



LAURENCE BONJOUR AND ERNEST SOSA

EPISTEMIC JUSTIFICATION

INTERNALISM vs. EXTERNALISM,
FOUNDATIONS vs. VIRTUES

 **Blackwell**
Publishing

Epistemic Justification

Great Debates in Philosophy

Series Editor: Ernest Sosa

Dialogue has always been a powerful means of philosophical exploration and exposition. By presenting important current issues in philosophy in the form of a debate, this series attempts to capture the flavour of philosophical argument and to convey the excitement generated by the exchange of ideas. Each author contributes a major, original essay. When these essays have been completed, the authors are each given the opportunity to respond to the opposing view.

Personal Identity

Sydney Shoemaker and Richard Swinburne

Consciousness and Causality

D. M. Armstrong and Norman Malcolm

Critical Theory

David Couzens Hoy and Thomas McCarthy

Moral Relativism and Moral Objectivity

Gilbert Harman and Judith Jarvis Thomson

Atheism and Theism, Second Edition

J. J. C. Smart and J. J. Haldane

Three Methods of Ethics

Marcia W. Baron, Philip Pettit, and Michael Slote

Epistemic Justification

Laurence Bonjour and Ernest Sosa

Epistemic Justification

*Internalism vs. Externalism,
Foundations vs. Virtues*

Laurence Bonjour
and Ernest Sosa

© 2003 by Laurence BonJour and Ernest Sosa

350 Main Street, Malden, MA 02148-5018, USA
108 Cowley Road, Oxford OX4 1JF, UK
550 Swanston Street, Carlton South, Melbourne, Victoria 3053, Australia
Kurfürstendamm 57, 10707 Berlin, Germany

The right of Laurence BonJour and Ernest Sosa to be identified as the Authors of this Work has been asserted in accordance with the UK Copyright, Designs, and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by the UK Copyright, Designs, and Patents Act 1988, without the prior permission of the publisher.

First published 2003 by Blackwell Publishing Ltd

Library of Congress Cataloging-in-Publication Data

BonJour, Laurence, 1943–

Epistemic justification: internalism vs. externalism, foundations vs. virtues / Laurence BonJour and Ernest Sosa.

p. cm. – (Great debates in philosophy)

Includes bibliographical references and index.

ISBN 0-631-18283-7 (alk. paper) – ISBN 0-631-18284-5 (pbk.: alk. paper)

1. Justification (Theory of knowledge) I. Sosa, Ernest. II. Title. III. Series.

BD212 .B66 2003

121'.2 – dc21

2002015309

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

Set in 10 on 12.5 pt Melior
by SNP Best-set Typesetter Ltd, Hong Kong
Printed and bound in the United Kingdom
by MPG Books Ltd, Bodmin, Cornwall

For further information on
Blackwell Publishing, visit our website:
<http://www.blackwellpublishing.com>

Contents

<i>Acknowledgments</i>	vii
<i>Introduction</i>	1
I A Version of Internalist Foundationalism	3
<i>Laurence Bonjour</i>	
1 The Regress Problem and Foundationalism	5
2 Externalist Accounts of Justification	24
3 In Search of Coherentism	42
4 Back to Foundationalism	60
5 The Conceptualization of Sensory Experience and the Problem of the External World	77
II Beyond Internal Foundations to External Virtues	97
<i>Ernest Sosa</i>	
6 Knowledge and Justification	99
7 Does Knowledge Have Foundations?	119
8 Skepticism and the Internal/External Divide	141
9 A Virtue Epistemology	156
III Replies	171
10 Reply to Sosa	173
11 Reply to Bonjour	201
<i>Bibliography</i>	229
<i>Index</i>	233

Acknowledgments

Laurence Bonjour

My main essay is a development and piecing together of a number of earlier papers and presentations. I am grateful to many different audiences and critics, including many students, for suggestions and criticisms. I am also especially grateful to Ann Baker for reading and commenting extensively on both the main essay and my reply.

Ernest Sosa

In my main essay I have tried to make a coherent whole out of a diversity of materials found until now only in widely scattered publications. It has been helpful and rewarding to discuss these matters, over many years, with colleagues and students. Special thanks go to John Greco, Peter Klein, David Sosa, and Jim Van Cleve.

Introduction

Laurence Bonjour and Ernest Sosa

Epistemology is the theory of episteme, of knowledge. Ever since Plato it has been thought that one knows only if one's belief hits the mark of truth and does so with adequate justification. The issues debated by Laurence Bonjour and Ernest Sosa concern mostly the nature and conditions of such epistemic justification, and its place in our understanding of human knowledge.

Bonjour defends a traditional, internalist epistemology, according to which epistemic justification derives from the subject's (1) taking what is given to his conscious awareness, and (2) accepting claims or steps of reasoning on an a priori basis. Rejecting the emphasis of epistemology on the concept of knowledge, he is mainly interested in the question of whether we have or could have good reasons to believe in an external world of the sort that we normally take ourselves to inhabit, and in the question of what could possibly constitute such reasons. His answer to the latter question is internalist and foundationalist, in that it takes the justification for claims about the external world to begin from apperceptions of present states of consciousness (mainly sensory consciousness) and to proceed from there on the basis of (allegedly) a priori reasoning, specifically an argument that the truth of our beliefs about the external world constitutes the best explanation of our sensory experience.

Bonjour also rejects recent proposals according to which justification can derive from contingent factors external to the consciousness of the believer: factors involving how that belief is caused, or how well it tracks the facts, or how reliably it is formed. While he grants some lesser epistemic status to beliefs that do satisfy such external requirements of causation, tracking, or reliability, he insists that the more important issues for epistemology, and certainly the more prominent and important issues in the tradition, are the questions that he wishes

to address, concerning the internally accessible reasons that one might have for one's beliefs about the world around us.

Sosa had in earlier work (as BonJour points out) drawn a similar distinction, between animal knowledge and reflective knowledge, so on the issue of whether there are two importantly different kinds or levels of epistemic assessment they are in agreement. But there is still a relevant difference in focus and emphasis. Sosa is interested in understanding the conditions of animal knowledge, and not only those of reflective knowledge. His cognitive virtues account of animal knowledge is reliabilist. About such knowledge his views are thus in line with contemporary externalism. In distinguishing between animal and reflective knowledge, however, and in requiring reliability for animal knowledge his views agree surprisingly with Descartes's. Recall the passage early in Meditation Three where the *cogito* is said to derive its high epistemic standing from its clarity and distinctness, which, we are told, it could not possibly do if clarity and distinctness were to the slightest degree unreliable (and could ever lead us to a false belief). Sosa likewise takes reliability to be necessary in a source of justification, but of course not sufficient.

Sosa rejects the sort of internalist foundationalism favored by BonJour, while agreeing to put aside issues of knowledge and its conditions, in order to focus on epistemic, rational, justification. He agrees that a belief's having a reliable source is not enough to render it justified. The source must be a cognitive virtue seated in the subject. This already yields a kind of internalism. Moreover, the source must operate fundamentally through the promptings of experience, through either introspective or perceptual belief formation. Reflective justification goes beyond such unreflective rational justification in requiring a coherent epistemic perspective that underwrites the belief thus justified. What the externalist virtue theorist will add, in sharp disagreement with any kind of internalism, including BonJour's, is that there is no way to delineate what a cognitive virtue is in general, if we prescind from all contingent relations that such belief formation might bear to our external environment. In understanding rationality, having a reason, being reasonable, and the like, as these notions apply to empirical beliefs, we must make proper allowance for such external factors.

I

A Version of Internalist Foundationalism

Laurence Bonjour

The Regress Problem and Foundationalism

1.1 Introduction

The aim of this essay is to investigate one main aspect of what I take to be the central question of epistemology. That question concerns the rational status of our beliefs about the world in relation to the independent world that they purport to describe: Do we have any good reasons for thinking that our beliefs about the world, at least the main ones that we hold most firmly, are *true* or at least approximately true – any rational basis for thinking that they succeed in describing the world more or less correctly? And if so, what form do these reasons take? It is fairly standard to describe a belief for which such truth-conducive reasons exist as being *epistemically justified*; and I will adopt this usage here (often omitting the qualifier “epistemic” for the sake of brevity), though with the warning that the term “epistemic justification” has also been employed in somewhat different ways that we will eventually have to take note of.

Here and throughout, I will assume the correctness of the realist conception of truth as correspondence or agreement with the appropriate region or chunk of mind-independent reality.¹ Thus the issue to be discussed is what reasons we have for thinking that our beliefs stand in such a relation to the world.

For present purposes, I will confine my attention almost entirely to the epistemic justification of putatively *empirical* beliefs: those contingent beliefs seemingly held on the basis of reasons that derive, in ways that we must try to better understand, from sensory or

¹ Where the relevant sort of mind independence is only in relation to the specific cognitive act in question, thus allowing for the possibility that beliefs about mental matters may also be true in this sense.

perceptual *experience*.² In brief, then, our question is: How, if at all, are empirical beliefs about the world epistemically justified? (It is common to associate this issue with that of whether, and under what conditions, such beliefs constitute *knowledge*; but, for reasons briefly indicated at the end of this chapter, I will mostly focus on reasons or justification alone, setting the issue of knowledge aside.)

It is obvious that we ordinarily take ourselves to have good reasons or justification of this sort for a wide variety of seemingly empirical beliefs: beliefs about our immediate physical environment, about our personal past, about things and events elsewhere in the world, about history, about various results of science, and of course about our conscious experience itself. At the level of common sense, there is no trace of a general doubt about the accuracy of our empirical beliefs, nor any suggestion that our confidence in this area might be unfounded or fundamentally irrational. With relatively rare and localized exceptions, we also act with great confidence on the basis of these beliefs, and here too there is in general no hint of any serious uncertainty or doubt. But the effort to explain what our reasons for beliefs of these kinds actually involve or even how such reasons are possible turns out to be fraught with familiar and extremely recalcitrant difficulties.

