would countenance the idea of naturalistic and psychological norms for belief-formation either. Now, I think it can be argued that, without the idea of such *sui generis* principles, the idea that rational explanation is of a different character from other, ordinary or scientific kinds of explanation cannot be upheld.³⁵ But even without doing this, it seems in any case reasonable for NP to trade on the mooted identification of explanation and evaluation. To begin with, McDowell often talks as if there are identifiable rational principles, for example, when he speaks of 'a priori rationality constraints'.³⁶ Moreover, his disagreement with Davidson does not seem to concern the existence of such principles, but rather their status; and Davidson, or at least many of his followers, refer quite freely to 'rational principles' – not least in connection with the rationality debate and the argument from charity for the rationality thesis (cf. Section 5.2).³⁷ In sum, whether or not McDowell is committed to the existence of such principles, he certainly seems to see it as compatible with his position; and in any case it appears to be compatible with his position. *Prima facie*, then, it would appear NP can exploit McDowell's understanding of both the natural and the psychological explanations in order to motivate its conception of naturalistic, genuine epistemic norms. (McDowell is also, needless to say, no relativist and sees the fundamental aspects of human nature as universal, albeit not purely biological in character.)

There is no doubt many things one might want to say about this kind of proposal, and in particular about the conception of naturalism that McDowell recommends. Some might feel that this conception – whereby nature becomes, as McDowell puts it, 'partially re-enchanted' – simply makes no clear sense. I am not unsympathetic to this line, though I think it needs to be argued more systematically. I will not attempt a general argument to that effect here. Instead, I will concentrate on repudiating NP by way of repudiating the coherence of the idea that psychological explanation has the *sui generis* character that McDowell claims it to have – at least insofar as it

involves principles. In doing this, I will have shown that McDowell's naturalism can offer no succour to a supporter of NE; for it is precisely the closely related understanding of (principle-based) psychological explanation that makes available the conception of norms which NP needs. This does not show that McDowell's naturalism is incoherent *tout court*, for it is not clear, as we have noted, that it is committed to the idea of codifiable rational principles. Nor, perhaps, is it even obvious that his naturalism as a whole depends on their being *sui generis* forms of explanation of the human world at all. Nevertheless, insofar as it is meant to save NE (via NP), it seems clear that it must have these implications (I think, as noted, that it has them anyway, indeed, that it gains most of the intuitive plausibility it has from the idea of such *sui generis* explanation; but I will not argue that here.).

My argument against McDowell's conception of psychological explanation is that is either incoherent – makes no sense – or not substantive, that is, not clearly distinct from that of ordinary or scientific explanations that are non-normative.³⁸ Recall first the quote from above. According to McDowell, when we explain actions and beliefs, we are making them intelligible by showing that they conform to what rationally ought to be the case. This is contrasted with making something intelligible by, roughly, assuming it under a general law. Factoring out the talk of intelligibility, this comes down to saying that intentional explanations explain a belief or an action by revealing it to be as it rationally should be.

On the face of it, however, this suggestion does not seem coherent. Let us imagine person X comes to believe B on the basis of believing A and if A, B. McDowell's suggestion as to the explanation of this episode would seem to be this:

- (1) X believes A
 X believes if A, B
 If one believes A and if A, B, then one should (rationally) believe B
 B
 X believes B.

The problem is just that the third 'premise' seems to be, at the very best, redundant.³⁹ How can the fact that someone *should* do something – in and of itself – help to explain *the fact that they do something*? Note my

objection is not that the explanation would only work were we to add something about it generally, that is empirically, being the case that people believe *B* when they believe *A* and *if A, B*; or that they generally do what it is rational to do (though I am not denying this either).⁴⁰ Perhaps the explanation is fine just as long as one deletes the normative statement. My point is simply a registration of confoundment at the idea that such a statement can be independently relevant in a psychological explanation.⁴¹

But does (1) really get across the form of the kind of explanation that McDowell intends? Another passage of McDowell might suggest something slightly different:

Intentional explanation makes an action unsurprising, not as an instance of the way the world works (though of course it does not follow that an action is not that), but as something which the agent can be understood to have seen some point in going in for (emphasis added).⁴²

This suggests that the right form of the explanation of X's belief in proposition *B* would be the following:

- (2) X believes *A*
 X believes *if A, B*
 X believes that if one believes *A* and if *A, B*, then one should believe *B*⁴³

 X believes *B*.

This is intuitively much better than (1). Moreover, it gets across the idea that when one forms a belief or acts rationally, one sees oneself as doing what is rationally required of one. However, the question arises now as to whether (2) is any way special as a kind of explanation.

In (2) we take as part of the explanation that X sees her belief-*B*-formation as in some sense 'the thing to do'. But now how do norms figure in this explanation? One might say that X has a belief about a norm of inference, about what the relevant norms are: she has a belief about what one *should* do in certain circumstances. But clearly this is not the same as making something 'intelligible by being revealed to be [...] as [it] rationally ought to be'. If X had some other, *false* belief about how things rationally ought to be, then a similar

explanation of why she formed some other belief, *C*, would still work, viz.:

- (3) X believes *A*
 X believes *if A, B*
 X believes that if one believes *A* and if *A, B*, one should believe *C*

 X believes *C*.

Here we can hardly say that the *explanandum* is revealed to be as it rationally ought to be, even though it is – plausibly – explained. Hence it is hard to see that there is anything essentially normative about either it or (2).⁴⁴

At this point, one might seek to defend McDowell by denying that one needs to mention any general rational principle, either as such or as the object of a belief, in a psychological explanation, and that it is only this that gives rise to the dilemma I have presented. However, as noted, this line, even if it were cogent (which I do not think it is), is not germane to our concerns, since the prospects for NP depend precisely on their being identifiable principles used in explanation.

An alternative way of reacting might be to insist that I am simply begging the question against McDowell and NP. Both these lines involve the idea that there is no sharp divide between the normative and the descriptive, whereas in my objections it might be felt that I am insisting on precisely such a divide. However, this would not be a fair appraisal of the dialectical situation. As I see it, NP and McDowell's naturalism, to the extent this can support NP, rest on the *intuitive* plausibility of the idea of normative explanation in intentional psychology, which is taken as a kind of datum. Now what is intuitive can no doubt vary from person to person. However, it strikes me that insofar as intuitions are relevant to philosophical discussions, they should at least be amenable to being coherently spelled out. The two quotes we have given from McDowell do not even seem to express the same idea. Moreover, given that neither presents a coherent notion of a special kind of (principle-based) psychological explanation – something which, as far as I can see, is as obviously the case as anything could be – it seems the intuitive support that McDowell and NP seek from the nature of psychological

explanation simply does not exist. Perhaps this does not show completely definitively that NP is incoherent. However, insofar as there are many who would say this anyway, without further ado, and insofar as it lacks the intuitive support it is meant to have, I think I have said enough to justify dismissing it from further consideration.

In this chapter, we have discussed psychologism generally in epistemology and critically assessed what I call 'mixed' varieties of NE (including weakly naturalistic theories more generally). The idea that psychology can normatively constrain our reasoning and physical theories of the world appeared defensible, a view which is also consistent with Quine's idea that substantive epistemology should be replaced by a purely descriptive psychological science. Neither of these ideas, however, entails that a theory of epistemic norms can partially or wholly draw on the facts of some natural, psychological capacity. This is not, I have argued, because we probably are other than naturally irrational. Rather it concerns the fact that it is only dubiously coherent to view principles of our reasoning capacity as anything other than descriptive facts underlying our reasoning performance, and the fact that there is no motivation for doing so. In essence, this is the problem for both P and NP: P seeks to regard descriptive, psychological principles as norms, but it is not clear how it can do this coherently, and there is at least no motivation for doing so, given our general naturalistic background assumptions. NP seeks to motivate a view of the natural as already norm-imbued, and thus to view natural psychological principles as normative; but the only coherent motivation for doing so is the supposed normative nature of psychological explanation, and this turns out to be a chimera. In sum, psychological NE is in at least as bad a state as the non-psychological versions we considered in Chapters 3 and 4.

Conclusion: Science without Norms

In this book, I have been arguing for my overarching thesis that, to the extent we can make sense of them as genuine, we have no need for naturalistic epistemic norms, either in everyday life or in science. The arguments assume a certain understanding of naturalism that can be motivated as a reaction to the failures of what I called TE. The project of NE builds on naturalism, but seeks to go further in specifying concrete norms for enquiry. I have considered several broad varieties of theories of such norms – different varieties of NE – which I take to be exhaustive of the possibilities: AP, AF and P (and, supplementarily, NP). In relation to AP and AF, I have not, for the main part, argued that their conception of norms is incoherent, or that the norms they propose cannot be justified in the ways they suppose they can be, but simply that, given the assumptions necessary to justify the norms, it turns out that they are not in fact required for optimal belief-formation. In relation to P and NP I have argued that it is only dubiously coherent to see natural psychological principles as norms; and that there is, partly because of this, no motivation for doing so against the backdrop of naturalism. I have also given reasons for thinking that any more weakly naturalistic theory of norms that seeks to maintain a role for distinctively apriori thinking in epistemology at the same time as taking account of empirical resources can only make sense of norms necessary for optimal belief-formation insofar as these correspond to the apriori information.

In this concluding discussion, I want to underline the significance of these claims and to discuss their implications for our actual epistemic practice and for philosophy.

In the first place, it is very important to be clear about the role of the underlying assumption of naturalism. As I have presented naturalism, it is essentially a negative doctrine that rejects the pretensions of first philosophy, that is, TE and the idea of apriori epistemic norms. What is left is the idea that our epistemic situation is necessarily embedded in ongoing practice – in our current theories of the world and our cognitive or epistemic practices, which must be assumed to be fundamentally in order (though not necessarily susceptible to a systematic account). Given a typical set of beliefs, it would seem ineluctable that there are some things that one *ought* or *ought not* to believe on the basis of these beliefs; and that, generally, we do as we ought. Surely, if any normal person comes to believe it is raining outside, then it is *reasonable* for her to believe – and unreasonable not to believe, or believe it is not the case that – she will get wet if she goes outside without, in some way, taking precautionary measures. We cannot make sense of a thoroughgoing relativism or (something I take relativism to entail) nihilism about epistemic normativity, nor (something which entails the first thing) a thoroughgoing scepticism about our actual epistemic ability. I do not think it is possible to analyse these commitments further – that is precisely one reason why TE fails. Nor, concomitantly, are they meant as substantive philosophical commitments, but merely to register – in the spirit of the first great naturalistic philosopher, Hume – certain aspects of our mode of operations which we *de facto* will never seriously be able to renounce.

Given this general background, one might, reasonably, think there is still scope for offering empirically justified, fallible, but nevertheless substantive and specifiable guidance to people and scientists in the form of general epistemic norms; or perhaps that we can transpose the traditional notion of apriori norms to talk of the principles of a natural psychological competence. I have examined what I see as the space of possible proposals that would articulate these ideas, some of which exist as extant proposals in the literature, and shown that, given their conception of what would justify their proposals, the norms they propose would not be necessary (or that it is not necessary to see what are proposed as norms as such). I have also, in

Section 4.6, tried to spell out how this could be so given the apparent reliance on norms in scientific activity. My own understanding of the outcome of this discussion is that, when it comes to scientific activity in the sense of making epistemic inferences and thereby progress towards the truth about the natural world,¹ scientists can make do with descriptive statements of various kinds, along with a serious and concentrated application of their scientific competence, in concert with similar efforts from their fellow scientists. When it comes to ordinary people, much the same applies, though they are restricted to use of a natural or at least non-scientific reasoning capacity, and presumably work less systematically towards their distinctively epistemic goals.

Having said all that, these conclusions must be read with certain qualifications in mind. The following points, most of which we have already broached, do not, I believe, compromise the significance of my overarching thesis, but are nevertheless worth stressing:

1. I have nothing in particular against continued talk in terms of 'epistemic norms' insofar as certain deeply entrenched propositions in science may play the role of non-negotiable standards which will normatively condition any inference they are relevant to. These propositions may concern fundamental theoretical tenets, like $E = mc^2$, or well-known methods for obtaining certain kinds of observational data, or even (in some future scenario) information about our psychological capacities of perception, memory and/or reasoning. It is a trivial matter to convert these propositions into norms, viz. 'Respect proposition/fact *P* in your reasoning!'. All the same, since we must, according to the web-of-belief model of science, in any case take account of these factual propositions, it is not necessary to exalt them as explicitly normative statements.
2. Instructions or rules of a procedural nature may still be necessary for practical matters, for example, for the instruction of technical assistants who are not concerned with drawing scientific inferences; or more generally for people whom I called 'philistines' that are indifferent to the goals of cognitive enquiry, but whom we need to perform cognitive tasks. However, as suggested in Section 4.5, these rules would not play an epistemic role for such people, and would arguably not class as epistemic in such a role.