It is largely in response to these difficulties that epistemology has lately found itself in a state of almost unparalleled ferment. The decades since the early 1980s have witnessed what may be fairly described as an explosion of epistemological discussion, with ever more new positions being suggested, elaborated, discussed, criticized – and then dismissed as untenable by at least large portions of the philosophical community. This situation might be taken to show that the subject is healthy and flourishing, despite the admitted failure of epistemologists to agree on very much, but many have instead drawn the pessimistic conclusion that epistemology is in its death-throes and should be abandoned as hopeless by all philosophers of good sense.³

2 There is also, of course, the traditional category of beliefs justified a priori, those for which the apparent reasons for thinking them to be true derive, not from sensory experience, but rather from pure reason or rational reflection alone. Reasons of this sort raise a largely different set of issues, which I have discussed elsewhere but have no space to consider here. (For a defense of a largely traditional conception of a priori justification, see my book *In Defense of Pure Reason* (London: Cambridge University Press, 1998). Hereafter cited as *IDPR*.)

3 Perhaps the best-known and most radical advocate of this sort of view is Richard Rorty in his book *Philosophy and the Mirror of Nature* (Princeton: Princeton University Press, 1979), but there are many others who have advanced views in the same general direction.

This latter conclusion strikes me as wildly premature at best, but it is hard to deny that it is one possible interpretation of the widespread and seemingly intractable disagreements that are presently to be found in this area.

Much of this recent discussion has been organized around two main dichotomies. On the one hand, there is the dichotomy between *foundationalist* and *coherentist* accounts of epistemic justification, and especially of empirical epistemic justification. Does such justification derive ultimately from “foundational” beliefs whose justification somehow does not depend at all on that of other beliefs, or does it derive instead from relations of coherence or agreement or mutual support among beliefs, with no appeal to anything outside the system of beliefs? On the other hand, there is the dichotomy between *internalist* and *externalist* accounts of such justification. Must epistemic justification depend on elements that are internal to the believer’s conscious states of mind in a way that makes them accessible to his conscious reflection (at least in principle), or might it derive instead from factors that are external to those states of mind, entirely outside the scope of his conscious awareness? These two dichotomies cut across each other, so as to generate four *prima facie* possible overall positions: internalist foundationalism, externalist foundationalism, internalist coherentism, and externalist coherentism. It is these positions, or rather the first three of them, around which the present discussion will be organized.⁴

The historically standard and seemingly obvious view of empirical justification, reflected in a great tradition stretching from Descartes through Locke, Hume, Kant, and many others, up to recent figures like

4 Externalist coherentism, as will be explained later, combines what turn out to be the less attractive sides of each of the two dichotomies. What results is a view that has very little in the way of intuitive or dialectical appeal, and that, not surprisingly, has rarely if ever been explicitly advocated. (Though it can be viewed as a dialectical pitfall into which would-be internalist coherentist positions have a rather alarming tendency to fall.)

The current epistemological literature also contains defenses of a variety of further alternatives, most of them falling into two main groups. First, there are views that are hybrids of various kinds between the main alternatives listed in the text. Such hybrids seem to me to inherit all of the difficulties pertaining to the main views that they attempt to combine, thus being in general less attractive than any of them. Second, there are also views, such as contextualism and the various forms of “naturalized epistemology,” whose *modus operandi* is in effect to evade the central epistemological issue formulated in the text. All such further alternatives will be set aside in the present discussion, though what is said here will be strongly relevant to the assessment of the first of these two groups.

Ayer, Lewis, and Chisholm, is, of course, internalist foundationalism. But a central theme of recent epistemological discussion has been a widespread retreat from, and repudiation of, this historically dominant view. Indeed, almost the only point on which large numbers of otherwise widely disparate epistemologists agree is the conviction that internalist foundationalism is an untenable, indeed hopeless, position and must be abandoned if epistemological progress is to be made. There are serious reasons for this view, the most important of which will be considered shortly. But it is worth pointing out that it is far from clear that there is any general agreement among such epistemologists about the specific deficiencies of foundationalism. Indeed many of those who reject it seem to have no very definite argument in mind. As happens with alarming frequency in philosophy, the recent movement away from internalist foundationalism often looks less like a reasoned dialectical retreat than a fashionable stampede. And it is of course the rejection of internalist foundationalism that provides the primary motivation for both internalist coherentism and externalist foundationalism, neither of which could plausibly be claimed to be intuitively very plausible if they were not viewed by their proponents as the only viable dialectical alternative, once internalist foundationalism has been rejected.

I myself have played a role in this development, offering some of the arguments against foundationalism of both the internalist and externalist varieties and attempting to develop and defend the internalist coherentist alternative. But I am now convinced that the rejection of internalist foundationalism is a serious mistake, one that is taking epistemology very much in the wrong direction and giving undeserved credibility to those who would reject the central epistemological project altogether. My initial reasons for this judgment are dialectical. Both coherentism and externalism seem to me to be quite unsatisfactory as responses to the deepest epistemological issues, albeit in quite different ways, and there seems to me to be no further alternative to internalist foundationalism that does any better. But I also think that I can now see the way around the most serious objections to internalist foundationalism, and that this also brings with it some idea of what a viable internalist foundationalist position might look like.

Thus my eventual, albeit still tentative, thesis will be that internalist foundationalism, in something approximating its classical Cartesian form, is defensible and essentially correct as an account of empirical epistemic justification – though some of the ideas reflected in

coherentist views will also play a subsidiary role in the position that I will ultimately defend. As already noted, this view represents a very substantial departure from my previous epistemological views.⁵ Since part of my goal in this essay is to correct what I now regard as my own previous and very serious mistakes – to confess my epistemological sins, as it were – I will be compelled to make somewhat more reference to those previous views than might otherwise be appropriate.

In the rest of the present chapter, I will offer a fuller account of the main issue and consider in a preliminary way some of the reasons why the initially appealing internalist foundationalist solution has been so widely rejected in recent epistemology.

1.2 The Regress Problem

As we have seen, our initial question can be formulated with what may turn out to be a deceptive simplicity: What reasons, if any, are there for thinking that our various empirical beliefs⁶ are true, or at least likely to be at least approximately true – qualifications that I will normally omit? How, if at all, are our beliefs *epistemically justified*?

Perhaps the least complicated answer to this question is the one offered by the most obvious form of skepticism, which claims that there are in fact no such reasons, no epistemic justification, for any empirical belief. Such an extreme version of skepticism is obviously extremely implausible from an intuitive or commonsense standpoint – and also something that the extreme skeptic himself cannot, on pain of contradiction, claim to have any reason for believing to be true. Clearly this skeptical answer to our question is quite unappealing; it would be foolish to even consider accepting it until all other alternatives have been thoroughly explored. But at the same time, even extreme skepticism cannot be simply *assumed* to be false if the original epistemological issue is to be taken seriously.

5 Mainly as presented in my book, *The Structure of Empirical Knowledge* (Cambridge, Mass.: Harvard University Press, 1985). (Hereafter referred to as *SEK*.)

6 I will conform to the fairly standard practice of recent epistemologists by using the term “belief” to stand for any state, whether dispositional or occurrent, whose content is the acceptance of a proposition. Thus “beliefs” will include conscious acts of acceptance or assent as well as the formed disposition to engage in such acts when the appropriate issue is raised. But see Paul Moser, *Knowledge and Evidence* (London: Cambridge University Press, 1989), pp. 13–23, for a plausible critique of this standard usage.

But what other alternatives are there? A skepticism that is confined to empirical beliefs and their justification will be likely to grant that many of my putatively empirical beliefs are logically and probabilistically interrelated in such a way that if a particular belief or conjunction of beliefs were somehow either justified already or perhaps merely assumed to be true, this would provide a proximate reason for thinking that some further belief was true. An explicit statement of such a reason would take the form of an argument or inference from the former belief or conjunction of beliefs as premise to the latter belief as conclusion. The inferential connections involved in such *conditional* or *inferential reasons*, as I shall call them, obviously raise justificatory questions of their own, and more would have to be said about them in a complete epistemological account.⁷ But they can reasonably be taken for granted where, as on the present occasion, it is empirical justification that is our concern. Thus it will be useful to imagine a skeptic who is willing to accept conditional reasons of this sort and see how far we can get on this basis.

Conditional reasons by themselves do not, however, speak very directly to our original problem, since the vast majority of the things that we ordinarily think we have reason to believe to be true are obviously not in this way conditional in form. And it seems obvious at once that the existence of a conditional reason can provide a reason or justification for its non-conditional consequent only if there is some further reason or justification, which must seemingly be epistemically prior, for accepting the truth of its antecedent.

In this way the issue of epistemic justification for one belief may be in effect transformed, via an appropriate conditional reason, into the issue of justification for one or more other beliefs. Clearly this process can be repeated, in principle at least through many stages, yielding an epistemological tree-structure in which a belief at one level is conditionally justified in relation to those at a successive level, those at that level in relation to others at a still further level, and so on.⁸ Equally clearly, however, the delineation of an epistemological structure of this sort does nothing by itself to show that *any* of the non-conditional beliefs that appear in it are true. It remains open to a would-be skeptic,

7 The traditional view, which seems to me essentially correct, is that the justification of these inferential connections is ultimately a priori in character. See *IDPR*, cited in note 2.

8 For an elaboration of this picture, see Ernest Sosa, "How do you know?," *American Philosophical Quarterly*, 11 (1974), pp. 113–22.

even if he⁹ concedes all of the conditional reasons involved, to reject any non-conditional belief in the structure simply by rejecting some or all of the premises upon which that belief's justification conditionally depends.

The foregoing picture leads directly to a version of the classical *epistemic regress problem*. At each node of the tree the issue of justification for the previous non-conditional beliefs is conditionally answered by appeal to a new set of premises; at the next level, the issue of justification for those new premises is conditionally answered by appeal to yet further premises; and so on. The obvious problem is to say how this regress of levels or stages of justification, each dependent on the next, finally ends, assuming (as I shall here) that the finding of new sets of premise-beliefs (beliefs that have not previously appeared in the overall structure), each adequate to conditionally justify the premise-beliefs of the previous stage, cannot and does not go on infinitely. At first glance, at least, there seem to be only three general alternatives:

- 1 The final stage of a particular branch of the regress may invoke premise-beliefs for which no further reason or justification of any sort is available. In this case, it seems to follow that the inferential connections reflected in the epistemological tree-structure, no matter how complicated and ramified they may be, offer no reason or justification for thinking that *any* of the component

9 The skeptic in question might, of course, be either a man or a woman. It is a fact about the English language, or was at least until very recent times, that the pronouns "he" and "him" can be correctly used in a generic or neutral sense to refer to persons of either gender. Many, many recent writers have found this usage to be politically incorrect and so have attempted in various ways to undermine it, such as by using "she" and "her" as generic pronouns, either exclusively or in alternation with the more standard ones. My sense of this is that we may well be getting close to the point where "he" and "him" are no longer naturally understood in a generic way, but are still very far from having "she" and "her" assume the generic role (and still further from having both sets of pronouns function naturally at the same time in this way), with the attempts to use the feminine pronouns in this way serving only as flags displaying the political correctness of the author and serving to distract the reader, at least momentarily, from the main issues under discussion. This seems to me to simply be bad writing (as does the use of the plural "they" and "them," where only a single person is involved), but the question is what to do instead, given that generic pronouns are frequently needed in philosophical writing to refer back to the person holding a certain belief or view (such as skepticism). Though in writing the present work I at first used the disjunctive "he or she" and "him or her," the clumsiness of this seemed eventually too great to be tolerated. Thus I have chosen to adhere to the perhaps still viable generic use of "he" and "him," adding this footnote to remove any uncertainty about what is going on.

non-conditional beliefs that are essentially dependent on those unjustified beliefs are true.¹⁰ Those connections tell us, in effect, only that some things would be true *if* other things were true, and that those other things would be true *if* still further things were true, and so on, ending with things that there is no reason to believe to be true. Thus, if all relevant justificatory relations are captured by a structure whose branches terminate in that way, the view of the skeptic is apparently vindicated with respect to all non-conditional beliefs.¹¹

- 2 The final stage of a particular branch of the regress may invoke premise-beliefs that have occurred somewhere earlier in that branch, so that the justificational structure in effect loops back upon itself. In this case, the result seems once more to be skeptical (assuming again that all justificatory relations are captured by the structure), since the justification for all of the non-conditional beliefs in that branch is either directly circular or else dependent on premise-beliefs that are justified only in this circular and apparently question-begging manner. A justificational structure whose branches all terminate in this way again seems to provide no reason for thinking that any of the component non-conditional beliefs are true.
- 3 The only alternative apparently remaining is that the premise-beliefs at the final stage of a particular branch are indeed justified, but in some fashion that does not involve any further appeal to conditional or inferential reasons and thus does not require new premise-beliefs that would themselves be in need of justification. According to this third alternative, while there is still indeed a reason or at least a rational basis of some sort for thinking that each of these ultimate beliefs is true (so that this alternative differs from the first one), this reason or basis does not appeal to any sort of argument or inference from further premise-beliefs about which further issues of justification could be raised. For obvious reasons, these ultimate premises are standardly referred to as *basic* or *foundational* beliefs, and the epistemological position that advocates them as *foundationalism*.

10 I ignore here, for the sake of simplicity, the possibility of a situation of epistemic overdetermination, in which there are multiple independent justifications for a particular belief.

11 Obviously a particular epistemological structure might realize at different places two or even all three of the alternative outcomes of the regress discussed in the text. Explicit consideration of these further possibilities will be left to the reader.

1.3 Foundationalism, Coherentism, and Externalism

Historically, the foregoing dialectic, together with the assumption that global skepticism about non-conditional beliefs is false – that is, that we do have reasons for thinking that at least many of the non-conditional things that we believe are true – has usually been taken to show that some version of foundationalism and indeed of internalist foundationalism must be correct. According to the most historically standard version of internalist foundationalism, foundational beliefs are justified by appeal to sensory and introspective *experience*, and it is a version of this position that I will eventually attempt to defend.

Internalist coherentism (though sometimes without much stress on the internalist aspect) has been widely discussed (though not really very widely advocated) in recent epistemology. Such a theory may be viewed, albeit perhaps somewhat misleadingly, as growing out of the second of the three alternatives considered above about the eventual outcome of the regress of justification. It rejects empirical foundationalism and holds instead that *coherence*, roughly the agreement and mutual inferential support of empirical beliefs (a relation that is reflected in conditional reasons of the sort discussed above), is the primary or even exclusive basis for empirical justification. (It is obvious at once that making clear how such coherence-based justification can even count as *empirical* will be a major problem that such a view must address.)

As indicated above, there is also a third main alternative to be considered. The empirical foundationalism briefly adumbrated above as well as the most common versions of empirical coherentism are *internalist* in character: they appeal for justification to something of which the believer in question is allegedly aware, to which he has direct cognitive access. But many recent epistemologists have been attracted by views that are *externalist* in character: views that appeal for justification to factors, most commonly the reliability of the belief formation process, that may be and normally are largely or entirely outside the cognitive grasp of the believer. Most such views can in fact be viewed as alternative versions of foundationalism, with the external factors supplying the justification of the basic beliefs – though this structural feature becomes relatively unimportant from an externalist standpoint.

Thus we have three main views on the table: internalist foundationalism, internalist coherentism, and externalist foundationalism.

(For the sake of brevity, I will henceforward refer to the first two views as “foundationalism” and “coherentism,” and to the third as simply “externalism.”) I will devote the balance of the present chapter to a consideration of some of the reasons that have led to the widespread rejection of foundationalism in recent epistemology. Chapters 2 and 3 will then be devoted to externalism and coherentism, respectively. It will emerge that these two views are also afflicted with very serious – indeed, in my judgment, fatal – problems, when advanced as solutions to the central epistemological issue formulated above.¹² This will motivate a reconsideration of foundationalism in chapter 4, in which a version of traditional foundationalism will be argued to be defensible after all. Unfortunately, however, the specific foundationalist position in question will turn out to greatly aggravate the further issue of how an inference from the foundation to various beliefs about the physical world can be justified. Thus the final chapter will be devoted to a reconsideration of this venerable problem.

1.4 The Case against Traditional Foundationalism

In the most standard sort of empirical foundationalist position, as we have seen, justification is claimed to rest on a foundation of “basic beliefs,” beliefs that are alleged to be justified or at least epistemically acceptable without that justification or acceptability being itself dependent on inference from other beliefs (or on anything else that would itself require justification). It is upon these basic beliefs that the justification of all other empirical beliefs is supposed to depend, in the way already indicated: if the empirical beliefs that constitute the justifying reasons for any particular, non-basic empirical belief were specified, and then the further beliefs supporting any non-basic belief cited as such a reason also specified, and so on, all of the branches of the resulting justificatory structure would terminate sooner or later with basic beliefs.

There are two main kinds of objections that have been raised against such foundationalist views. The first focuses on the justificatory relation between the foundation and the superstructure of non-basic beliefs, as those two components are specified by any particular version of foundationalism. It questions whether it is in fact possible on the

¹² Externalism, as we will see, may still be a defensible and indeed valuable view for other purposes, some of which still fall within the general aegis of epistemology.

basis of the foundation thus specified to arrive at an adequate justification for the various sorts of beliefs that we ordinarily regard as justified, or at least for a reasonably high proportion of such beliefs. Here the most important beliefs whose justification is alleged to be problematic are beliefs about the physical world (assuming that, as is the case for many foundationalist views, these are not already part of the foundation). Clearly a foundationalist view that falls seriously short in this area will itself amount to a fairly severe and hence intuitively implausible version of skepticism.

It is obvious that the seriousness of this first general sort of problem will vary widely with respect to different foundationalist views, depending in large part on just how much is included in the set of basic or foundational beliefs. In particular, a less traditional foundationalist view according to which at least some beliefs about physical objects count as foundational will clearly have much less difficulty arriving at a reasonably plausible account of the overall scope of justified belief than will a more traditional view that restricts the foundations to beliefs about subjective states of experience. Since, as already noted, the general sort of foundationalist view that now seems to be otherwise defensible is of the more traditional sort, this problem will eventually become quite urgent.

For the moment, however, I want to focus on a second and to my mind much more fundamental kind of objection to foundationalism, one that asks how the supposedly basic or foundational beliefs are themselves justified or rendered epistemically acceptable. The basic beliefs in a foundationalist account of empirical justification are, after all, themselves contingent beliefs, beliefs that are true in some possible worlds and false in others. It thus seems obvious that if they are to serve as the justificatory premises for all the rest of empirical justification, then some sort of reason or rational basis for thinking that they themselves are true or at least likely to be true in the actual world is required. And the problem is that it seems initially impossible for there to be such a reason or rational basis for these allegedly basic beliefs (of an internalist character) that does not at the same time impugn their status as genuinely basic.

In fact, the characterization so far offered here of basic beliefs is almost entirely negative: though they are justified *somehow*, which I have taken to mean that there is a suitable reason or basis of some sort for thinking them to be true, this reason does *not* appeal to conditional reasons that would invoke further premises that would themselves be in need of justification. But how is this possible? How can there be a

reason or basis for thinking that a given claim is true that involves no inference or argument and no further premise of any sort? What might such a reason or basis consist in?