(In this connection, it should also be remembered that I have not been arguing against the need for rational norms *tout court*, but only in relation to theoretical rationality, i.e. forming beliefs and drawing inferences.)

3. One might have supposed that quasi-foundational naturalistic norms in science (i.e. AP-norms conceived of as derived through DRSPS) could be of use in instructing ordinary people who want to become scientists (or more scientific), even though the scientists themselves, in obeying them implicitly, have no need for them. I have argued against this idea by drawing on recent criticism of the idea of a monolithic set of rules for science, and have backed this up with Kuhn's conception of scientific competence as something that does not consist in following a set of abstract, specifiable rules, even at a tacit level. If one demurred at Kuhn's ideas, and could mount a strong case against the recent studies suggesting the variability of methods in science, then there could perhaps be a limited role for norms of science to play, namely that of instructing and educating neophyte scientists. However, even if one accepts this – which is a big 'if', in view of the implausibility of the monolithic conception of scientific norms – my thesis would still hold good to the extent that norms would be necessary neither for doing science, nor in everyday epistemic activity – a conclusion that strikes me as highly significant.

4. I have argued that norms are not necessary for everyday reasoning or science. But what about for philosophy or philosophers? Well, I have been assuming that, insofar as philosophers are concerned with specifying norms, they have viewed their activity as groundwork for science – as part of 'methodology', as it is (or was) commonly called – or for normative guidance to the man- or woman-in-the-street.² If someone insists there is some inherent interest in or point to uncovering epistemic norms other than that of offering guidance – perhaps that of yielding a kind of self-understanding of epistemic activity – that is not something I will quarrel with here. Nevertheless, it seems clear that my arguments should at least prompt naturalised epistemologists to consider more thoroughly the point and conception of their activity. At the end of Section 4.5, I also noted that some who call themselves naturalised epistemologists might be

satisfied with the role of detailing empirical connections that could be useful to individuals in decision-making aimed at securing their epistemic goals. However, as I argued, if this is all it takes to be a naturalised epistemologist, then most ordinary scientists should also be classed among their ranks; moreover, it seems clear that the common conception of naturalised epistemology is that it goes beyond this – that it is engaged in specifying something essentially normative.

Summa summarum: we can, and must, because the alternative is unthinkable, assume the viability of our epistemic practices, first and foremost those that concern or connect to modern science; but we do not need norms. Thus the slogan: *science without norms*.

For some, this conclusion will no doubt be reminiscent of that of other philosophers who have reflected on the nature of rules and rule-following. I am thinking here in particular of Wittgenstein, and of Lewis Carroll through his famous article 'What the Tortoise said to Achilles'.³ Do these philosophers' ideas and arguments have any bearing on what I have argued? In Carroll's witty and consternating tale, Achilles finds himself launched on an infinite regress of finding rules to justify concluding a proposition of the form q on the basis of premises *if p, then q* and p . The Tortoise asks why he should accept q on the basis of these premises, to which Achilles eagerly replies that if one accepts propositions of the form *if p, then q* and p , then one must accept q – to which the Tortoise reacts by again registering scepticism about the validity of the inference on the basis of these augmented premises, to which Achilles responds by adding yet another, and so on and so forth, the ardent hero never appearing to grasp the futility of his attempts to convince his 'slow-witted' companion. There is some debate over exactly what the story is meant to show. Some have seen the point as being that explicit knowledge can never eventuate in any actual inferential *move*, or in action more generally for that matter. This is not something that would seem to impinge on our discussion.⁴ Another central lesson would seem to be that one cannot fully *justify* a logical inference, such as the above, without in some way already assuming its validity. Insofar, Carroll's point can be seen as part of a criticism of TE. NE accepts that we must, in some sense, assume the validity of our existing inferential practices, but takes it that there is still a point in,

and meaning to be made of, the idea of epistemic norms. I do not see that Carroll's argument has anything to say to this, as those of mine precisely do.

The subject of Wittgenstein on rules is far more complex and difficult to evaluate, and not one I can go into in any depth here. A core thought in Wittgenstein's writings on rule-following⁵ is that when one follows a rule in behaviour or inference, there will always be a further question, not decided by the rule itself, as to whether the rule has been applied correctly or not in relation to the particular circumstance. Positing an interpretation of the rule or else another rule to determine this leads only to the question of the correctness of the interpretation or of the new rule's application; and a regress thus threatens. Wittgenstein's diagnosis seems to be not that we must somehow reconstruct the idea of correctness out of materials that do not presuppose this idea, as Kripke famously has suggested,⁶ but that the very idea of correctness, of meaningful application of explicit rules or norms, presupposes some kind of non-explicit normative background: 'there is a way of grasping a rule which is not an interpretation of it, but which is exhibited in what we call 'obeying the rule' and 'going against it' in actual cases.'⁷ This is a conclusion that most naturalists, of my brand, would probably take no exception to. Indeed, much of Wittgenstein's agenda would seem to be aimed at refuting the belief in and the search for apriori rules in meaning and in gaining knowledge. This leaves open the prospects of a substantive NE, and thus my arguments against it are independently significant.

Perhaps some will feel that what has been said in these pages does not penetrate the deeper and more exciting philosophical questions that Wittgenstein, Carroll and others have connected to the question of norms or rules in inference and rational behaviour. For my part, it strikes me that at least one worthy pursuit for philosophy is to find room for strong stands on contemporary issues *without* prejudicing issues at the very heart of the discipline. At the same time, I see my ideas as falling into place as part of a recognisably Wittgensteinian tradition insofar as they seek to *deflate* the pretensions of (a certain kind of) philosophy. In this book, I have been arguing for a conclusion that can be seen as showing that there is no easy way for a modern naturalist to co-opt the traditional conception of epistemic enquiry as norm-steered. Wittgenstein saw

the traditional project as fundamentally incoherent and misguided. Within a naturalistic setting, the project's incoherence is largely shaken off, but the misguidedness remains: the traditional project was not just seriously deluded about the materials available to construct with, but ultimately about whether there was any point in constructing in the first place.

Notes

Introduction

1. More precisely, his seminal article of 1969 'Epistemology' naturalised' (Quine, 1969a).
2. At least two distinct reasons for holding that this project is most central suggest themselves. On the one hand, of the three notions traditionally viewed as constitutive of knowledge – belief, truth and justification – it is arguably only the last, justification, that is the peculiar concern of epistemology; the others being treated of in psychology/philosophy of mind and semantics/metaphysics respectively (cf. Kim 1988: 33 ff.; Pacherie 2002: 300). If justification requires, in some sense, norms, as many would still hold, then a theory of the latter is clearly of the essence in epistemology. On the other hand, many today, and not least naturalists, are doubtful about the significance of giving analyses of epistemological concepts (or any other concepts for that matter), and seek instead to offer accounts of how to improve our concrete belief-forming practices – that is, precisely to justify and specify epistemic norms. I will not go into any argumentation for either of these views here, nor is their tenability something I presuppose. It is enough that the project I focus on is at least admitted to be a central project in contemporary epistemology, which I take it is uncontroversial.
3. This 'scepticism' towards substantive results and/or theories in philosophy is something we first find clearly expressed in Wittgenstein (though not consistently, I believe). It finds expression today in the so-called 'minimalist' programmes in philosophy of language (cf. e.g. Horwich 1990, 1998), whilst Richard Rorty has sought to apply it to a broader field of philosophical endeavour (cf. e.g. Rorty 1980, 1998).
4. I will not, as already suggested, attempt any systematic definition of knowledge – for example, as true, justified belief – nor offer any commentary on traditional or naturalistic attempts to account for what this might amount to. My focus is not on a definition of knowledge, but on naturalistic conceptions of norms for rational belief – whether or not this should count as knowledge (I will not in particular have much to say about truth). For most of the book I will talk of 'knowledge' in the sense of *rational* (or: *rationality formed*) belief, a usage which is not uncommon (viz. the way in which we talk of 'scientific knowledge' without assuming (absolute) truth).
5. See for example, Chalmers (1990) for an overview. The ideas presented in the current paragraph are discussed more fully in subsequent chapters.
6. As it seems have later manifestations of Quine himself; see note 8.

7. When, in what follows, I talk of an 'epistemological theory', I will mean a theory that seeks to give an account of how epistemic norms are to be understood, derived and justified, with a view to actually specifying concrete norms. A naturalistic epistemological theory is one which draws on only naturalistic resources (in a sense I will be specifying).
8. Rejecting naturalised epistemology in this sense does not of course commit one to there being no study or theory worthy of classification as 'a contribution to naturalised epistemology', for example, a naturalistic definition of knowledge, or, for that matter, Quine's conception of science as an ever-evolving web of belief. *A propos* Quine, it is also worth noting that he talks of himself as a naturalised epistemologist and has stressed continuities between his view and the normative aspirations of traditional epistemology. Indeed, many of Quine's writings subsequent to the 1969 article might suggest more of a leaning towards ideas that would be classed under 'naturalised epistemology' as I understand the term in this book. For further discussion of the complexities, see Sections 1.4, 2.1, 4.5 and 5.1.
9. This is in many ways just a handy slogan, and should not be seen as indicating a delimitation of my thesis to science in a narrow sense. The relation between everyday belief-forming practices and those of science will be a recurrent theme throughout the book.

1 Traditional Epistemology and Naturalism

1. As suggested, the points made in this chapter are widely acknowledged and discussed, but my exposition is especially influenced by the work of Harold Brown (especially 1988a: Chapter 1, 1988b, 1996). Two other review articles that I have made particular use of are Matfje (1990) and Kitcher (1992). The precise characterisation of traditional epistemology (and its problems), and of naturalism and naturalised epistemology are, however, my own.
2. Exactly where and how one draws the distinction between types and tokens and/or particulars and non-particulars in a philosophically circum-spect manner is a complex issue that I trust I do not need to go into here. For further reassurance, see the following discussion of my understanding of epistemic norms.
I allude to another thorny philosophical issue in saying norms are 'propositions, principles or statements', which I do in order to remain non-committal about their ontological status. Are norms a kind of truth, perhaps one which we might be able to discover, a special category of *sui generis* object, or just part of our language (or mind)? I will touch on this issue later on in this chapter and in subsequent chapters, without reaching any firm conclusion. Again, none of my original arguments should depend on taking a principled stand on it.
3. I will often refer to norms as 'rules' or 'prescriptions' (and sometimes 'principles'), intending these as notational variants of 'norm'. The important point is that the statements (propositions or whatever) that are norms be

- essentially normative, not just statistical regularities or otherwise descriptive in nature. I am not concerned to defend the choice of the expression 'norm' as the most apposite term to express this idea, and choose it mostly because of its widespread use in the philosophical literature.
4. Norms of the form 'If you believe Y, then you should (or may) believe X too', or 'In situation S, do X' are not, on this understanding, hypothetical imperatives. Many epistemic norms might be seen as of this kind, but since they merely refer, directly or obliquely, to prior *states of belief*, not ulterior *goals or values*, their normative force is not, as in the case of genuine hypothetical imperatives, conditional or derivative.
 5. Though not really an English expression, because of the frequency with which I will use it in the following, I will henceforth write '*a priori*' in one word and in non-italics form.
 6. For example, Putnam (1983), Siegel (1990), Bonjour (1998), Janvid (2001). I am not concerned to defend in detail these thinkers' classification as traditionalists in relation to the characterisation I will be offering.
 7. Even though it is no doubt true that many actual empiricist philosophers were less sanguine about the idea of the *a priori*, and at least tended towards naturalism insofar as derivation of first principles was also meant to proceed empirically.
 8. For the idea that we might have such innate knowledge, cf. Speke (1995).
 9. Given this strong notion of the *a priori*, it is also natural to hold that the truths thus known must be necessary. In fact, I think the traditional idea of necessity meshes naturally with the ideas behind TE, insofar as the validity of its norms should not be beholden to the way the world turns out, though scientific investigation, to be. However, cashing out what this means will arguably only return one to the idea of *a priori*, whilst if one stresses a more 'metaphysical' notion of necessity, of the kind made popular by Kripke (1980), the connection will lapse because necessity of this kind is not meant to be limited to truths knowable only *a priori*. Finally, I should point out that I am open to the possibility of weaker notions of the *a priori*, at least in principle, such as those which would allow *a priori* knowledge of truths that can also be known *a posteriori*, such as 'I exist' (cf. Kitcher, 1980, for development of such an account of the *a priori* that is also meant to be naturalistic). These are not, however, notions that are relevant to TE; whether they are compatible with naturalism is not something I will investigate (a very weak notion of the *a priori* as simply *innate*, of course, is).
 10. Again, the traditional idea of necessity is close to hand; insofar as this is essentially the strong notion of the *a priori* outlined above (see previous note), the foundational aspect of TE norms is plausibly coeval with their *a priori*ity.
 11. A further, different sense in which one might say *a priori* norms could be fallible is the possibility of coming, through non-empirical enquiry, to revise our opinions about what they are (as some might hold has been the case with logic, *viz.* the debate between classicism and intuitionism). Again this idea might seem to be in some tension with the strongly

- a priori* nature of TE, but does not, as far as I can, seem to be ruled out in principle.
12. Cf. Brown (1988a: 5 ff.).
 13. I will not defend this claim further here, and hope there are not many readers who feel I need to. How relativism relates to the traditionalist/naturalist divide is accordingly not an issue I will explore; apologists for relativism will have to read the rest of the book with this caveat in mind. (Relativism is, however, briefly discussed again towards the end of Section 4.4 in relation to Laudan's pluralistic conception of epistemic goals.)
 14. This in contrast to Brown (1988a: ch. 1), who sees algorithmic rules as central to what he calls 'the classical conception of reality', which is otherwise very close to my TE (in respect of matters epistemic).
 15. See note 6.
 16. I say 'most naturalists' because there are those who describe themselves as naturalised epistemologists who nevertheless see a necessary role for the distinctively philosophical kind of intuition that TE sees as the sole basis for epistemic norms, even though they would also stress the need to take account of empirical data (and, connectedly, concur in some of the criticism of TE I will be detailing). I shall have more to say about such 'moderate' naturalists in Section 1.4.
 17. The idea behind the analytic account of the *a priori* is that *a priori* truths are secured in virtue of the *meanings* of the terms they involve. Thus, to take the classical example from everyday life, 'All bachelors are men' is seen as *a priori* in virtue of the meaning of 'bachelor'. For the logical positivists, following the lead they found in Hume, analytic *a priori* truths concerned essentially language. Since they also held that what meanings we assign terms is ultimately something we are free to stipulate, this led them to the position known as *conventionalism*, whereby *a priori* truths assume the status of conventions, adopted for some or other merely pragmatic aim. The problem with this is that it threatens to undermine the cognitive credentials of the *a priori*, and thus of the cognitive enquiry it is meant to provide a foundation for – a problem the positivists were acutely aware of, and which led to accusations of relativism. As Brown (1988: 56 ff.) relates, one response to the problem considered by Carnap was to allow that analytic truths may evolve in relation to the extent to which they are empirically or, in some cognitively relevant sense, pragmatically fruitful. However, to embrace this line consonant with denying relativism is to buy into a form of naturalism.
- Recently, Paul Boghossian (1996) has resuscitated the analytical account of the *a priori*, seeking to view logical truths as *a priori* in virtue of their expressing implicit definitions of the terms they involve, but at the same time renouncing conventionalism about meaning. The suggestion has, however, met with many critiques, of which Horwich's (1998: ch. 6) is regarded as one of the most devastating; he argues that Boghossian's account fails to secure a notion of *a priori* justification beyond something conventionalist in nature. This is too large a debate to

rehash here, but it at least seems reasonable to hold that Boghossian's line must at some point break out of the realm of the purely apriori to justify the account of our grasp of meanings and how this gives us access to logical truths – thus rendering his views irrelevant to our discussion of TE. Moreover, in concerning first and foremost logic, the account fails to address the wider ambit of epistemic norms which traditional epistemologists have been concerned to derive and defend apriori (see below).

18. Someone might object that it is enough if the norm seems pretty convincing, rather than be absolutely indubitable (note this is not the same as being fallible in the sense discussed in the previous section). On the basis of this starting point, one might then seek to 'test' the norm against concrete cases of reasoning, adjusting intuitions about the rule and concrete cases mutually in relation to one another until one reached some stable conception of the correct rule. This is in outline the *method of reflective equilibrium* first espoused by Goodman (1965), something which I will be discussing at greater length in Section 2.2 in relation to naturalised epistemology. In the present context, the point is that I do not see such a method as available to TE in order to justify its norms. The reason is that such a method, to be viable, must assume that our ordinary reasoning practices are to some significant extent already in order – in some sense given – whereas the guiding idea behind TE is that apriori norms are needed precisely to tell us which pieces of reasoning are in order and which are not. There is a necessary circularity in the method of reflective equilibrium which is anathema to TE, even though the former does not bring in substantive world-knowledge.
19. Two relatively obvious problems are that it presupposes a truth-table definition of the logical constants, plus an arguably arbitrary and dubious relevant constraint on what one can be said capable of believing. These objections, however, enunciate large issues in themselves which I do not have space to explore here.
20. An important exception to this will be when discussing *moderate* or *weak* naturalist proposals that I would want to be touched by my arguments, but that might seek to wriggle free from them by bringing in a dose of the apriori (these might also reject the paranthetical qualification of the first reason for not finding my project less interesting, above). See Section 1.4, and Section 5.1.
21. Cf. Section 4.1 for more on the justification of inductive rules.
22. I use these latter two prescriptions only as examples of norms that could plausibly be viewed as suggestions within a TE-framework; I do not mean to imply that their actual progenitors subscribe to TE, which is debatable in view of their conventionalist leanings (cf. note 17).
23. Cf. Brown (1988b: 55 f.).
24. I also see the metaphysical views of Kant – when all's said and done – as involving the idea of something non-natural as the repository for apriori norms (even though God does not play the role of a guarantor). This will no doubt be seen by Kant scholars as an execrable over-simplification, but this is not the place to explore the vast complexities of Kant's system.

25. Note that none of this is meant to impact directly on issues about the so-called *abstract objects* – for example the idea that numbers, grammars or other things (including even norms themselves) might be seen as objects outwith the world of space and time. The metaphysical objection to TE builds on the idea that TE would want to countenance *truths* or *states of affairs* (however these are conceived) that go above and beyond the truths that constitute the natural universe.

26. Cf. Brown (1996: 20 ff.).

27. For presentations of similar lines of thought, cf. ibid.: 20; Kircher (1992: Sections 2–3). It is worth noting that in ethics it has been objected against Kant that 'ought' does not straightforwardly imply 'can', in particular, it may be that one ought to do X and ought to do Y, but cannot do both X and Y. However, this point does not really undermine Kant's claim since there seems no intuitive reason to accept that being behaved to do X and being behaved to do Y imply being behaved to do both X and Y. Moreover, even if one accepted this, there would remain a sense in which it would have to be the case that one could do X for X to be obligatory, since what one *does* in any given case is reasonably regarded as an 'atomic' action, and the Kantian principle would hold for cases where X is an atomic action.

28. For a more detailed presentation of the following point, see Chemiak (1986: ch. 4); cf. also Stein (1996: 234–235, 238–239).

29. The term derives from Lakatos (1970).

30. Thus the following do not spell out precisely how I understand naturalism, for that, see the commentary of N1 and N2 in the following text.

31. Cf. Quine (1960) for famous deployment of physicalistic ideas in philosophy of mind/language.

32. I said in note 11 that apriori norms might be fallible in the sense of being revisable. However, if this is possible, it will be in some special, restricted sense; the suggestion is not that new developments in science or empirical evidence generally could effect such a change. According to naturalism, the latter is the only kind of way in which revision takes place, and all statements are in principle susceptible to it.

33. Fodor (1983: 105). In his discussion of these ideas, Fodor distinguishes two properties of scientific belief-formation: what he calls *isotropy*, which is the idea that any part of the belief-system is in principle relevant to any other; and *Quineanism*, which is the idea that evidential constraints impact on the belief-set taken as a whole, not on individual beliefs. Though they are distinct, however, Fodor admits that it is only natural to see them as generally coinstantiated. The reason would seem to be that if any belief is in principle relevant to anything else, then it is really the whole belief-system that the belief impacts upon. In relation to my distinction between N1 and N2, one might perhaps see isotropy as reflecting the naturalistic world-view (N2) and Quineanism as expressing the nature of our cognitive practices in relation to this (N1), and the fact that they generally cooccur as expressing the dovetailing of N1 and N2 that I am trying to get across in this paragraph.

34. I do not mean to be committed to these possibilities, but record them since it has become fashionable in recent philosophy to see nothing but the vaguest kind of pragmatic constraint as applying to viable systems of beliefs (cf. the works of Richard Rorty, *op. cit.*).
- From a rather different angle, someone might point out that if witchcraft, say, can form a coherent world-view, why not apriori philosophy (*viz.* 'Are you that anti-philosophy?')? However, one must recall that apriori philosophy (traditional epistemology) sees itself as grounding, and thus in some sense serving science. It seems it has failed in this task, and perhaps must fail. It is thus in relation to its *own* goals that TE has failed: What my naturalism retains is its commitment to the centrality of scientific belief.
35. Cf. for example, Fodor (1987), Kim (1998), Papineau (1993).
36. Cf. Crane (1991), Crane and Mellor (1990).
37. Cf. Dupré (1993).
38. On this point in relation to intentional psychology, cf. Stich and Laurence (1996), Knowles (1999), (2001). In a broader context see Steward (1997).
39. Cf. Knowles (2000b: Section 3.4) for a defence.
40. Quine (1969a: 76).
41. A prejudice that is perhaps on a par with his equally dogmatic physicalism.
42. This is evinced both in the fact that it never seems to be anything other than *justification* that Quine sees as taking place between beliefs in the holistically interconnected web, and, in more recent work, in a positive conception of epistemic norms as instrumentalities (cf. 1986: 664–665, 1992: 20–22). For more on the role of the normative in Quine's thought, see Sections 2.1, 4.5 and 5.1.
43. That is, normative for belief in a non-relativistic sense.
44. By all of which expressions I thus mean a theory of how, naturalistically, norms for belief-formation are to be conceived of, derived and justified, with a view to actually specifying these norms (cf. the Introduction, note 7).
45. This divergence from Quine's original usage is quite widespread; see for example Kim (1986), Stein (1996: 15 ff.). See also below for more on normativity in relation to NE.
46. Kim (1986: 43).
47. For what I see as an explicit example of a theory which takes into account aprioristic ideas, cf. Stein (1996: ch. 7) (discussed in Section 5.2). Goldman's (1986) stress on the concept of truth as the natural goal for epistemic activity seems also aprioristically justified. There may be other examples; it is, as noted above, often unclear what specific role naturalised epistemologists envisage for philosophical insight. (A notable exception to this is Kornblith, 1993b, discussed in Chapter 3, note 24.)
48. Cf. for example, Quine (1981), Brown (1988: Section III) for fuller presentations of this kind of response to the circularity objection.
49. I remain agnostic about the antecedents here, that is whether there is such a thing as normativity 'in the world' or not. It seems likely, though perhaps not certain, that NE would founder if there were not; it is

perhaps also arguable that a fully committed naturalist must reject the idea of normativity. This, however, is not a discussion I can (or need to, given the remit of the book) go into depth about here.

2 Naturalised Epistemology

1. This famous formulation derives from another paper: Quine (1953: 41). The details of Quine's disagreement with Carnap are complex, but need not detain us here; see Hookway (1988: ch. 2) for some useful commentary.

2. Quine (1969a: 84).

3. In fact, there are several different ways in which Quine has continued to express an interest in normative matters, in spite of the central message of 'Epistemology naturalised'. Two more of these are discussed in, respectively, Sections 4.5 and 5.1.

Yet a further sense is his allusion to what he calls 'evolutionary epistemology' (1969a: 90), in particular the idea that our belief-forming practices are valid in virtue of the survival value they have conferred upon us. Organisms whose mechanisms fail to latch on to truth, as he puts it in a paper explicitly on the theme, 'have a pathetic but praiseworthy tendency to die before reproducing their kind' (1969b: 126). Thus are *our* belief-forming practices vindicated (the argument goes). Whether this kind of inference – from survivability to truth – is in fact sound will be taken up in Section 5.2, in relation to our discussion of psychological NE. Here we will also see that, even if it could be established that our reasoning mechanisms secure largely true beliefs, this does not *ipso facto* secure a place for normativity. (A related but distinct suggestion would be that we ought to follow those procedures that most reliably lead to truth [and hence, perhaps, happiness or survival]. This instrumental view of norms is something that Quine also has espoused – cf. 1986: 664–665, 1992: 20–22 – and will be discussed explicitly in Section 4.4.)

4. Psychology corresponds roughly to what another commentator on naturalised epistemology, Hilary Kornblith, calls 'the strong replacement thesis' (cf. Kornblith 1997: 3 ff.) – the idea that epistemology is exhausted by empirical psychology – and *not* to what he calls 'psychologism', which is a less radical position that sees psychological evidence only as in some way *relevant* to epistemology (see Section 5.1). The view we identify below as *psychological naturalistic epistemology* subscribes to the stronger idea, and is thus a subvariety of psychologism as we understand this. (It is important to bear in mind that one can espouse psychologism but also deny *p*, as I understand Quine as doing.)
5. This assessment does not, however, hold in relation to that later work where he espouses the idea of epistemic norms as instrumentalities (cf. Quine 1986: 664–665, 1992: 20–22). See Section 4.5.
6. Goodman (1965: 63–64).
7. Cf. Rawls (1971).
8. Cf. Rawls (1974–75).