Foundationalists have responded to this challenge, sometimes only by implication, in a variety of ways. Some have claimed in effect that the issue of justification for the basic or foundational beliefs somehow does not arise or at least for some reason cannot be correctly or meaningfully raised. This sort of view seems, however, to be difficult or impossible to understand, especially given the already noted contingent character of the beliefs in question. The only intelligible way that a belief that is to serve as a foundation for the whole structure of other beliefs can itself be “not in need of justification” is if it already possesses something tantamount to justification (whether or not that term is employed), in which case this status needs to be further explained.

Other foundationalists have appealed instead to the idea that such beliefs are “self-justified” or “self-evident” (or “intrinsically justified” or “justified in themselves”), but it is difficult to attach a content to these characterizations that is both clear and defensible. A basic belief cannot be literally *self*-justifying unless the foundationalist accepts circular reasoning as a source of justification, a view that seems obviously wrong (and that would also undercut one of the main objections to coherentism). Nor can it be plausibly claimed that the foundational beliefs are self-evident in the sense that is sometimes claimed to apply to beliefs justified a priori. Whether or not that conception is finally defensible, there is obvious and substantial intuitive plausibility to the idea that when I understand a simple necessary truth, e.g., the proposition that $2 + 2 = 4$, I am able to directly and immediately apprehend on the basis of that understanding and without appeal to any further premise or argument that the claim in question must be, and so is, true. Such a proposition is naturally described as justified or evident in virtue of nothing more than its own intrinsic content, and in this way as *self*-evident. But whatever the ultimate merits of this essentially rationalist conception of self-evidence,¹³ nothing at all like it can be invoked for the sort of belief that is our immediate concern here, namely a belief having as its content a contingent proposition requiring empirical justification. Being contingent, true in some possible worlds and not in others, such a proposition cannot be seen or appre-

13 For a defense of a moderate version of rationalism, one that allows for the possibility that justification of this sort is fallible, see *IDPR*, cited in note 2.

hended to be true simply on the basis of its content.¹⁴ Moreover, to say that such a belief requires empirical or experiential justification is to say that it is precisely not *self-evident*, not justified merely by virtue of its intrinsic character or content, but rather, if at all, by something, *experience*, that is obviously external to that content.

Thus the obvious and, I now believe, correct thing to say is that basic or foundational beliefs are, after all, justified by appeal to experience. But the difficulty, which turns out to be very formidable, is to give a clear and dialectically perspicuous picture of how this is supposed to work.

Foundationalists such as C. I. Lewis and Richard Fumerton,¹⁵ among many others, have spoken at this point of basic beliefs being justified by the “direct apprehension” or “direct acquaintance” with the relevant experiential content. On the surface, however, this answer is seriously problematic in the following way. The picture it suggests is that in a situation of foundational belief, there are the following three distinguishable elements. First, there is the relevant sensory experience itself. Second, there is the allegedly basic or foundational belief, whose content, I will assume, pertains to some feature or aspect of that experience. And third, there is what appears to be a further mental act of some kind that is distinct from the belief, an act of direct apprehension of or immediate acquaintance with the sensory experience and its relevant features. And it is this further mental act that is supposed to provide the person’s reason for thinking that the belief is true.

Thus, for example, we might have, first, the actual presence in my visual field of a red triangular shape; second, the allegedly basic belief that there is a red triangular shape in my visual field; and, third, the direct apprehension of or immediate acquaintance with the red triangular shape. It is this third element that is apparently required for the view under consideration to differ from what would otherwise seemingly be a purely externalist view of the justification of basic beliefs. The suggestion is that the basic belief is justified, not merely because it in fact describes the experience correctly in a non-accidental way (which would be an externalist account), but rather because the character of the experience in virtue of which the description is correct is

14 Indexically formulated claims such as “I am here now” may be an exception here, but a basis for justification that goes beyond its mere content is needed for any claim with non-indexical descriptive content.

15 See C. I. Lewis, *An Analysis of Knowledge and Valuation* (La Salle, Ill.: Open Court, 1946) and Richard Fumerton, *Metaepistemology and Skepticism* (Lanham, Md.: Rowman & Littlefield, 1995).

cognitively *given* or *presented* to the person in question via the act of direct apprehension or immediate acquaintance, and is thereby cognitively accessible. At the same time, however, this direct apprehension or acquaintance is claimed to require no further justification itself, thus allegedly bringing the regress of justification to a close.

The immediate and obvious problem is to understand the nature of this apparently essential third element, the other two being at least reasonably unproblematic. Even if it is somehow not strictly a *belief*, is it still an assertive or judgmental cognitive act that involves something like a *conceptualization* or *classification* of the experiential element in question? Is what is directly apprehended or “given” something like the truth of the conceptually formulated *proposition* that there is a red and triangular experiential element present (rather than no such element at all or only one that is green and square or blue and oval or some other combination of color and shape)?

If the answer to this question is “yes,” then it is easy to see how this second cognitive act can, *if it is itself justified*, provide a reason for thinking that the belief in question is true. On most conceptions of direct apprehension, the content of the direct apprehension and that of the basic belief would not indeed be strictly identical, as the discussion so far might suggest, since the former would be more specific or determinate than the latter. But the truth or correctness of the direct apprehension that there is a red triangular shape in my visual field would nonetheless be sufficient for the truth of the basic belief that there is a red triangular shape in my visual field and hence would apparently provide an impeccable reason for accepting it – on the assumption, once again, that the direct apprehension is itself somehow justified or acceptable.

Just here, however, lies the apparent difficulty. Since on this construal a direct apprehension has as its content a contingent propositional thesis or assertion concerning the classification of my experience, some reason seems to be required for thinking that such a direct apprehension is *itself* true or correct. Such a reason obviously cannot be provided by the basic belief and to appeal merely to the first of the elements enumerated above, the sensory experience itself as distinct from any reflective awareness or apprehension thereof, would seemingly amount to a collapse into externalism. But having as its content a contingent claim, the direct apprehension cannot, as we have already seen, be strictly self-evident. And to say simply that acts of direct apprehension, unlike ordinary beliefs, somehow do not by their very nature require any further justification is merely to stipulate that

the foregoing problem is not genuine without offering any clear account of how and why this is so. Thus it is very hard to see why a direct apprehension or immediate acquaintance does not itself require some further sort of justification, presumably by appeal to some further sort of cognitive state, in which case the regress either apparently continues, even if it is perhaps no longer strictly a regress of belief, or else terminates in the first of the seemingly unsatisfactory ways considered above – with the result in either case being that the supposedly basic belief turns out not to be genuinely basic after all.

If, on the other hand, the answer to the question raised three paragraphs back is “no,” if the act of direct apprehension or immediate acquaintance is in no way assertive or judgmental in character, if it has no content that amounts to or approximates the proposition or thesis that the person’s experience has one set of features rather than another, then any clear reason for demanding epistemic justification for such a state vanishes. If such an awareness has as its content no claim or assertion that is even capable of being true or false, then the notion of epistemic justification, as understood so far, simply does not apply to it. At the same time, however, it becomes difficult to see how the occurrence of such a state can in itself provide any reason or other basis for thinking that the original allegedly foundational belief is true. If the direct apprehension of the experience involves no claim or assertion regarding its character, so that who thus has such an apprehension is apparently not thereby aware *that* it has such-and-such features, then in what way is his belief *that* he has an experience with those features justified by that apprehension? The basic belief, after all, *is* judgmental: it has the assertive content that something, in this case a sensory experience, has one set of features rather than one of the various others that it might have had. How can a state whose content does not in any way say or indicate that things are one way rather than another nonetheless provide a reason or any sort of basis for thinking that the propositional content of a belief that they are one specific way is true?

It is this dilemma,¹⁶ together with a related argument that will be discussed later, that seems to me to constitute the most basic difficulty for those traditional versions of empirical foundationalism that do not

16 The original source of the dilemma is Wilfrid Sellars’s “Empiricism and the philosophy of mind,” reprinted in Sellars, *Science, Perception and Reality* (London: Routledge & Kegan Paul, 1963), pp. 127–96, esp. pp. 131–2. See also Sellars, “The structure of knowledge,” in Hector-Neri Castañeda (ed.), *Action, Knowledge, and Reality: Critical Studies in Honor of Wilfrid Sellars* (Indianapolis: Bobbs-Merrill, 1975), pp. 295–347. For my own previous account and further elaboration of the dilemma, see *SEK*, ch. 4.

resort to externalism – a difficulty that once seemed to me to be clearly fatal. I will conclude this initial discussion of it by considering a possible rejoinder on behalf of the traditional foundationalist, one that now seems to me to point in the right direction but that nonetheless faces problems of its own.

One natural response for a foundationalist who takes the apparent dilemma seriously is to attempt to “go between the horns” by claiming that a state of direct apprehension or acquaintance is somehow *neither* fully assertive or judgmental nor entirely and unproblematically nonassertive and nonjudgmental. Rather such a state is, as it were, *semi*-assertive or *semi*-judgmental in character: it has a kind of content or cognitive significance, but not in a way that would raise a further issue of justification. Such states would thus allegedly resemble judgments or beliefs in having the capacity to confer justification on judgmental states proper, while differing from them in not requiring justification themselves.

But if this is to be more than a bare stipulation that the problem is somehow solved but without giving any hint of what the solution might be or how it is even possible, some further account is needed of how a state can have both of these properties. Some philosophers, perhaps most notably Husserl, have appealed at this point to the idea of a rudimentary cognitive state, prior not only to language but even to anything that is properly called conceptualization. Such a “pre-predicative awareness” would still *represent* or *depict* something, presumably experience, as being one way rather than another, but that representative content would be nothing like a propositional thesis or assertion, nothing that could be strictly true or false. In this way, it might be suggested, it could intelligibly provide a reason for a basic belief, while still being itself immune to the demand for epistemic justification.