9. An example of the kind of empirical evidence that might be relevant would be the limitations in our computational capacity discussed in Section 1.3. Thus, whilst a narrow process of MRE might endorse a norm prescribing an exhaustive consistency check of one's beliefs, taking into account this empirical limitation in us would rule such a norm out. For extended discussion of the potential impact of psychological limitations and theories from cognitive science generally on normative proposals, cf. Chermiak (1986); Goldman (1986, 1992).

Sometimes wide reflective equilibrium is presented as a process that will include other kinds of empirical factor, in particular considered intuitions, plus scientific data from disciplines other than psychology. For what I have to say about the former, see note 18, below. As for the latter, I choose here to view such data as not relevant to a process of reflective equilibrium. It is in fact unclear to me how such data is meant to contribute to a theory of norms based on how we reason (i.e. on our basic cognitive practices); moreover, to do so here would only needlessly complicate the presentation of my arguments, not affect their substance. For those who would want to see how essentially the same arguments can be developed assuming scientific data of all kinds are allowed to enter into the process of wide reflective equilibrium, see Knowles (2002a). I will also indicate through notes in the present work the points at which viewing wide reflective equilibrium in this way might affect the form of my arguments.

10. This kind of objection is sometimes known as the 'garbage-in, garbage-out objection'. Various versions of it have been put forward, notably by Stich (1991: Chapter 4). My point here is not to rebut his objections to specific versions of MRE (e.g. those building on considered intuitions, on expert intuitions etc., some of which we shall consider in Chapter 3), but to the general drift of all of these objections: that MRE is not guaranteed to produce correct norms because its starting points might be, as Stich puts it, 'quite daffy'. If MRE is, by contrast, seen in my naturalistic perspective, as an ongoing process of theory-construction, any general complaint of this kind necessarily falls.

11. Here we are relying on our assumption, made above, that MRE can be used to derive and justify epistemic norms generally, not just deductive and (in a narrow sense) inductive norms.

12. On statistical reasoning, cf. Tversky and Kahneman (1982); on logical reasoning, cf. Wason (1969).

13. Cf. Cohen (1981).

14. Cf. e.g. Chomsky (1965: ch. 1), (1986: ch. 1 and 2), (2000).

15. For fuller expositions of the field, cf. Hirschfeld and Gelman (1994); Pinker (1997).

16. Cf. Stein (*op. cit.*: ch. 2) for arguments to this effect. Clearly, P presupposes the existence of a natural reasoning competence in some sense, though it can presumably remain relatively agnostic about what exactly such a competence is meant to involve. Thus, one central issue is whether the states of cognitive modules are propositional attitudes or instead some kind of *subdoxastic* intentional state; another is how they relate to

other 'modules' of the mind (cf. Knowles 2000c for discussion of these and other issues in relation to the grammar module). The viability of P would not, however, seem to hinge on the outcome of such debates. Indeed, it would not even seem necessary to P that 'the reasoning competence' be a single, unitary module – or consist in intentional structures at all, for that matter. For example, perhaps reasoning competence could be viewed as encompassing general structural features of our thinking, such as the tendency to conceptualise the world with the help of natural kind concepts (cf. Markman, 1989; Kornblith, 1993a). In the following I shall just talk of our 'reasoning competence', assuming something like the traditional Chomskyan picture in doing so, but presupposing the qualifications made in this note.

17. Note that this does not have any direct implications about the *ontology* of norms insofar as one might hold that norms *per se* are, for example, abstract objects, *a fortiori* not psychological objects. Nevertheless, a theory of what the correct norms are along the lines of P would be one according to which any interest we might have in identifying abstract norms would be a function of the fact that their structure corresponds to that of some psychological capacity (cf. Higginbotham, 1983; George, 1989 for elaboration of this point in relation to grammars and grammatical competence). We may say that, according to P, epistemic norms are at least *indexed* by our reasoning competence, even if not literally contained within them (the locution 'indexed by' is due to Stein *op. cit.*: 75). (See also the introduction to Chapter 5.)

18. This objection to Cohen is made by Stein (*op. cit.*: Section 4.3).

Another objection made by Stein against Cohen is that MRE must use *considered* intuitions and judgements about concrete cases to be plausible as a justification of norms, whereas such considered intuitions are at best irrelevant and at worst misleading in empirical psychology (*op. cit.*: Section 4.2). However, Stein also suggests, somewhat confusingly, that considered intuitions are also used in linguistics (*ibid.*: p. 146 ff.). I choose here to pass over the complexities of sorting out the impact this might have on my argument, since it is in any case rather unclear what is meant by a 'considered intuition', and how it differs in principle from ordinary intuitions or judgements, which presumably themselves are not just random guesses, but involve a degree of thought.

19. It should be noted that what I present here as P is in fact only one of the two versions of psychological naturalised epistemology that we shall ultimately be considering; it corresponds to what I will term *descriptivist* psychological NE. The other, which is motivated differently, I shall term *normativist* psychological NE. I will use 'P' to refer to the former, and 'NP' to refer to the latter; see Chapter 5 for details.

20. Note, however, that AF can be – and, for it main proponents, is – motivated in other ways than *via* MRE; see below and Chapter 4.

21. Cf. Chalmers (1990: Section 2.3); also Kosso (1991: 351 ff.). The idea behind the strategy is also referred to by Laudan as the *metamethodology thesis* – see Section 4.4.

22. Cf. his (1974), (1978).
23. The parenthetical remark alludes to the problems surrounding in particular Lakatos' proposals for a scientific methodology, which are inherently vague concerning how *long* a research programme gets to establish strong heuristic guidance and predict novel facts (cf. Feysabend 1981: 202–230). However, as many have pointed out (e.g. Couvallis 1996: 74 ff.), the lack of precision in Lakatos' prescriptions do not seem to rule them out as having normative force entirely. In this book, it should also be noted that we are concerned, first and foremost, not with the viability of particular norm-proposals, but with the more *meta-epistemological* questions of what a theory of norms looks like, how it should be justified, whether a justification of that kind is cogent and so on.
24. See for example, Kosso (1991) for a similar taxonomy motivated through a different analysis.

3 Anti-psychologism

1. Some might think that a form of AP could be motivated through something other than MRE. I consider an alternative form of AP briefly at the end of this chapter.
2. Cf. Chapter 2, note 16.
3. I should also point out that, except in the sense that one can 'learn' to become a scientist (see Sections 3.3 and 3.4), I do not actually think the replies in defence of AP here are valid (i.e. that one can learn to become more rational, in the sense of an ability). (Some relevant discussion of this point is given in Section 4.3.)
4. This line of thought has also been used against a psychological interpretation of the object of grammatical theory and of other theories putatively aimed at outlining the structure of cognitive competences: cf. Quine (1972), Stich (1972), Stich and Ravenscroft (1996). I do not subscribe to these views, but that is not relevant to the present point.
5. For more systematic treatment of the idea of variable methods, which also provide other references, cf. Chalmers (1990: ch. 2), (1999: ch. 13), Laudan (1977), (1984); and for sociological studies documenting a cavalier attitude amongst scientists to principles cherished by philosophers, cf. Collins and Pinch (1993).
6. Cf. Laudan (1984: ch. 3), (1996: ch. 8); Freedman (1999). Whether Laudan would subscribe to universal norms for science on the basis of this commitment is, however, unclear; see Sections 4.4 and 4.5 for fuller discussion of Laudan's views.
7. Chalmers (1999: 171).
8. Cf. his (1970: 185 ff.), (1977: 320 ff.). Note this goes beyond the idea that such rules fail to constitute algorithms, which is how Kuhn – misleadingly to my mind – sometimes expresses his view.
9. I thus do not mean to be presupposing a cognitivist construal of grammar as something actually known or internally represented.

10. In this connection, see the discussion of *philistines* towards the end of Section 4.5.
11. This is of course an oversimplified example, but I believe it serves to get across the essential point. Details of our perceptual and memory systems, for example, would seem also to have relevance to normative matters (cf. Cherniak 1986; Goldman 1986, 1992), but the way in which they do this appears to be, as with the computational limitations, primarily a matter of curtailing otherwise attractive epistemic norms. Note also that I am not going to argue that knowledge of our psychology should not have normative implications, only that it does not make a theory based on wide MRE any more likely to deliver norms that need to be taken account of (cf. Section 5.1 for further discussion of these last points).
12. Some might feel that the impact of knowledge of things like our computational limitations is more naturally viewed as applying to *instrumental norms* for achieving certain epistemic goals, norms which would not be derived through a process of MRE (as I understand this – see note 11 from Chapter 2, also note 23 below and note 15 in Chapter 5). For discussion of this suggestion, see Section 5.1 (instrumental norms as such are discussed in Section 3.5 below, and more extensively in Sections 4.4 and 4.5).
13. Cf. Stich (*op. cit.*: ch. 4) and Stein (*op. cit.*: Sections 3.3 and 3.5) for more on the general idea.
14. For different versions of the circularity objection, cf. Stich (*op. cit.*: 186); Stich and Nisbett (1980); Stein (*op. cit.*: 148 f.).
15. Indeed, they would then not even be derived using a process of MRE, as I understand this. This point is, however, of lesser significance: the idea that some specialised group has something to teach lay people by way of imparting knowledge of norms does not have anything specially to do with MRE, or with AP for that matter, and will be a recurrent theme in the following chapter as well.
16. Note, then, that in rejecting the idea that there are specifically *philosophical* reasoning experts, I am not thereby rejecting the idea that there might be mathematical or logical experts in some sense, and whose knowledge could enhance ordinary people's reasoning abilities. I am merely stipulating (a stipulation justified by the assumption of naturalism) that such experts, if they exist, be viewed as scientists, and that any norms they countenance be viewed as implicit in their scientific ability, or perhaps as based on investigation into the natural world. What I will argue, in this and the following chapter, is that such norms, to the extent they exist, are not necessary for optimal rational belief-formation (perhaps because they merely correspond to factual statements that could just as well be communicated to the less enlightened – see Chapter 4, note 80).
17. Cf. Kuhn (1970), Polanyi (1967).
18. Though I seek to uphold the following Kuhnian account of scientific competence as part of the defence of my overarching thesis, I should point out that, even without it, my position would not simply collapse

(all I really need is the idea that science does not have a monolithic foundation of rules, which is quite widely acknowledged independently of allegiance to Kuhn). In the following, I will therefore try to indicate to what extent objections to Kuhn would also affect the rest of what I argue for.

19. Cf. Kuhn (1970: 192)
 20. Cf. e.g. Chomsky (1980: 70).
 21. Cf. Kuhn (1977).
 22. For why the strict reading is untenable, cf. Laudan (1984: ch. 4).
 23. In some sense, they might be seen as based on a *wide* form of MRE that takes into empirical information of all kinds, not just from psychology; however, as noted in Chapter 2, note 11, it strikes me this is not a terribly helpful way of understanding the idea of wide reflective equilibrium. In any case, the point is that the empirical information about connections between means and ends would directly inform the proposed instrumental norms, not merely be fed into some further balancing process along with our ordinary intuitions (as our cognitive limitations are meant to be – see Section 3.3).
 24. Some might feel that AP could be defended by giving a naturalistic account of a common human goal or set of goals that was not based on anything like MRE and thus would be not susceptible to my criticisms of an account based on such. I am thinking perhaps first and foremost of someone who saw, say, happiness, or reproduction and survival, or satisfaction of our desires as the ultimate human goal or goals – as a simple matter of naturalistic fact. However, to this suggestion there are at least three good objections. Firstly, it is highly unclear that happiness (etc.) can count as *epistemic* goals, even if they are, in some broader sense, rational ones (note that saying this does not imply that they cannot function as part of an *explanation* as to why putative epistemic norms can be normative for belief, as both EE [a version of AF] – see Section 4.2 – and perhaps P – see Chapter 5 – would claim is the case). Secondly, it is not clear that happiness (etc.) just obviously, as a matter of naturalistic fact, *are* ultimate human goals – some further justification for this seems to be required. Thirdly, even if we accept that this simply is a matter of naturalistic fact, it is highly unclear that we would need to specify the goals as explicit norms, for surely the claim would precisely be based on the fact that humans naturally – i.e. in any case – strive towards these things.
- Hilary Kornblith (1993b) has argued that we can explain naturalistically why we (should) all value truth: since satisfaction of desire depends on true belief, then, given that we desire anything at all, we must value truth. Does this, just in and of itself, succeed in showing that truth is or should be an epistemic goal for all human beings? One might object that desire satisfaction, in some sufficiently narrow sense of desire, does not depend on true belief; we will consider this claim in relation to the idea that our beliefs are the product of a naturally evolved cognitive system in Section 5.2. However, even if one accepts that true belief is, in general, necessary for desire satisfaction, it does

not follow that truth *thereby* emerges as a universal goal. It seems reasonable to suppose that, on an ordinary understanding of desire (that on which we want a nice house, nice spouse, nice kids, nice car, etc.), what we desire includes a desire for *knowledge* of whether what we desire is obtained. If it is important for me that *p* be true, I will, in general, also want to know I've succeeded; that I have in fact achieved what I set out to. Otherwise, it surely would not matter to me whether in fact I was, say, an envatted brain.) In other words, my desires generally include a desire to know the truth about certain matters of fact. But then clearly Kornblith has not explained the norm of truth by reference to desire, but simply presupposed it.

A second – and perhaps more obvious – problem with Kornblith's account in relation to establishing truth as a universal aim is that according to this account, truth is not inherently valuable, so even if true beliefs are conducive to satisfying desires and happiness, truth is not normative as such. The norm he suggests is in effect *If you want to satisfy your desires, seek true beliefs*. I have suggested that this conceals a tautology, but even if it did not, the idea that satisfying desires is a naturalistic epistemic goal that needs to be specified as such is, as already argued, implausible.

Note finally that it is also possible to hold to an instrumental construal of epistemic norms whilst seeing the aims as establishable *a priori*, on which see Section 5.1.

4 Anti-foundationalism

1. Cf. for example, Kim (1988), Siegel (1996b).
2. A small exception to this claim will be taken up in Section 4.3.
3. Having said that, I leave open that particular varieties of AF might collapse into a form of relativism – cf. Section 4.4 on Laudan's normative naturalism.
4. Alternatively, one might hold that these are insufficient as a basis for deriving *all* the epistemic norms; for more on this idea see Section 5.1 on 'mixed' positions.
5. The idea of progress in science is traditionally understood as movement towards a specified goal, such as the whole truth about the natural world, or complete predictive adequacy. Others hold that we do not need to be that specific to make sense of scientific progress, for example Kuhn and Laudan. Laudan's views on this matter will be discussed and criticised below; apart from what I say about him, I will not go into the debate between the different camps on this issue (indeed, I am somewhat sceptical as to whether there is really a substantive debate to be pursued). Instead, I will assume that progress in some very broad, unspecified sense is something which a naturalist (in my sense) can help herself to (at least, without it, I see no way of making sense of AF-norms as genuinely normative).

6. Cited by Skorupski (1995: 567). (I have been unable to find this passage in Mill's original texts, and so defer to Skorupski's scholarship here.)
7. Mill (1895: 209).
8. Couvalis (1997: 77–78). The example is based on one that Mill himself uses (*loc. cit.*); Couvalis also acknowledges Skorupski's (1990) study of Mill as the basis of his own summary. For a similar presentation, and endorsement of a Millian approach to epistemic norms, see Papineau (1993: ch. 5). Whilst on the subject of similar views, it is also arguable that Nelson Goodman's (1965) idea of *entrenchment of predicates* as the means whereby certain inductions become acceptable and others not can be seen as basically the same kind of account of Mill's (how this should be seen as relating to Goodman's espousal of MRE; as the means for gleaming norms is not something I will explore here; perhaps his overall view can be seen as a 'mixed' variety of NE, cf. Section 5.1). Finally, I take it that the kind of justification Mill proposes could be applied to inductive canons in a wide sense, including those sometimes called *abductive*. (I do not suppose the only plausible naturalistic way of justifying abductive or inductive norms is the way Mill envisages.)
9. In fact, Mill also held that inductive inference is more fundamental than, and indeed the rational basis of, deductive inference, but one need not follow him in this to appreciate the brand of AF the views of his we are concerned with suggest. Thus one could see logic as justified through something like MRE, or perhaps even *a priori* – along with, possibly, the fundamental principle of reliance on observation. See further Section 5.1 on 'mixed' versions of NE.
10. Note that though picture may seem to assume a theory-independent observational basis for deriving hypotheses and norms, there is (as far as I can see) nothing in the Millian suggestion that requires this; in accord with our holistic view of the web of belief, we should not assume it.
11. It is perhaps worth noting that Couvalis calls this conception a *coherentist* theory of justification of norms (*ibid.*: 79). For us, the idea of 'coherentism' would be confusingly unspecific in the present context, since it has been used to describe both Mill's approach and also those which stress the use of MRE to derive epistemic norms (cf. Chapters 2 and 3). Thus it is perhaps useful to underline the central differences between these two approaches. To start with, for Mill, it is the same observations used to form hypotheses that constitute the data for choosing the norms we also use to form and evaluate the (first-order) hypotheses; whereas with MRE the data are concrete judgements of ordinary people and/or scientists. Further, as we have seen, MRE need not be used to give an anti-foundational theory of norms, in the sense of a theory that need never reach, or even have, a well-defined end-point. For discussion of a theory that does see MRE in such an anti-foundational light, see Section 4.3.
12. The expression derives from Goodman's new riddle of induction, in which the predicate 'grue' means *green before [a certain date in the future], blue thereafter*. Cf. Goodman (*op. cit.*: 80 ff.).

13. It has also been argued, by Quine (1969b) and Kornblith (1993a), that induction could only hope to work if there were natural kinds at all. However, I do not see how this idea can be justified in a strong sense – i.e. such that the assumption of natural kinds is an absolutely *essential* condition for inductive success – and at the same time naturalistically, unless it simply amounts to the assumption that our basic cognitive practices are on track. Note also that in the text we are concerned first and foremost with inductive principles being proposed on the basis of knowledge about *what* the natural kinds are.
14. A logically possible alternative is that the Millian AF-norms rest on AF-norms of a different kind, as described in subsequent sections of this chapter. Since the Millian version understands itself as the *correct* version of AF, this is not really a meaningful option, but in any case it just pushes over the question of the legitimacy and necessity of AF-norms over to these other sub-varieties.
15. This seems to be Harold Brown's (1988b) chief understanding of what naturalistic epistemic norms are (cf. pp. 70–71), though he also seems to accept norms as hypothetical imperatives in the manner of Laudan (see Sections 4.4 and 4.5); in a later paper (Brown, 1996), he also extends his conception of norms in a psychological direction (cf. Section 5.1).
16. One might wonder whether such a norm should be classed as an AP- or AF-norm. I tend towards the latter view, for presumably the advice offered will be *pro tem* given the ever-evolving nature of scientific knowledge. Some might nevertheless be inclined to see it as an AP-norm – after all, P, which, like AP, is only a weakly anti-foundational theory (see Section 2.1), uncovers norms that are presumably *pro tem* insofar as the theories of psychology are. However, the real reason why P is not strongly anti-foundational is that it involves the idea of a specifically epistemological subject-matter. Thus the kind of norm cited in the text would seem apt to be classed as (strongly) anti-foundational, for its status as such is owed only to its playing a certain role in a certain context; not its partly constituting the subject-matter of 'the theory of knowledge'. The point, however, is not of decisive importance for anything substantive I want to argue for.
17. This conception accords roughly with Popper's understanding of evolutionary epistemology (cf. Popper, 1962), as well as that of Donald Campbell (1974). Note then that I mean to exclude theories based on the idea that biological evolution has in some way or other ensured that our belief-forming practices are essentially correct (as we shall see in Chapter 5, this is an idea that has relevance to P). On the other hand, I would want to include theories that would see development in norms as based on some genuine yet – in a narrow sense of 'scientific' – non-scientific development, in the manner of what is sometimes called *developmental* or *genetic* epistemology (i.e. a theory based on the idea of learning in children). Clearly, there are many issues connected to the range of theories alluded to here, which I cannot go into. Nevertheless, the fundamental objection I have against EE (see below) would, I believe, carry-over in

- principle to an analogous theory based on the broad idea of, as it is sometimes put, conceptual evolution.
18. NN has been developed in several distinct stages over the last twenty years. The central works are Laudan (1984) and Laudan (1996), especially part IV of the latter, which collects together his essays on NN and replies to his critics.
 19. Laudan characterises NN as a *metaepistemological* theory because it is not concerned, in the first instance, to specify particular norms or values for enquiry, but rather account for how such values and norms in general are to be derived, justified and used. In this sense, the different varieties of NE we have been concerned with would also class as *metaepistemological*.
 20. Laudan (1996: ch. 7 and *passim*).
 21. Of goals or values or ends, as he alternatively refers to these things. I regard these as notational variants and will use whichever seems most natural in the context. Instrumental norms I will tend to talk of as 'rules', 'methods' or 'methodological rules'. Note that both aims and rules qualify as epistemic norms in the broad sense relevant to our conception of NE.
 22. This is a recurrent theme in Laudan (1981), (1984), (1996). It was also discussed in Section 3.2.
 23. Cf. Laudan (1984: ch. 3–4).
 24. *Ibid.*: ch. 3.
 25. *Ibid.*
 26. Cf. Laudan (1996: 158, ch. 10).
 27. *Ibid.*: 138.
 28. *Ibid.*: ch. 8; 162.
 29. The present section involves an extended discussion of Laudan's views and several interpretations of these which some readers, in particular those already convinced of the deficiency of Laudan's theory, might want to skim or skip. These readers are advised to at least read the conclusion of the section before proceeding to the discussion of instrumental norms.
 30. Cf. for example, Doppelt (1990). Harvey Siegel (1990, 1996a) has a related but distinct objection to NN, to the effect that justifying methodological norms of the kind Laudan envisages presupposes epistemic rules that are not instrumental in nature. Though the dialectic involved in answering Siegel's objection is somewhat more complicated than that involved in answering the one we are considering, Freedman's *scientific naturalism*, which I will be deploying against the latter objection (see below), seems ultimately apt to buttress Laudan against Siegel's too.
 31. Laudan (1996: 160); cf. also Laudan (1977).
 32. Siegel (1990: 311).
 33. Freedman (1999).
 34. *Ibid.*: 535.
 35. *Ibid.*: 157 (emphasis added).
 36. Cf. Freedman (*op. cit.*: 537).
 37. Laudan (1996: 128).

38. *Ibid.*: 126; cf. also Lakatos (1974, 1978), Feyerabend (1975).
 39. Laudan (1996: 129).
 40. *Ibid.*: 137.
 41. *Ibid.*: 138.
 42. *Ibid.*: 139.
 43. *Ibid.*: 162.
 44. *Ibid.*: 158.
 45. *Ibid.*: ch. 8 *passim*. Though this seems to be Laudan's considered opinion, it is somewhat at odds with his statement that '[i]f we were today espousing cognitive aims which had not been progressively realised in the development of science, we should be forced to do methodology largely [...] apriori' (*ibid.*: 139), insofar as the latter suggests we could have bona fide scientific aims that are not to be found in the 'scientific canon', as he puts it – which seems to be precisely what the statements cited in the text exclude. I will not go into this tension here, though I do take up the idea of doing methodology without reference to the history of science below.
 46. Laudan (1996: 132).
 47. It is important to bear in mind, as noted above, that I am not committed to the correctness of this conception of progress.
 48. In this connection, it is worth pointing out that a central example of a divergence in aims in later science that Laudan uses – the debate between supporters and opponents of the method of hypothesis in the eighteenth century (1984: ch. 3) – seems clearly to be a divergence of methods, not aims. For criticism of Laudan's assumption of a sharp distinction between aims and methods in science, cf. Brown (1986).
 49. Laudan (1996: 138).
 50. Freedman makes clear, as do the passages from Laudan she cites, that we are talking about *scientists'* and not just *science's* constraints – as must be the case if we are seeking to *justify* these constraints (cf. also note 52).
 51. For those who are sanguine about the prospects of 'weakly' naturalistic theories, cf. Section 5.1 on 'mixed' versions.
 52. Note that satisfying the axiological constraints (assuming these can be naturalistically justified) would not be sufficient to justify a set of values (since, as noted above, they underdetermine a choice of values; see also the discussion below).
- One might wonder whether Laudan could point to anything other than scientific naturalism/MMT to justify axiology, within the boundaries defined by naturalism generally. As noted above, any axiological theory, in addition to satisfying the anti-utopian constraints Laudan mentions, must also accord with the aims realised in important scientific work, though not necessarily the aims of great scientists. However, building an axiology in accord with the aims actually achieved by science cannot justify one's axiology: an axiology of science would not be an axiology of science unless it did the former. It is only when we relate our aims to some independent factor – according to MMT, the aims of past scientists in doing great scientific work – that anything like a

justificational structure comes into view. Nor could Laudan plausibly distance himself from MMT by seeing aims as justified only by reference to the aims of scientists today, since for a start these facts are already history, and in any case the restriction would be completely arbitrary given that aims are shared between present and past scientists.

Could Laudan appeal, in place of MMT, to something like the general pragmatic success of theories steered by certain values as justification for those values? At least one point (1996: 166–167) Laudan seems to suggest his allegiance to such a pragmatist line. This idea is reminiscent of the varieties of AF discussed in the earlier sections of this chapter, which we have already shown do not entail a need for norms. Moreover, it is doubtful that the idea could make a substantive contribution within Laudan's overall metaepistemological framework. For him the justification of an action must always be relative to it procuring some *specific* goal or other; there would seem, therefore, to be no room for a general notion of pragmatic success by reference to which one could justify values.