The problem with this move is not that the idea of such a pre-conceptual cognitive state is untenable or even especially implausible (though many philosophers who are prone to identify intelligible thought-content with what can be linguistically expressed would surely be unhappy with it). The main difficulty is rather that any representative state that is capable of justifying a belief must somehow have as at least part of its content the *information* that the relevant state of affairs is one way rather than another, the way that the belief says it is rather than some way that would make the belief false. But for any representation that has an informational content of this sort, whether it is strictly conceptual or propositional or not, it will seemingly be possible to ask whether the information it presents is *correct* or *incor-*

rect in what it depicts, even if perhaps not strictly true or false. And once the issue of correctness has found a foothold, the issue of justification, of whether there is any good *reason* to think that the representation is correct rather than incorrect, will apparently follow immediately behind, and the regress will break out all over again.

As will emerge in chapter 4, I now think that a view in very roughly this general direction is defensible and indeed correct, if formulated and explained in the right way. But the present version, which attempts to avoid the need for justification simply by reducing or attenuating the conceptual or representative content of the direct awareness, plainly cannot succeed in itself. For no matter how far the representative character of the direct apprehension is thus diluted or attenuated, so long as it retains the capacity to justify a basic belief, there will be the same apparent reason for thinking that it itself requires justification. This is so because the very same aspect of such a state that allows it to justify a belief, namely its involving as its content the information that things are one way rather than another, also creates the apparent need for justification.

1.5 The Concept of Knowledge

There is one further topic that needs to be considered before turning to a fuller discussion of the alternatives to foundationalism. Though the present discussion falls within the general area usually referred to as “the theory of knowledge,” the concept of knowledge itself has itself barely been mentioned so far. This is not an accident and will indeed largely continue (with one important exception) to be the case, and I want to conclude this introductory chapter with a brief explanation of why this is so. In fact, for all of its prominence, both philosophically and commonsensically, the concept of knowledge is, in my judgment, a seriously problematic concept in more than one way. So much so that it is, I believe, best avoided as far as possible in sober epistemological discussion – as paradoxical as that may sound.

To begin with the most obvious difficulty, it is generally though not universally agreed that one necessary condition for knowledge is the possession by the belief in question of an adequate *degree* of epistemic justification or warrant in at least roughly the sense adumbrated above, that of there being a reason or basis for thinking that the belief is true (or likely to be true). But *what* degree of justification? How strong does such a reason have to be to satisfy this requirement? To require with

many historical philosophers that the reason be strong enough to *guarantee* the truth of the belief seems to restrict knowledge to a few simple necessary truths, such as simple propositions of mathematics and logic, together *perhaps* with simple claims about one's own private sensory and introspective experience.¹⁷ Faced with the obvious incompatibility between this result and the vastly more extensive knowledge ascriptions of enlightened common sense, epistemologists have generally adopted the view (sometimes referred to as the "weak conception" of knowledge) that there is some lesser degree or level of justification, lower than a guarantee of truth but presumably higher than mere 51 percent probability, that is required for a belief to count as "knowledge." The obvious question, however, is just what this crucial level of justification actually is or how it might be determined or specified. And the striking fact is that there is very likely no attempt actually to specify this favored level of justification that would be agreed by anyone beyond its author to have succeeded and very few that have even been very seriously attempted.¹⁸ Indeed, the sole reason for thinking that there actually is a specific level of this sort is that its existence is apparently the only way in which there can be a reasonably precise concept of knowledge that does not lead at once to the skepticism generated by the more traditional view (the "strong conception" of knowledge).

Moreover, a further problem is that if the levels of justification are thought of in the seemingly obvious way as something like degrees of probability in relation to the justifying premises or evidence, then the idea of a definite level of justification short of a guarantee of truth that is sufficient for knowledge seems to generate a serious conflict between the ordinary use of the concept of knowledge and the demands of the probability calculus: while someone who knows both proposition *A* and proposition *B* would ordinarily be thought to be able on that basis to come to know the conjunctive proposition *A and B*, the possession by both *A* and *B* of any proposed level of probability less than 1 in

17 For a persuasive recent defense of this traditional view as to the required degree of justification, see Panayot Butchvarov, *The Concept of Knowledge* (Evanston, Ill.: Northwestern University Press, 1970).

18 The main sustained attempt here is Roderick Chisholm's, in the three editions of *Theory of Knowledge* (Englewood Cliffs, NJ: Prentice-Hall, 1966, 1977, 1989) and in various other writings in the same period. Following the vicissitudes of Chisholm's attempts in this area is perhaps the best way to convince oneself that what might be referred to as the magic level of justification sufficient for "knowledge" as opposed to mere rational belief cannot be successfully specified and very probably does not exist.

relation to any specified body of evidence or justifying premises does not guarantee that their conjunction will possess at least that level of probability, making it hard to see how the envisaged “weak” conception of knowledge could work, even if a definite level of justification could somehow be specified.

These problems suggest strongly that to the seemingly elementary but obviously fundamental question of what degree of justification or warrant is required to satisfy the concept of knowledge, there is not only no satisfactory answer presently available but no real prospect of finding one. And that in turn suggests that our grip on the supposed concept of knowledge itself is anything but sure, if indeed there is even a clear and univocal concept there to be understood. For reasons like these, which will be augmented in chapter 2 and could easily be expanded yet further,¹⁹ I will largely concern myself here with justification rather than knowledge. My conviction, which cannot be further defended now (but for which the success of the following discussion would constitute good evidence), is that such an approach is adequate to the issues arising out of the foregoing dialectic and indeed to all of the central issues of traditional epistemology.

19 An important further source of such reasons is the notorious “Gettier problem,” which is, I believe, largely if not entirely an artifact of the idea of a weaker level of justification that is still adequate for “knowledge.” It is also worth asking what the significance of a “magic” level of justification short of a guarantee of truth could possibly be. Even if we had reached such a level in a particular case, there would still be no reason not to seek still higher levels of justification for any claim whose truth was a matter of serious interest, nor would increases in justification become in any clear way less valuable once the “magic” level had been obtained. This again seems to call into question whether the concept of knowledge as understood by the weak conception could possibly have any real importance.

Externalist Accounts of Justification

The fundamental problem for empirical foundationalism, as it has emerged so far, could scarcely be simpler or more straightforward. The foundationalist's basic beliefs are claimed to be epistemically justified, so that it is incumbent on him to explain how there can be a reason for thinking them to be true that is nonetheless still compatible with their status as basic. Since the basic beliefs we are presently concerned with are contingent and empirical in character, they cannot, as we have seen, be strictly self-evident; hence the justifying reason must appeal to some condition or state of affairs beyond the belief itself, such as the occurrence of an experience of an appropriate kind. But then the *internalist* foundationalist needs to explain how that further justifying element can be cognitively available to the person in question, so that it can provide a reason *for him* to think that the basic belief is true. Any attempt at such an explanation must apparently involve the invocation of a further cognitive state of some sort that represents or depicts the element in question as being of the right sort: a cognitive state about which a further issue of epistemic justification can seemingly be raised. The result is that the supposedly basic belief is apparently not basic after all, so that the regress of justification either continues or else ends with beliefs or belief-like states that are not justified at all and so cannot be genuine sources of justification.

2.1 The Externalist Gambit

It is apparent that one key ingredient in the foregoing line of argument is the internalist idea that the justifying reason for a basic belief, or indeed for any belief, must somehow be cognitively available to the believer himself, within his cognitive grasp or ken. At first glance, the

credentials of this idea seem quite clear and straightforward: If the belief is to be justified *for that particular person* (rather than perhaps for someone else), then it seems at least initially obvious that the reason for thinking it to be true must be one to which that person himself has access. How, it might reasonably be asked, could *that* person be justified in accepting a belief in virtue of a reason or basis for thinking it to be true that is cognitively available only to others – or, as we will see, on most versions of externalism ultimately to no one at all? Why would a reason that is unavailable to that person be even relevant to the epistemic justification of *his* belief? The foregoing rationale for a requirement of cognitive availability reflects, of course, the *internalist* view of epistemic justification briefly discussed in chapter 1. Despite occasional suggestions to the contrary,¹ it seems clear that such a conception was taken utterly for granted by virtually all epistemologists until very recent times.

It is this internalist requirement that the justifying reason be cognitively available to the believer in question that externalist views propose to discard. To be sure, the suggestion is not that just *any* reason for thinking that a belief is true that is not thus available can justify the acceptance of the belief – a suggestion that would seemingly mean that virtually all true beliefs and very many false ones are justified, and indeed basic, for everyone. Rather the idea, in first approximation, is that certain special sorts of reasons, ones that although cognitively unavailable to the person in question are nonetheless intimately connected in a different way with the operation of his cognitive processes, can play this sort of role.

Though externalist views can take many different forms, the versions most widely discussed and advocated have been versions of *reliabilism*, and it is on reliabilist views that I will mainly focus in this chapter. According to the reliabilist, the main requirement for epistemic justification is roughly that a belief be produced or caused in a way or via a process that makes it objectively likely that the belief is true. Such a mode of belief production is thus a *reliable* source of true beliefs. As we will see further below, reliabilist views differ among themselves with regard to whether a belief's being produced in a reliable way is by itself *sufficient* for epistemic justification or whether there are further requirements that must be satisfied as well. But what is *not* required for justification on any such view is that the person for whom the belief is justified be in any way *aware* (whether justifiedly

1 See, e.g., Frederick Schmitt, *Knowledge and Belief* (London: Routledge, 1992).

or not) that the belief is produced in a reliable way. In the absence of such an awareness, that person will also in general be aware of no reason of any sort for thinking that the belief is true. It is the insistence that the cognitive availability of such a reason is unnecessary for epistemic justification that is the distinctive – and problematic – feature of externalism.

An externalist view need not be offered merely as an account of the justification of beliefs that are basic or foundational in the sense explained above, but may instead claim to provide an account of all empirical justification or even of epistemic justification generally. But the plausibility of externalism is not particularly enhanced by extending it more widely; and it will simplify the discussion and also clarify the comparison with other views to limit our consideration here to the attempt to use the externalist conception of epistemic justification to account for the justification of empirical, contingent beliefs that are allegedly basic or foundational in the way explained above. Such an account can offer an explanation of how basic empirical beliefs can be genuinely justified without that justification depending on a further cognitive state in a way that would undermine their status as basic, and can thereby provide a solution to the epistemic regress problem.

The most obvious and widely held versions of such a view appeal to the (*apparent!*) reliability of our ordinary perceptual and introspective modes of belief formation. Their main claim is roughly that if those processes are reliable in the way that we ordinarily believe them to be, i.e., if the beliefs that they produce are in fact mostly true, then those beliefs are thereby, possibly given also the satisfaction of certain further requirements, epistemically justified – and that this is so whether or not our commonsensical belief in the reliability of those processes is itself justified in any way or indeed whether or not such a further belief is even held by the individual believer in question.²

2.2 Objections to Externalism

But although the foregoing dialectical motive for externalism is abundantly clear, it is nevertheless far from obvious that what results is a

2 The concept of reliability embodied in such views is somewhat less straightforward than it initially seems and can in fact be further spelled out and refined in a number of different ways; but the differences among these are relatively unimportant for the main issues that will mainly concern us and thus can safely be ignored here. For discussion of this issue, see Alvin Goldman, *Epistemology and Cognition* (Cambridge, Mass.: Harvard University Press, 1986), pp. 43–51.

plausible account of epistemic justification. I have already remarked that externalist views represent a major departure from the main epistemological tradition, and indeed such views may well be suspected of being merely ad hoc in relation to the difficulties arising from the epistemic regress problem. Why, the internalist will ask, should a reason that is outside the cognitive grasp of a particular believer nonetheless be taken to confer epistemic justification on his belief? Is this not indeed contrary to the whole idea of epistemic justification, which surely has something to do with selecting one's beliefs responsibly and critically and above all *rationally* in relation to the cognitive goal of truth? How can the fact that a belief is reliably produced (or indeed any sort of fact that makes a belief likely to be true) make my acceptance of that belief rational and responsible when that fact itself is entirely unavailable to me? No doubt I am quite fortunate if my perceptual and introspective beliefs happen to be caused or produced in such a way, for then such beliefs will in fact be mostly true and my actions on the basis of them will tend, other things being equal, to be successful; indeed, such cognitive serendipity might, at least if it were thoroughgoing enough, be in the end preferable from a purely practical standpoint to anything that mere reason could be very confidently expected to deliver. But none of this has any clear bearing on the issue of epistemic justification, which has to do, it will be argued, with the rationality or irrationality of one's beliefs, rather than with what appears to be, from the standpoint of the believer, mere cognitive luck, however practically helpful it may be.

This general intuitive objection to externalism may be developed in two quite different ways. Some philosophers have tried to show that external reliability is not *necessary* for epistemic justification by describing examples in which such justification seems intuitively to exist even in the face of modes of belief production that are not in fact reliable.³ Here the favorite example is that of persons in a world controlled by a Cartesian evil demon: though their reasoning processes and collection of evidence may be quite impeccable, so that they seem intuitively to be justified in their resulting beliefs, their belief-forming processes are in general quite unreliable, owing to the pervasive interference of the demon.

But this sort of objection, though forceful against a general externalist position, has no direct force against one that appeals to externalist justification only as *sufficient* in certain cases, like that of

3 See Stewart Cohen, "Justification and truth," *Philosophical Studies*, 46 (1984), pp. 279–95.

allegedly foundational beliefs, where other sorts of justification are lacking. Thus the more fundamental sort of objection will be one that challenges the *sufficiency* of external reliability for justification in any case, even these. One way to do this is to describe apparently parallel examples in which beliefs that are in fact reliably produced nonetheless seem from an intuitive standpoint clearly not to be justified. It is this latter, ultimately more fundamental objection to externalism that I want to briefly develop and discuss here. (As we will see shortly, the divergent externalist responses to this objection lead to a major bifurcation in the externalist position.)

The objection in question depends on describing possible modes of belief production that are reliable in the indicated sense without that reliability being cognitively accessible to the believer in question. Suppose then that Amanda is a reliable clairvoyant under certain specific conditions with respect to a particular range of subject matter. Owing perhaps to some sort of causal process that has so far eluded scientific investigators, beliefs about that subject matter now and then occur spontaneously and forcefully to Amanda under those conditions, and such beliefs are mostly or even perhaps invariably true. Amanda, however, though she generally accepts the beliefs in question, has never checked empirically to see whether any of them are true, nor has the potentially available empirical evidence for the truth of any of the specific claims and in consequence for her general reliability been supplied to her by others.⁴

Consider some particular small set of beliefs that Amanda arrives at in this way. Are these beliefs epistemically justified, as the most straightforward version of reliabilism must say? The answer to this question will depend, to some extent at least, on the further specification of the example, with respect to which there are two main alternative possibilities. First, suppose that Amanda, in addition to having no reason in her cognitive possession that positively supports the thesis that the beliefs in question are reliably produced or otherwise likely to be true, also has a substantial body of seemingly cogent evidence *against* this thesis. Such evidence might take a number of different forms, including at least the following: general scientific evidence apparently showing that there is no possible way in which beliefs concerning that subject matter could be reliably produced under those conditions; seemingly compelling (though in fact mistaken)

4 Here I am more or less following the case of Norman, developed in *SEK*, pp. 41–5.

empirical evidence that Amanda's own past beliefs of the kind in question have been mostly in error; and empirical evidence seeming to show independently that the various specific claims presently at issue are false.

From an intuitive standpoint, it seems abundantly clear that Amanda would be highly irrational and irresponsible if she were to continue to hold beliefs of this sort in the face of apparently cogent evidence (which may be made as strong as one likes) that shows in one of these ways that the beliefs in question are very unlikely to be true. And this epistemic irrationality does not seem to be in any way mitigated by the fact that entirely unbeknownst to her, those beliefs are in fact caused in a way that makes them likely to be true. Thus if there is, as seems initially apparent, a close connection between the idea of epistemic justification and the idea of seeking rationally and responsibly to find the truth, we must conclude that the beliefs in question are not justified.

This intuitive assessment may be reinforced to some extent by supposing also that Amanda is forced to make some critical, life-and-death choice, where the proper action depends on the truth or falsity of the beliefs in question. There is little or no plausibility to the claim that she would be *acting rationally* if she allowed her action to be guided by her in fact reliable clairvoyant beliefs, instead of letting it be guided by other beliefs for which reasonably strong justification is cognitively available to her – even if the latter beliefs are perhaps somewhat less likely to be true from the objective, God's-eye point of view to which she has no access. (Indeed, on this last point, it is hard to think of anyone who has ever seriously advocated a contrary view.) Thus the proponent of externalism who insists that the clairvoyant beliefs in question are nonetheless epistemically justified must apparently sever either the connection between epistemic rationality and epistemic justification or else that between epistemic rationality and rationality in action. He must say either that the belief that has the highest degree of epistemic justification need not be the one that it is, from an epistemic standpoint, most rational to hold, or else that the belief that it is epistemically most rational to hold need not be the one that it is most rational to act on. But each of these alternatives seems at least initially to be quite implausible from an intuitive standpoint.

This sort of objection has seemed to many epistemologists to be extremely telling, indeed more or less conclusive, against an unqualified externalism of the sort we have been concerned with so far. To be sure, when faced with this objection, there are occasional externalists who simply dig in their heels, bite the bullet, and insist that the exter-

nalist account of epistemic justification is correct nonetheless, that the apparent intuitions to the contrary are simply mistaken. Perhaps the most straightforward move here is simply to repudiate the idea that there is *any* connection between epistemic justification and epistemic rationality. But then the idea of epistemic justification seems to lose most or all of its intuitive content, with the term becoming little more than a label for *whatever* it is that is required to turn true belief into “knowledge” (perhaps given the satisfaction of certain further conditions as well).⁵ But in addition to a considerable burden of intuitive implausibility, such views seem to me, for reasons that will be indicated more fully at the end of this chapter, to have lost any real bearing on the main epistemological issues.

2.3 Qualified Externalism

The other, perhaps rather more plausible externalist response to cases like the Amanda case is to concede that Amanda’s clairvoyant beliefs are not justified when she possesses strong evidence of one or more of the kinds specified against their truth, while insisting at the same time that this does not show that the basic externalist idea must be abandoned. Instead, it is suggested, such cases show only that the satisfaction of the reliability requirement is not sufficient by itself for justification, but must be supplemented by a further requirement. This further requirement can be formulated in a number of different ways, but the basic idea is roughly that the believer in question not have in his possession strong reasons of an *internalist* sort for thinking that the belief is false or that the cognitive process is unreliable.⁶ It is this further requirement that Amanda, in the version of the case specified above, fails to satisfy. But, it is alleged, as long as countervailing reasons of this sort are not present, the satisfaction of the reliability

5 Indeed, it is not surprising that some of those whose views lie in this general direction have indeed chosen to downplay the idea of “epistemic justification” and simply make the satisfaction of the externalist requirement itself one of the necessary conditions for the application of the term “knowledge.” This is in effect what Alvin Plantinga does in his *Warrant and Proper Function* (Oxford: Oxford University Press, 1993). Such views will be considered further below.

6 See, e.g., Alvin Goldman, “What is justified belief?” in George Pappas (ed.), *Justification and Knowledge* (Dordrecht: Reidel, 1979), pp. 1–23. Goldman’s added requirement is more general than this, but both his leading example and his subsequent discussion make it clear that it is intended to exclude the possession of strong internalist reasons of the sort indicated and unclear that it excludes anything else.

requirement will still be enough by itself for epistemic justification. Thus if Amanda's clairvoyant belief is reliably caused in the way already suggested, and if in addition she possesses no reasons of the various sorts indicated for doubting it, then her belief is justified. And this is enough to make possible an externalist version of foundationalism, even if the resulting set of basic beliefs turns out to be somewhat narrower than it would otherwise have been.