53. Laudan (1996: 162).

54. *Ibid.*: 168. The charge has been levelled by *inter alia* Worral (1988).

55. Stephen Stich has also defended the idea of pluralism about cognitive goals in a way that resembles Laudan's (cf. Stich 1990: ch. 6), and is, I think, susceptible to the same critique as that offered here.

56. Laudan hints at several points at a distinction between methods and how these are described (e.g. 1996: 162), a distinction that might seem to be open for distinguishing at least between how methods are conceived now and how they were conceived in the past. However, this distinction, even if made out with sufficient clarity, does not seem relevant to the point in the text: if the method gets the job done, how you think about it or talk about it is not going to change this fact.

57. Indeed, he seems to suggest we do not need to do this: cf. the citation in note 45.

58. Thanks to Øyvind Linnebo for bringing these necessary points of clarification to my attention.

59. In addition to Laudan, cf. Giere (1989), Kornblith (1993b), Kitcher (1992), Goldman (1986), Maffie (1990), Stich (1990), Henderson (1994). It is also part of Quine's more recent views on naturalised epistemology; as we shall shortly see.

60. Quine (1986: 664).

61. For more on Quine on truth, cf. his (1992).

62. Cf. Rosenberg (1990: 36). The objection is discussed by Laudan at pp. 163 ff. of his (1996).

63. *Ibid.*: 164–165. I abstract away from recent discussions of the chicken-and-egg problem from evolutionary theory suggesting the egg must have come first!

64. And which of course are not, in any case, *epistemic* norms.

65. Note that this point is essentially unaffected if one instead sees instrumental norms as of the form 'Since you want Y, do X' (or something similar). In asserting the antecedent, we can of course detach the normative

consequent. But one who is apprised of the empirical basis for the connection no more needs to have this consequence made explicit than he needs a hypothetical imperative. Since the latter is not needed to reason oneself to what is rational to do – *given* one's aims and information about the means for attaining those aims – then neither is a statement which simply in addition records explicitly the existence of the aim – and, *a fortiori*, detaches the relevant normative conclusion.

66. Previously I thought one could go even further and mount a case to the effect that hypothetical imperatives concealed, beneath their surface normative form, a semantic essence that was essentially descriptive in character – to the effect of something like *assuming you might have some interest in goal Y, I can inform you that X leads reliably (or better than anything else) to Y*. I still think the proposal is plausible, but am not as convinced of it as I used to be. Moreover, since it plays only a subsidiary, and ultimately inessential role in my argument, I have chosen to omit it here (for a full presentation, cf. Knowles 2002b).

67. In discussion, Helge Høibraathen coined the term 'science slaves' to characterise such people, a phrase which I think aptly summarises their envisaged role! In fact, I think it is unlikely that there are many such people, for most involved in scientific work would, it seems, have some idea about the aims their work is meant to further, and would be performing what they are doing at least partly for the sake of furthering those aims. (For the register: I do not mean to be implying anything of a demeaning or disparaging character about any particular professions, nor of a moral character concerning the organisation of scientific communities.)

68. Indeed, since the envisaged norms would then be purely hypothetical or instrumental in character, one could also run an argument of the kind I have offered in this section against their necessity. My concern here, however, is only with epistemic norms.

Some might feel that my division of human beings into 'scientists' and 'philistines' is too crude. For example, Philip Kitcher has argued that in a scientific community, it is optimal to endorse a variety of goals that different practitioners might pursue in the overall furthering of our knowledge of the world (cf. Kitcher 1992). Thus, for example, whilst most should perhaps concentrate on refining the existing paradigm, it is also desirable that a smaller group seeks to find problems for the paradigm of the hope of developing new ones and thus 'pushing back the frontiers of knowledge'. An individual scientist should be free to an extent to choose his or her goals, but if there begin to be too many concerned, say, with ground-breaking ideas (and, perhaps, winning Nobel prizes), rather than consolidating the knowledge-base as it stands, it could presumably be reasonable to compel them to behave otherwise. Thus it might seem as if norms would be required even though such 'people could hardly be classed as philistines. However, I do not find this 'counterexample' to my thesis convincing at all. The idea of a hierarchical structure in the aims and methods of a scientific community (something I have already discussed in Section 4.4) is perfectly consistent with the use of just

- descriptive knowledge of the relevant connections to facilitate optimal research. Thus, a scientist who was genuinely interested in the growth of science (i.e. in its most superordinate aim: truth, or whatever) would only need to be told that, if there are not enough people working on some particular sub-task, then the scientific enterprise as a whole will suffer. Exactly who gets delegated to do what is another matter, but though this process may be normative (in the sense of involving orders or imperatives from superordinates, or even practical rules), it will surely not, or at least need not, involve epistemic norms.
69. This objection was put to me by David Henderson.
70. Cf. Collins and Pinch (1993) for an array of examples.
71. Though perhaps not all, for example, internal consistency. I return to such criteria below in discussing common sense maxims.
72. Cf. Kuhn (1977: 262, 322, esp.).
73. Laudan (1996: 90 ff.) reads him this way.
74. A similar construal of the point was made in Chapter 3.
75. *Ibid.*: 92.
76. Laudan connects the constraint of anti-utopianism directly to the question of realism: for him, pursuit of the truth is a utopic, and thus unworthy, aim for science (cf. Laudan 1984: ch. 5). If one has, like me and many others, qualms about the meaningfulness of the realism/anti-realism debate, one will also, I think, be suspicious of the possible significance of norms of utopianism/anti-utopianism that are anything other than wildly implausible.
77. Something I would class as an Af-norm; see note 16.
78. Cf. Brown (1988b: 70 f.) for a very explicit statement of this idea and several more specific examples from physics (as he puts it on p. 70: '...scientists evaluate theories on the basis of powerful imperatives, but these imperatives are part of science' – to which I say 'Precisely!', and infer that we should jettison the idea that they are essentially normative statements).
79. Perhaps because it has been argued that it is *essential* for us to have any chance of understanding nature inductively (cf. Quine, 1969b and Kornblith, 1993a). For criticism of this idea, see note 13.
80. I think a similar view for principles of mathematics and logic is plausible – at least to the extent that these are not viewable as part of some kind of systematisation cognitive ability. In other words, these principles would be seen as basically factual statements, embedded, albeit deeply, in the web of belief, and justified in relation to their coherence with the rest of our science. However, this view is highly controversial; as noted in Section 1.3, there are many who would still want to defend a view of logic and mathematics as *a priori*. (As also indicated there, the arguments of this book must accordingly be understood either as concerning putative norms other than those of logic and mathematics, or as drawing out consequences of the assumption that the latter should be understood naturalistically.)
- Brown (1988b: 71 f.) also operates with norms for observation in addition to what we might call theoretical norms (of the kind we have just

- been discussing). He claims that science operates with clear norms for what can count as a legitimate observation, about how one can legitimately gather such observations and so on. However, to the extent one holds this, it is clear that these rules are based on empirical knowledge about the nature of the physical world and our capacities for detecting it, and thus again it seems we can dispense with them so long as the theoretical knowledge remains in place. Any given scientist need not know, of course, the detailed background for the norm that, say, certain kinds of telescopic data are reliable whilst other such data are not. But she must, surely, be credited with the knowledge *that* certain data are reliable whilst other data are not (and what this involves in practice), which is enough to eclipse any norm *prescribing* the use of certain kinds of data. It is also reasonable to believe that many observational norms take the form of hypothetical imperatives (see below in the text).
81. This objection was put to me by Martin Carrer in relation to the Laudan material.

5 Psychologistic Naturalised Epistemology

1. It also plausibly subscribes to what Kornblith (1997: 3 ff.) calls the 'strong replacement thesis' (see Chapter 2, note 4), which is coterminal with our idea that epistemic norms are essentially something psychological.
2. Cf. Chapter 2, note 17 for an explanation of the notion of indexing.
3. In a suitably broad sense of 'competence', this can also be said to be the position of NP. For NP, though the idea of rational competence as a biological module to be studied empirically would probably be rejected, the general idea of rational competence as something natural and contrastable with performance would seem to be acceptable. Thus most of what I have to say about reasoning competence in Section 5.2 below applies to both P and NP.
4. For relevant discussion, cf. Knowles (2000b), Stein (*op. cit.*: ch. 2).
5. This is of course pivotal for P (and NP): if one denies that there is anything worthy of description as a natural reasoning competence, then the idea that psychology should have anything to do with cognitive norms for reasoning and enquiry is surely a non-starter (see also the qualifications in my understanding of what is meant by 'reasoning competence' in Chapter 2, note 16). For some considerations generally supportive of the idea that we possess such a competence, cf. *ibid.*
6. Note that in saying this and thus suggesting that MRE may be viewed as a scientific method, I am not of course saying that we need to follow any explicit rules. Given what I have argued so far, we can view MRE simply as a process whereby we generate hypotheses about the structure of some or other ability we possess by using our reasoning ability (which for P will be based in a natural reasoning competence).
7. Antony (1987: 255, first emphasis added). In this piece, I present P as a position most heavily influenced by Chomsky's neo-rationalistic project

- in linguistics (though I do not mean to imply that it follows from accepting this project, indeed, I personally think it does not). Certain views that are more directly inspired by an evolutionary perspective on cognition, viewing our reasoning practices as *ipso facto* optimal from such a perspective, might also be seen as cleaving to P (cf. Stanovich and West, 2000; Stein *op. cit.*: Section 7.2).
8. Cf. Kornblith (1997: 10 ff.). Note, however, that Kornblith uses the expression somewhat differently from this (for something more like the first precisification given below).
 9. Cf. Kornblith (1997: 11 f.) for examples of this kind of view.
 10. In the Kornblith presentation I am referring to, examples here would be supporters of what he calls 'the weak replacement thesis' and 'psychologism'. Examples of such views relevant to my concerns will be discussed below.
 11. What Kornblith calls 'the strong replacement thesis' (cf. Chapter 2, note 4).
 12. Cf. Chemiak (1986: ch. 3) and Goldman (1986), (1992) for some subtler examples involving details about the human memory system and perception. As far as I can see, however, the principled point that emerges is the same.
 13. For discussion of the idea of considered judgements, and why I pass over the qualification, see Chapter 2, note 18.
 14. Cf. Stein (*op. cit.*: ch. 7, especially p. 256). Stein tends to talk of a naturalised theory (or picture) of *rationality*, rather than just naturalised *epistemology*, a practice I will to an extent follow in this chapter (since Stein's work will be a central theme). This should not, however, be taken as indicating a broadening or shift of subject-matter: for us the focus is still epistemology, and references to acts and actions should be understood as actions of belief-formation or other 'epistemic moves'. For a view similar to Stein's in relation to questions concerning specifically practical rationality, cf. Bermúdez (1999).
 15. For example, Goldman (1986). As a matter of fact, it is not always made terribly clear exactly what status the goals of enquiry are meant to have in discussions in the literature, for an admirable exception to this, cf. Kornblith (1993b, discussed in Chapter 3, note 24). Note that one could also see the suggestion as building on a process of wide reflective equilibrium in which one takes into account philosophical intuitions plus empirical information of all (relevant) kinds, not just from psychology. In this piece, I have chosen to understand wide MRE such that it excludes this latter kind of information, though this has no substantive ramifications on my arguments (see Chapter 2, note 11).
 16. Cf. Brown (1996).
 17. In saying this, someone might wonder why a supporter of Stein's position, discussed above, could not also use information from psychology to arrive at a norm in the straightforward way I have suggested is possible here. The reason is that Stein's picture does not operate with hypothetical norms: if one does, then one can easily (in principle) decide that a certain method will be the best or second best, or whatever, for attaining a certain goal. But if one does not – if the *prima facie* norms are categorical in nature – then finding out that following these is unfeasible for us does not allow us to suggest some alternative norm without (much) further