Call the foregoing sort of view *qualified externalism*. One way to put it a bit more perspicuously is to distinguish explicitly between two main aspects of the opposing internalist view of epistemic justification. According to the internalist, justification requires both that the believer have available positive reasons (of sufficient strength) for thinking that the belief is true (call this the requirement of *positive internal justification*) and also that he not have available negative reasons (of sufficient strength) for thinking that it is false (call the requirement that such reasons be absent the requirement of *negative internal justification*).⁷ The modified version of externalism that we are now considering in effect accepts one part of the internalist view, the requirement of negative internal justification, while rejecting the other part, the requirement of positive internal justification. But is such a compromise view plausible and clearly motivated? Or do the intuitive problems that force the externalist to retreat this far threaten to make even qualified externalism untenable?

One way to approach this issue is to consider more specifically the relevant variant of the clairvoyance example. Suppose then that Bertha resembles Amanda in being a reliable clairvoyant under certain conditions about a certain range of subject matter and also in having no empirical evidence at all that this is so, but differs in that she also possesses no countervailing reasons or evidence of any sort to show that beliefs arrived at in this way are unreliable or that the particular beliefs in question are false. In addition to being reliably caused, Bertha's clairvoyant beliefs thus satisfy the requirement of negative internal justification. According to qualified externalism, these beliefs are therefore epistemically justified.

But is this the right result? My suggestion is that the same intuitive considerations that led to the conclusion that Amanda's clairvoyant beliefs are not justified apply with only slight modification to Bertha's

7 For the internalist, such positive and negative reasons are, in general, commensurable with each other, and so the two requirements can be combined into one requirement that the overall balance of reasons or evidence be strongly in favor of the truth of the belief.

as well. We may assume that Bertha does not believe either that she is clairvoyant or, more generally, that her beliefs arrived at in this way are likely to be true. (For these beliefs would presumably have no epistemic justification of even an externalist sort, so that reliance on them would only make matters worse.) But then it appears that Bertha is still being epistemically irrational and irresponsible in accepting beliefs whose provenance can only be a total mystery to her, whose status is as far as she can tell no different from that of a stray hunch or arbitrary conviction. Here again, externalism seems to sunder the concept of epistemic justification entirely from the concept of epistemic rationality or responsibility, leaving the former concept with no clear intuitive content. (Again, the basic intuition may perhaps be sharpened slightly by supposing that Bertha must make a critical decision as to how to act and must rely either on the clairvoyant beliefs or on competing beliefs for which she possesses good internalist reasons, albeit ones that yield a lesser degree of likelihood of truth. It seems clear that it would again be irrational and irresponsible for her to choose to act on the clairvoyant beliefs, a result that, here as in the earlier case, can be reconciled with the claim of epistemic justification only by breaking either the intuitively obvious connection between epistemic rationality and epistemic justification or that between epistemic rationality and the species of rationality that applies to action.)

A different way of making essentially the same point is to ask whether accepting part of the internalist requirement for justification while rejecting the other part does not amount to an untenable halfway house. Is there any intelligible rationale for the requirement of negative internal justification, which the modified version of externalism accepts, that does not also support the requirement of positive internal justification, which it rejects? My suggestion is that the only clear reason for the negative requirement is that accepting beliefs that are, as far as one can tell from one's own cognitive perspective, unlikely to be true is plainly irrational and irresponsible from an epistemic standpoint that aims at truth – even if those beliefs happen to be, unbeknownst to the person in question, reliably caused. But this reason plainly supports the positive requirement as well. To be sure, it seems somewhat worse from the standpoint of epistemic rationality to accept beliefs that one has specific reasons to think are unlikely to be true than to accept beliefs about which one has no specific information at all, so that their apparent unlikelihood of truth merely reflects the general fact that the vast proportion of possible beliefs at any given level of specificity (setting aside beliefs in simple negations and other

logical compounds) are false. But that does not prevent the latter sort of acceptance from being still quite obviously irrational and irresponsible, as even the externalist would presumably agree in the case where the externalist justification is not present. In this way, it is hard to see how the qualified externalist can defend one part of his position without undermining the other.

Is there any response available here to the qualified externalist? Assuming that there is no plausible alternative rationale to be found for the requirement of negative internal justification, it seems that the qualified externalist must go still further and admit that the absence of *positive* internal justification also counts against the justification of a belief, but hold that this negative feature is (sometimes?) capable of being offset by the presence of external reliability. This would amount to treating external reliability and internal justification as two independent components of the overall concept of epistemic justification, components that are often in conflict with each other and between which a balance must somehow be struck. On such a view, the concept of epistemic justification would be a very odd concept indeed. It is difficult to see the rationale for such a bifurcated concept and even more difficult to see how the proper balance between the two components is to be determined in a non-arbitrary way.

2.4 Additional Arguments in Favor of Externalism

Thus the intuitive objections to externalism appear to be quite strong. But there are also serious arguments on the other side that need to be considered. One dialectical argument for externalism was implicit in the discussion at the beginning of the chapter. It combines the argument for foundationalism that appeals to the regress problem with the claim that internalist foundationalism is unacceptable for reasons like the ones considered in the previous chapter. I believe that this argument can be effectively answered, but doing so will have to await the reconsideration of internalist foundationalism in the last two chapters of this essay.

In this section, I will look at two other arguments that purport to show, this time on an intuitive basis, that internalism is unacceptable.⁸ The first of these begins with the allegedly commonsensical premise

⁸ Arguments for externalism also sometimes appeal to the currently fashionable but ill-defined doctrine known as “naturalism.” For a general critique of naturalism in epistemology, see my paper “Against naturalized epistemology,” *Midwest Studies in Philosophy*, 19 (1995), pp. 283–300.

that “knowledge” of a wide range of ordinary claims about the world can be unproblematically ascribed to relatively unsophisticated adults, to young children, and even to higher animals. Given the widely accepted assumption that one condition for “knowledge” is an appropriate level of epistemic justification, it would follow from the correctness of such ascriptions that the beliefs of such epistemic subjects are epistemically justified. But, it is claimed, most or all of these epistemically unsophisticated subjects fail to meet the requirements of any otherwise plausible *internalist* account of epistemic justification, whether foundationalist or coherentist in character, since many of the beliefs and inferences appealed to by such accounts are simply too subtle or complicated or sophisticated to be plausibly ascribed to them. Thus only an externalist view is compatible with these commonsensical ascriptions of “knowledge”; and this, given the strong intuition that common sense is correct in these cognitive assessments, constitutes a strong argument in favor of externalism.

Obviously a full assessment of this argument would have to depend on an investigation of the internalist alternatives, in order to see if it is indeed true that epistemic justification as characterized by those accounts cannot be plausibly ascribed to such unsophisticated epistemic subjects. I am inclined, however, to think that the externalist is at least approximately right about this – as will indeed be partially borne out by the discussions of coherentism and foundationalism later on in this essay. Both the most plausible (though, as we will see, still not *very* plausible) versions of coherentism and the only version of foundationalism that seems to me ultimately defensible seem to require more subtlety and complexity of both thought and inference than unsophisticated subjects can, in general, plausibly be said to engage in. Thus this first argument seems, at least initially, to have considerable intuitive force.

The second externalist argument is that no internalist account of epistemic justification, whether foundationalist or coherentist in character, can avoid an extreme version of skepticism in which few if any of the empirical beliefs of even sophisticated adults turn out to be justified or to constitute “knowledge.” But, it is argued, such an extreme skepticism is massively implausible from an intuitive standpoint, thus providing a very strong reason for thinking that all internalist accounts of justification are mistaken and that some kind of externalist account must be correct.⁹

9 For a recent version of this argument, see Alvin Goldman, “Internalism exposed,” *Journal of Philosophy*, 96 (1999), pp. 271–93.

A full assessment of this argument must obviously depend on an evaluation of the success of opposing theories in avoiding skepticism, including both coherentism and the version of foundationalism to be outlined and defended below. But again it must at least initially be conceded to possess substantial force, in light of centuries of internalist epistemological efforts whose positive results are quite inconclusive at best.

2.5 Can Internalism and Externalism be Reconciled?

We seem to be threatened with something of an impasse. Both externalism and internalism face serious intuitive objections, and arguments between the two sides often amount to little beyond more and more vehement displays of the relevant intuitions, with little real progress being made. I have come to think that a more fruitful and constructive approach must begin by re-examining the idea that externalism and internalism should be viewed as genuinely contradictory or at least contrary views, between which a choice would accordingly have to be made.

As was implicit in some of the previous discussion, the opposition between the two views is sometimes thought to pertain to the correct specification of the concept of *knowledge*: is the third condition for knowledge, the justification or warrant condition that goes beyond the requirements of belief or assent and truth, properly understood or formulated in an internalist or externalist way? I have already expressed some doubts as to the clarity and univocality of the concept of knowledge, centering mostly on the question of the specific degree of justification that knowledge requires (see the end of chapter 1). But the point that is most relevant to the present issue is that even apart from issues pertaining to the degree of justification, it is far from clear that there are not in fact two (or perhaps even more) quite distinct and largely unrelated concepts of knowledge, one (at least) predominantly externalist in character and one (at least) predominantly internalist in character, a suggestion that has been made by more than one epistemologist and that seems to me quite plausible.¹⁰ And if this were correct, of

10 For versions of such a view, see J. L. Mackie, *Problems from Locke* (Oxford: Oxford University Press, 1976), pp. 217–20; and Ernest Sosa, *Knowledge in Perspective* (Cambridge: Cambridge University Press, 1991), p. 240 (and elsewhere in the same volume).

course, externalism and internalism could be plausibly taken to be describing features of quite different concepts, rather than giving competing descriptions of one concept.