- ado. For example, if having consistent or true beliefs is inherently epistemically desirable, but we as humans cannot hope to attain this, it is just not clear what kind of compromise we ought to accept.
18. Or at least need not. As noted in Chapter 4, one could construe any scientific truth as a norm, and this might even seem natural in cases of well-established general truths, like the theories of relativity and quantum mechanics (cf. the discussion of QR in Section 4.6). However, as also noted there, given the background naturalistic conception of belief and rationality, this is quite unnecessary.
 19. Note also that the notion of competence is being understood in a broad sense such that the rationality debate is also of relevance to supporters of NP (who see our competence as natural but not, at least wholly, biological): cf. also chapter 2, note 16, for general qualifications concerning the notion of competence.
 20. This is also true of NP. However, unlike P, NP can be said to incorporate an explanation as to why the rationality question is foreclosed: see below.
 21. The motivation for this is that it seems that many principles that we intuitively would not regard as rational seem to be in reflective equilibrium for many people, for example those that lead to the gambler's fallacy (cf. Stich and Nisbett, 1980; Stich, 1990: Section 4.3).
 22. Cf. Stein (*op. cit.*: Ch. 7).
 23. One might be wondering how these inferences are meant to be justified, if not by reference to some independent, non-psychological set of standards – thus presupposing the falsity of P. However, as far as I can see, they might be underwritten by the principles in our reasoning competence, as P maintains. That of course returns us to the superordinate conundrum of how P could ever be tested empirically in the way being assumed here, but that, as I said above, is something we are simply demurring at for present purposes.
 24. For those interested, cf., in addition to Stein (*op. cit.*: ch. 6), Stich (1990: ch. 3).
 25. Stein considers essentially the same kinds of objection to be cogent against an evolutionary epistemological defence of the rationality thesis, that is one that does not base itself on *biological* selection but rather some kind of conceptual evolution. I see evolutionary epistemology (or at least what I call 'EE') as a brand of Af rather than P (cf. Section 4.3), but Stein's objections would seem at least equally forceful against EE as against an evolution-based form of P ('at least' because, in the evolution of scientific knowledge, one would presumably not be considering beliefs that will interact with desires – see below).
 26. One angle on this would be to reflect that if what makes our reasoning competence rational is that it is conducive to survival, why should not a whole load of other natural capacities, both psychological and otherwise (e.g. digestion), also be counted as part of our reasoning competence?
 27. For example, Fodor (1998), Stich (1990: ch. 3).
 28. This can be compared to Chomsky's argument for grammatical principles as genuinely known – to the effect that, if they were apprehended

- consciously, we would not hesitate to say they were known (cf. Chomsky, 1980: 70).
29. Cf. in particular McDowell (1994), esp. ch. IV. Since this text is richly interconnected, I do not give detailed page references to it in what follows.
 30. Translated by Ross (1954).
 31. Cf. McDowell (1980) and, especially, (1985).
 32. *Ibid.*, p. 389.
 33. The following characterisation is an extremely simplified distillation of a large body of work, the principal foundations of which are now collected together in Davidson (1980) and (1994).
 34. I am, of course, just concentrating on theoretical rationality here.
 35. Something I attempt in Knowles (2002c: part 3). This issue is not, however, directly relevant to the prospects for NE.
 36. Cf. e.g. (1985: 390 f.). The idea of the constraints being, as he puts it, 'apriori' should not suggest a return to TE and its attendant metaphysics, but rather as distinguishing the psychological as non-empirical in relation to 'bald' nature. See also following note.
 37. This argument, as noted above, is based on the idea that abiding by rational norms constitutes an apriori condition for belief-possession. Glossing 'apriori' as 'non-baldly empirical', it seems that McDowell would want to accept this argument (*viz.* norms are simply part of our psychological nature).
 38. The following is, with slight modifications, taken from Section 5.2 of Knowles (2002c). I do not assume any sharp divide between scientific and other kinds of explanation here, but that does not affect my essential point, which is that the idea of specifically normative explanations make no sense. Thus I am not insisting that all real explanations must be deductive-nomological or cleave to some otherwise monolithic model; perhaps scientific ones do, though I also doubt that. (Thanks to Petri Ylikoski for getting me straight on this point.)
 39. Note that I do not mean to suggest, in presenting the putative explanations in this section in this way, that explanations necessarily take the form of arguments. The format used is merely a convenient way of trying to make clear McDowell's intentions.
 40. This was Hempel's (1965) objection to the idea of purely 'rational' explanations. Ostensibly normative explanations aside, both Hempel and McDowell assume that subsumptive explanations exhaust the realm of genuine explanations, an assumption I do not make use of in my arguments.
 41. Note also then that it should not be understood as begging any fundamental questions about the fact-value distinction: it merely records an intuitive reaction to the idea that a statement about what someone should do – in and of itself – can be relevant to explaining what they actually did. McDowell seems to think it can. (See also below for further discussion of this point.)
 42. Cf. McDowell (1980: 126).

43. Or perhaps a rule that covers inference patterns of this form (i.e. *modus ponens*). Note that the assumption that explanation (2) is a good one does not imply that X needs to believe what it is stated she believes in the third 'premise' in precisely that way, or as a rule. One way for it to be true, for example, would be for X to be consciously aware of the inference from A and *if* A, B to B being, as Peacocke (1991) puts it, *primitively compelling*, that is, something she simply feels is rationally mandated without further justification and which she is overwhelmingly inclined to make. The clause is meant simply to get across the idea that X *herself* sees what she comes to do as warranted, given her prior beliefs. This also make clear that countenancing explanation (2) in no way compromises anything I have said concerning the non-necessity of norms for rational belief-formation. There is no assumption that X needed to have any conscious awareness of anything explicit, *a fortiori* of a norm, in order to draw the conclusion in the way (2) gets across she did.
44. Essentially this line of thinking is also to be found in Henderson (1993) (cf. especially chs. 5 and 6); I take my point to extract the intuitive essence of Henderson's case.

Conclusion: Science without Norms

1. Or however one seeks to characterise the ultimate goal of science. I have perhaps, more often than not, assumed what would be called a *realistic* conception of science in this piece (in a broad sense) but in deference to anti-realists such as Laudan, I record the possibility of an alternative. Nothing of what I have said, as far as I can see, turns on this debate (if it is genuine).
2. Supporters of P might see their work as part of cognitive science. But since P is not tenable, cognitive science is not relevant to a theory of norms (something I take it most practising cognitive scientists will not be too concerned about).
3. Carroll (1895).
4. Perhaps this point has connections to the problem of normativity (this is plausible in the light of the fact that both can be seen as featuring in Hume's discussions of ethics and the precursors of rational action; cf. Blackburn 1995). Insofar, we can merely point out the problem of normativity is, at a completely general level, something which naturalism justifiably demurs at (cf. Section 1.5).
5. Cf. Wittgenstein (1953) and (1956).
6. Kripke (1982).
7. Wittgenstein (1953: Section 201).

Bibliography

- Abbreviations: OUP = Oxford University Press; CUP = Cambridge University Press; HUP = Harvard University Press.
- Antony, L. (1987) 'Naturalised epistemology and the study of language', in A. Shimony and D. Nails, eds, *Naturalistic Epistemology*. Reidel.
- Bermúdez, J. (1999) 'Psychologism and psychology', *Inquiry*, 42, 487-504.
- Blackburn, S. (1995) 'Practical tortoise raising', *Mind*, 104, 695-711.
- Boghossian, P. (1996) 'Analyticity', in C. Wright and B. Hale, eds, *A Companion to the Philosophy of Language*. Blackwell.
- Bonjour, L. (1998) *In Defense of Pure Reason*. CUP.
- Brown, H. (1986) Review of Laudan (1984). *Philosophical Review*, 95, 439-441.
- (1988a) *Rationality*. Routledge.
- (1988b) 'Normative epistemology and naturalised epistemology', *Inquiry*, 31, 53-78.
- (1996) 'Psychology, naturalised epistemology and rationality', in W. O'Donohue and R. Kitchener, eds, *The Philosophy of Psychology*. Sage.
- Campbell, D. (1974) 'Evolutionary epistemology', in P. Schlipp, ed., *The Philosophy of Karl Popper*. Open Court.
- Carroll, L. (1895) 'What the Tortoise said to Achilles', *Mind*, 4, 1895, 278-280.
- Chalmers, A. (1990) *Science and its Fabrication*. Open University Press.
- (1999) *What is this thing called science?* 3rd edition. Open University Press.
- Chemlak, C. (1986) *Minimal Rationality*. MIT Press.
- Chomsky, N. (1965) *Aspects of the Theory of Syntax*. MIT Press.
- (1980) *Rules and Representations*. Blackwell.
- (1986) *Knowledge of Language*. Praeger.
- (2000) *New Horizons in the Study of Language and Mind*. CUP.
- Cohen, J. (1981) 'Can human irrationality be experimentally demonstrated?', *Behavioural and Brain Sciences*, 4, 317-331.
- Collins, H. and Pinch, T. (1993) *The Golem: What Everyone Should Know About Science*. CUP.
- Couvalis, G. (1997) *Philosophy of Science: Science and Objectivity*. Sage.
- Crane, T. (1991) 'Why indeed? Papineau on supervenience and the completeness of physics', *Analysis*, 51, 32-37.
- and Mellor, H. (1990) 'There is no question of physicalism', *Mind*, 99, 185-206.
- Davidson, D. (1980) *Essays on Actions and Events*. OUP.
- (1994) *Inquiries into Truth and Interpretation*. OUP.
- Doppelt, G. (1990) 'The naturalist conception of methodological standards', *Philosophy of Science*, 57, 1-19.
- Dupré, J. (1993) *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. HUP.
- Feyerabend, P. (1975) *Against Method*. New Left Books.
- Fodor, J. (1983) *The Modularity of Mind*. MIT Press.
- (1987) *Psychosemantics*. MIT Press.
- (1998) *In Critical Condition*. MIT Press.
- Freedman, K. (1999) 'Laudan's naturalistic axiology', *Philosophy of Science*, 66 (Proceedings), S526-S537.
- George, A. (1989) 'How not to become confused about linguistics', in A. George, ed., *Reflections on Chomsky*. Blackwell.
- Giere, R. (1989) 'Scientific rationality as instrumental rationality', *Studies in the History and Philosophy of Science*, 20, 377-384.
- Goldman, A. (1986) *Epistemology and Cognition*. HUP.
- (1992) 'Epistemic folkways and scientific epistemology', reprinted in H. Kornblith, ed., *Naturalising Epistemology*, 2nd edition. MIT Press, 1997.
- Goodman, N. (1965) *Fact, Fiction and Forecast*, 2nd edition. Bobbs-Merrill.
- Hempel, C. (1965) 'Reasons and covering laws in historical explanation', in S. Hook, ed., *Philosophy and History*. New York University Press.
- Henderson, D. (1993) *Interpretation and Explanation in the Human Sciences*. State University of New York Press.
- (1994) 'Epistemic competence', *Philosophical Papers*, 23, 139-167.
- Higginbotham, J. (1983) 'Is grammar psychological?', in L. Cauman et al., eds, *How Many Questions? Essays in Honour of Sydney Morgenbesser*. Hackett.
- Hirschfeld, L. and Gelman, S., eds (1994) *Mapping the Mind: Domain Specificity in Cognition and Culture*. CUP.
- Hookway, C. (1988) *Quine*. Polity.
- Hookway, P. (1990) *Truth*. Blackwell.
- (1998) *Meaning*. OUP.
- (1998) 'Naturalism and the Status of Epistemology'. Doctoral dissertation, Stockholm University (published by Akademityck AB).
- Kim, J. (1988) 'What is "naturalised epistemology"?' in J. Tomberlin, ed., *Philosophical Perspectives*, 2: Epistemology. Ridgeview.
- (1998) *The Philosophy of Mind*. Westview.
- Kitcher, P. (1980) 'A priori knowledge', *Philosophical Review*, 86, 3-23.
- (1992) 'The naturalists return', *Philosophical Review*, 101, 53-114.
- Knowles (1999) 'Physicalism, teleology and the miraculous coincidence problem', *The Philosophical Quarterly*, 49, 164-181.
- (2000a) 'Can we give a foundationalist, non-psychologistic naturalised epistemology?', in B. Brogaard, ed., *Proceedings of the 23rd International Wittgenstein Symposium*, Kirchberg, 232-238.
- (2000b) *Vitenskap, Rasjonalitet og Common Sense*. Cappelen Akademisk Forlag.
- (2000c) 'Knowledge of grammar as a propositional attitude', *Philosophical Psychology*, 13, 325-353.
- (2001) 'Does intentional psychology need vindicating by cognitive science?' *Minds and Machines*, 11, 347-377.
- (2002a) 'Do we need norms in a naturalised epistemology?' *Croatian Journal of Philosophy*, 2, 283-297.

- (2002b) 'What's really wrong with Laudan's normative naturalism', *International Studies in Philosophy of Science*, 16, 171–186.
- (2002c) 'Is folk psychology different?', *Erkenntnis*, 57, 199–230.
- Kornblith, H. (1993a) *Inductive Inference and its Natural Ground*. MIT Press.
- (1993b) 'Epistemic normativity', *Synthese*, 94, 357–376.
- (1997) 'Introduction: What is naturalistic epistemology?', in H. Kornblith, ed., *Naturalising Epistemology*. 2nd edition. MIT Press.
- Kosso, P. (1991) 'Empirical epistemology and philosophy of science', *Meta-philosophy*, 22, 349–363.
- Kripke (1980) *Naming and Necessity*. Blackwell.
- (1982) *Wittgenstein on Rules and Private Language*. Blackwell.
- Kuhn, T. (1970) *The Structure of Scientific Revolutions*, 2nd edition. Chicago University Press.
- (1977) 'Objectivity, value-judgement and theory-choice', in *The Essential Tension: Selected Studies in Scientific Tradition and Change*. Chicago University Press.
- Lakatos, I. (1974) 'Falsificationism and the methodology of scientific research programmes', in I. Lakatos and A. Musgrave, eds, *Criticism and the Growth of Knowledge*. CUP.
- (1978) 'History of science and its rational reconstruction', in J. Worral and G. Currie, eds, *Imre Lakatos. Philosophical Papers, Volume 2: Mathematics, Science and Epistemology*. CUP.
- Laudan, L. (1977) *Progress and its Problems*. Routledge.
- (1981) *Science and Hypothesis*. Reidel.
- (1984) *Science and Values*. University of California Press.
- (1996) *Beyond Positivism and Relativism*. Westview.
- Matfite, J. (1990) 'Recent work on naturalised epistemology', *American Philosophical Quarterly*, 27, 281–293.
- Markman, E. (1989) *Categorization and Naming in Children*. MIT Press.
- McDowell, J. (1980) 'Physicalism and primitive denotation: Field on Tarski', in M. Platts, ed., *Reference, Truth and Reality*. Routledge.
- (1985) 'Functionalism and anomalous monism', in B. McLaughlin and E. Lepore, eds, *Actions and Events*. Blackwell.
- (1994) *Mind and World*. HUP.
- Mill, J.S. (1895) *A System of Logic*, 8th edition. Longman.
- Pachter, E. (2002) 'Naturalistic epistemologies and normativity', *Croatian Journal of Philosophy*, 2, 297–317.
- Papineau, D. (1993) *Philosophical Naturalism*. Blackwell.
- Peacocke, C. (1991) *A Theory of Concepts*. MIT Press.
- Pinker, S. (1997) *How the Mind Works*. Penguin.
- Polanyi, M. (1967) *The Tacit Dimension*. Routledge.
- Popper, K. (1962) *Conjectures and Refutations*. Basic Books.
- Putnam, H. (1983) 'Why reason can't be naturalised', in *Realism and Reason*, Philosophical Papers, Volume 3. CUP.
- Quine, W. (1953) 'Two dogmas of empiricism', in *From a Logical Point of View*. HUP.

- (1969a) 'Epistemology naturalised', in *Ontological Relativity and Other Essays*. HUP.
- (1969b) 'Natural kinds', in his *Ontological Relativity and Other Essays*. HUP.
- (1972) 'Methodological reflections on current linguistic theory', in D. Davidson and G. Harman, eds, *Semantics of Natural Language*. Reidel.
- (1974) *The Roots of Reference*. Open Court.
- (1975) 'The nature of natural knowledge', in S. Guttenplan, ed., *Mind and Language*. OUP.
- (1981) *Theories and Things*. HUP.
- (1986) 'Reply to Morton White', in L. Hahn and P. Schlipp, eds, *The Philosophy of W.V. Quine*. Open Court.
- (1992) *Pursuit of Truth*, Revised edition. HUP.
- Rawls, J. (1971) *A Theory of Justice*. HUP.
- (1974–75) 'The independence of moral theory', *Proceedings and addresses of the American Philosophical Association*, 48, 5–22.
- Rorty, R. (1980) *Philosophy and the Mirror of Nature*. Blackwell.
- Rosenberg, A. (1990) 'Normative naturalism and the role of philosophy', *Philosophy of Science*, 57, 34–43.
- Ross, W. (1954) *The Nichomachean Ethics of Aristotle*. OUP.
- Siegel, H. (1990) 'Laudan's normative naturalism', *Studies in the History and Philosophy of Science*, 21, 295–313.
- (1996a) 'Instrumental rationality and naturalised philosophy of science', *Philosophy of Science*, 63 (Proceedings), 1996, S116–S124.
- (1996b) 'Naturalism and the abandonment of normativity', in W. O'Donohue and R. Kitchener, eds, *The Philosophy of Psychology*. Sage.
- Svendsen, J. (1996) *Vitenskap og rasjonalitet*. Gyldendal.
- Skorupski, J. (1990) *John Stuart Mill*. Routledge.
- (1995) Dictionary entry on John Stuart Mill, in T. Honderich, ed., *The Oxford Companion to Philosophy*. OUP.
- Spelke, E. (1995) 'Initial knowledge: six suggestions', *Cognition*, 50, 433–447.
- Stanovich, K. and West, R. (2000) 'Individual differences in reasoning: Implications for the Rationality Debate?', *Behavioural and Brain Sciences*, 23, 645–665.
- Stein, S. (1996) *Without Good Reason: The Rationality Debate in Philosophy and Cognitive Science*. OUP.
- Steward, H. (1997) *The Ontology of Mind: Events, Processes and States*. OUP.
- Stich, S. (1972) 'Grammar and indeterminacy', reprinted, in J. Katz, ed., *The Philosophy of Linguistics*. OUP, 1985.
- (1990) *The Fragmentation of Reason*. MIT Press.
- and Laurence, S. (1996) 'Intentionality and naturalism', in S. Stich and Laurence, S. (eds), *Deconstructing the Mind*. OUP.
- and Nisbett, R. (1980) 'Justification and the psychology of human reasoning', *Philosophy of Science*, 47, 188–202.
- and Ravenscroft, I. (1996) 'What is folk psychology?', in S. Stich and Laurence, S. (eds), *Deconstructing the Mind*. OUP.

- Tversky, A. and Kahneman, D. (1982) 'Judgement under uncertainty: Heuristics and biases', in D. Kahneman, P. Slovic and A. Tversky, eds, *Judgement under Uncertainty: Heuristics and Biases*. CUP.
- Wason, P. (1969) 'Regression in reasoning', *British Journal of Psychology*, 60, 471-480.
- Wittgenstein, L. (1953) *Philosophical Investigations*, trans. E. Anscombe. Blackwell.
- (1956) *Remarks on the Foundations of Mathematics*, trans. E. Anscombe. Blackwell.
- Worrall, J. (1988) 'The value of a fixed methodology', *British Journal for the Philosophy of Science*, 39, 263-275.

Index

- Antony, L., 116
- Aristotle, 10, 136
- Bacon, F., 9
- Banham, G., viii
- Bernúdez, J., 172 n.14
- Blackburn, S., 175 n.4
- Boghossian, P., 153-4 n.17
- Bohr, N., 45, 91
- Bonjour, L., 152 n.6
- Brown, H., 13, 122, 151 n.1,
153 nn.12, 14, 17, 154 n.23, 155
nn.26, 48, 165 n.15, 167 n.48,
170 nn.78, 80, 172 n.16
- Campbell, D., 165 n.17
- Carnap, R., 35, 153 n.17, 157 n.1
- Carrier, M., 171 n.81
- Carroll, L., 5, 147-8, 175 n.3
- Chalmers, A., 54, 55, 109, 150 n.5,
159 n.21, 160 nn.5, 7
- Cherniak, C., 155 n.28, 158 n.9,
161 n.11, 172 n.12
- Chomsky, N., 42-3, 115, 116, 124,
158-9 n.14, 162 n.20, 171 n.7,
173-4 n.28
- Cohen, L.J., 42-4, 124, 126, 158 n.13,
159 n.18
- Collins, H., 160 n.5, 170 n.70
- Couvallis, G., 70, 160 n.23, 164 n.8
- Crane, T., 156
- Darwin, C., 45, 127
- Davidson, D., 1, 126, 136-7,
174 n.33
- Descartes, R., 2, 9, 14, 18, 19,
85, 129
- Doppelt, G., 166 n.30
- Duhem, P., 25
- Dupré, J., 156 n.37
- Ehring, H., viii
- Einstein, A., 19, 28, 45, 85, 91
- Euclid, 10, 16
- Feyerabend, P., 89, 160 n.23
- Fidjestøl, A., viii
- Fodor, J., 1, 25, 155 n.33, 156 n.35,
173 n.27
- Freedman, K., 87, 95, 160 n.6,
166 n.30, 167 n.50
- Galileo, G., 45, 53, 91
- Gelman, S., 158 n.15
- George, A., 159 n.17
- Goldman, A., 156 n.47, 158 n.9,
161 n.11, 168 n.59, 172 nn.12, 15
48, 114, 154 n.18, 157 n.6,
164 nn.8, 12
- Hempel, C., 174 n.40
- Henderson, D., viii, 168 n.59,
170 n.69, 175 n.44
- Higginbotham, J., 159 n.17
- Hirschfield, L., 158 n.15
- Høibraathen, H., 169 n.67
- Hookway, C., 157 n.1
- Horwich, P., 150 n.3, 153 n.17
- Hume, D., 9, 15, 17, 19, 32, 34,
153 n.17, 175 n.4
- Janvid, M., viii, 152 n.6
- Kahneman, D., 158 n.12
- Kant, I., 9, 22, 129, 154 n.24,
155 n.27
- Kim, J., 1, 31, 150 n.2, 156 nn.35, 45,
46, 163 n.1
- Kitcher, P., 151 n.1, 152 n.9,
155 n.27, 168 n.59, 169 n.68

- Knowles, J., vii, 156 nn.38, 39,
158 n.9, 159 n.16, 169 n.66,
171 n.4, 174 nn.35, 38
- Kornblith, H., 119, 126, 156 n.47,
157 n.4, 159 n.16, 162-3 n.24,
165 n.13, 168 n.59, 170 n.79,
171 n.1, 172 n.8
- Kosso, P., 159 n.21, 160 n.24
- Kripke, S., 152 n.9, 175 n.6
- Kuhn, T., 9, 14, 55, 60-4, 83, 85,
108, 146, 160 n.8, 161 n.17,
162-3 nn.18, 19, 21, 170 n.72
- Lakatos, I., 46, 89, 155 n.29,
160 n.23, 167 n.38
- Laudan, L., vii, 5, 54, 64, 69,
83-99, 100, 101, 104, 108, 111,
153 n.13, 159 n.21, 160 nn.5, 6,
162 n.22, 163 n.3, 5, 165 n.15,
166 nn.18, 19, 20, 22, 23, 26,
29, 30, 31, 37, 167 nn.39, 45,
46, 48, 49, 50, 52, 168 nn.53,
55, 56, 59, 62, 170 nn.73, 76,
171 n.81, 175 n.1
- Lavoisier, A., 45
- Leibniz, G., 19
- Linnebo, Ø., viii, 168
- McDowell, J., 1, 118, 135-41, 174
- Maffie, J., 151 n.1
- Markman, E., 159 n.16
- Maxwell, C., 45, 91
- Mellor, H., 156 n.36
- Mell, J.S., 5, 9, 68, 69-78, 80, 122,
164 n.6
- Moore, G.E., 32
- Newton, I., 45, 53, 85, 90, 91
- Nisbett, R., 173 n.21
- Pacherie, E., 150 n.2
- Papineau, D., 156 n.35, 164 n.8
- Peacocke, C., 175 n.43
- Pinch, T., 160 n.5, 170 n.70
- Pinker, S., 158 n.15
- Plato, 19, 51, 136
- Polanyi, M., 60-2, 64, 161 n.17
- Popper, K., 9, 18, 165 n.17
- Putnam, H., 152 n.6
- Quine, W.V.O., 1, 3-4, 9, 16,
24-31, 34-7, 41, 47, 48, 64,
100, 114, 115, 116, 123, 126,
142, 150 nn.1, 6, 151 n.8,
155 n.31, 33, 156 nn.40, 42,
45, 48, 157 nn.1, 2, 3, 4, 5,
160 n.4, 165 n.13, 168 nn.59,
60, 61, 170 n.79
- Ravenscroft, L., 160 n.4
- Rawls, J., 39-40, 157 nn.7, 8
- Roll-Hansen, N., vii, viii
- Rorty, R., 150 nn.3, 156 n.34
- Rosenberg, A., 100, 168 n.62
- Ross, W., 174 n.30
- Sand, A., viii
- Siegel, H., 87, 152 n.6, 163 n.1,
166 n.30
- Sivertsen, J., viii
- Skorupski, J., 164 nn.6, 8
- Stanovich, K., 172
- Stein, E., 117-18, 120-1, 123, 124-9,
155 n.28, 156 nn.45, 47,
158 n.16, 159 nn.17, 18,
161 n.13, 14, 171 n.4, 172 n.7, 14,
17, 173 nn.22, 24, 25
- Stich, S., 58, 156 n.38, 158 n.10,
160 n.4, 161 nn.13, 14, 168 n.55,
59, 173 nn.21, 24, 27
- Strawson, P., 1
- Tversky, A., 158 n.12
- Wason, P., 158 n.12
- West, R., 172
- Whewell, W., 9
- Witgenstein, L., 5, 147, 148,
175 nn.5, 7
- Worrall, J., 168 n.54
- Ylikoski, P., 174 n.38