Perhaps then the issue between externalism and internalism is better formulated as an issue about the correct account of the concept of *epistemic justification*: is the correct understanding of *this* concept externalist or internalist in character? The immediate problem with this suggestion, of course, is that epistemic justification is perhaps most standardly specified as that species of justification that is required for knowledge. But a deeper problem is that even if an alternative specification can be found, the specifically epistemological notion of justification is to a significant extent a technical philosophical notion, one that is not clearly and unquestionably present in common sense. Thus the door seems open to the possibility, again advocated by some, that there may simply be different and incommensurable concepts of epistemic justification, one or more internalist and one or more externalist in character (and also perhaps some that are hybrid in character), leaving it once again unclear in what way these are competitors between which a choice has to be made. (Thus the internalist can perhaps concede that unsophisticated epistemic subjects like those discussed earlier may perhaps have either knowledge or justified belief or both in externalist senses, while still denying that this has any tendency to show that his quite different conception of epistemic justification – and perhaps also of knowledge – is thereby shown in any clear way to be mistaken.)

But the best case for the view that externalism and internalism need not be regarded as irreconcilably opposed can, I think, be made by showing that there are in fact important and clearly epistemological questions for which a predominantly externalist approach seems preferable, and others for which only an internalist approach will do. If this is so, then perhaps externalism and internalism are best regarded as complementary approaches to largely different issues, rather than as competitors.

Most of the issues that seem to suggest or even require an externalist approach fall within the confines of what Philip Kitcher has aptly labeled “the meliorative epistemological project,”¹¹ that is, the general project of assessing and improving the reliability of human cognitive efforts. Thus, for example, suppose that the issue is raised as to which of several methods of organizing or structuring scientific inquiry in

11 In his paper “The naturalists return,” *Philosophical Review*, 101 (1992), pp. 64–5.

a certain area is most likely to lead to correct results and to do so reasonably promptly and efficiently. The natural way to investigate this question would be to study many cases of research organized in the various ways in question and see how frequently and how readily cognitive success is attained. Such an investigation would be naturally conducted from a *third-person* perspective, looking at the people employing the various methods from the outside and assessing their success from that perspective. And it would be possible, though hardly essential, to formulate the results of such an investigation by saying that the more successful methods and the beliefs that they lead to are more *justified* in what would be essentially an externalist (reliabilist) sense. Thus there is clearly room in epistemology for the sorts of investigations whose results could be formulated by using an externalist conception of justification (or perhaps instead an externalist conception of knowledge).

But while such investigations are obviously legitimate and valuable, and also obviously of epistemological, though not merely epistemological, significance, there are two important and closely related features of them that have an important bearing on the main issues that concern us here. First, as already suggested, such investigations are normally and naturally conducted from a third-person standpoint, looking at the epistemic agents in question from the outside; and, second, these investigations do not and must not employ the specific methods or epistemic sources whose reliability is at issue. The persons studying the reliability of the various methods of inquiry examine the use of those methods by others, not by themselves; and they must not employ one or more of those same methods in their investigation, on pain of obvious circularity.

In contrast, the internalist approach becomes essential, I suggest, when the issue is, not the third-person question of whether someone else's beliefs are true or reliably arrived at, but instead the first-person (singular or plural) question about the truth (or reliability) of my own or our own beliefs, especially the relatively global version of this question in which it is all of a person's beliefs that are in question. In relation to this global question, no externalist approach is available to the person or persons in question without begging the very question at issue (though some other person or persons could, of course, conduct such an investigation from the outside).

One immediate upshot of this is that the "internal" of "internalism" means primarily that what is appealed to for justification must be *internal to the individual's first-person cognitive perspective*, that is, some-

thing that is unproblematically available from that perspective, not necessarily that it must be internal to his mind or person in the way that mental states are. Thus, I suggest, a person's conscious mental states play the role that they standardly do in internalist conceptions of justification, not simply because they are internal to him in the metaphysical sense of being his individual states, but rather because it has been thought that some (but not all) of the properties of such states, mainly their specific content and the attitude toward that content that they reflect, are things to which the person has a first-person access that is direct and unproblematic – a view that will be further investigated below.

The other side of this point is that, contrary to what has sometimes been supposed,¹² there is nothing at all about the fundamental rationale for internalism that automatically limits what is available for and relevant to internalist justification to facts about conscious mental states and their properties as such. If there are facts of some other sort that are directly and unproblematically available from a person's first-person cognitive perspective, then these are equally acceptable for this purpose. Thus, to take the most important such possibility, if some facts about logical and probabilistic relations among propositions can, as the rationalist holds,¹³ be directly discerned via a priori insight, then these facts would also be available from the first-person perspective in which the global epistemological issue is raised, even though these are obviously not facts about conscious mental states.

My suggestion is thus that both internalist and externalist approaches are legitimate in relation to genuine epistemological issues and hence that there is no clear reason why one has to be chosen in preference to the other. There is intellectual room for lots of different kinds of epistemological issues, including many that are naturally approached from the third-person perspective in a way that is at least largely externalist in character, together with some that are essentially internalist issues, especially relatively global issues having to do with whether one has good reasons for one's own beliefs. From this standpoint, the intuitive objections to externalism lose most of their sting (since no claim need be made that externalist justification brings with it first-person rationality). And, as already noted, the intuitive objection to internalism on the basis of unsophisticated epistemic subjects can be defused by taking it to reflect externalist conceptions of

12 E.g., by Goldman in "Internalism exposed," cited in n. 9.

13 See *IDPR* for a defense of this rationalist view.

justification and knowledge, leaving only the skepticism objection, which will be set aside for later reconsideration.

2.6 The Indispensability of Internalism

Where does this leave us? Having been reconciliatory to this extent, I still want to insist that there is a clear way in which an internalist approach, in addition to being intellectually legitimate on its own, has a fundamental kind of priority for epistemology as a whole, so that externalist views, whatever their other merits, do not constitute satisfactory responses to the general issue with which this essay is concerned: that of whether we have any good reasons to think that any of our beliefs about the world are true (and what form these reasons might take).

This is so because externalist justification simply does not speak to this global and essentially first-person issue. One way to see this is to note that if an epistemologist claims that a certain belief or set of beliefs, whether his own or someone else's, has been arrived at in a reliable way, but says this on the basis of cognitive processes of his own whose reliability is for him merely an external fact to which he has no first-person, internalist access, then the proper conclusion is merely that the belief or beliefs originally in question are reliably arrived at (and perhaps thereby are justified or constitute knowledge in externalist senses) *if* the epistemologist's own cognitive processes are reliable in the way that he believes them to be. Of course there might be a whole series of hypothetical results of this sort: cognitive process *A* is reliable if cognitive process *B* is reliable, cognitive process *B* is reliable if cognitive process *C* is reliable, and so forth. But the only apparent way to arrive at a result that is not ultimately hypothetical in this way is for the reliability of at least some processes to be establishable on the basis of what the epistemologist can know directly or immediately from his first-person, internalist epistemic perspective.

This point also has an important bearing on the problem of skepticism and the way in which it seemed to provide an intuitive argument for externalism. Our fundamental commonsense conviction, I suggest, is not that we have "knowledge" in some unspecified sense, nor that our empirical beliefs are "epistemically justified" in a sense that is loosely enough specified to admit externalism as a possible interpretation. It is rather the conviction that in general we actually do have

good reasons within our cognitive grasp for thinking that our various beliefs about the world are true: good reasons to think that there is a physical world, that it contains various kinds of familiar objects, that the world did exist in the past, that nature follows regular laws, etc. Indeed that anyone who denies such things is flying in the face of reason and good sense. If we did not have such a conviction, there would be nothing particularly implausible about skepticism and no particular reason to think that our beliefs are justified in *any* sense, including the externalist one.

There is, of course, no guarantee that this commonsense conviction is correct. But it is the attempt to make philosophical sense of it, to show how we can really have the reasons that we think we have, that provides both the main motivation and the only important criterion of success for the most fundamental sort of epistemological theorizing. And making sense of this conviction is something that an externalist view cannot in principle do, since externalist justification is not available to the subject in question, and indeed, assuming a general externalist solution to the regress problem, is available to no one. On an externalist view, though there *may* in a sense be reasons why our various beliefs are true or likely to be true, these reasons are in principle inaccessible, so that in particular no one can ever have any reason to think that they exist. This would mean that the anti-skeptical conviction just described could not even in principle be well founded (though it too, for all that we can tell, might be externally justified). And such a result by itself constitutes a very strong and intuitively implausible version of skepticism, one that externalism not only cannot answer but is indeed positively committed to.

Externalists often write misleadingly as though from a perspective in which the reasons that are unavailable to the ordinary believer are apparent to them: from which, for example, it is obvious that our perceptual beliefs about medium-sized physical objects are reliably caused and so mostly true. But in fact, if externalism is the only solution to the regress problem, *there is no such perspective available to anyone*, no perspective from which anyone ever has good reasons to think that anyone's beliefs of any sort are in fact reliably caused. Thus the externalist should speak instead of the mere *possibility* that beliefs are, in ways that are inaccessible to anyone, reliably caused; and hence of the *possibility*, which may or may not be realized, that they are, in the externalist sense, justified. But putting things in this way would, I submit, utterly destroy the appeal of externalism as a response to skepticism.

And this is why internalism is indispensable to epistemology as a whole. Though there are many other legitimate questions and issues, only an internalist approach will ultimately do when, paraphrasing Bishop Butler, “I sit down in a cool hour” and ask whether I ultimately have any good reasons for thinking that my beliefs are true or indeed that they are reliably arrived at.

In Search of Coherentism

The second main putative alternative to internalist foundationalism is a coherence theory of justification: a view according to which (1) there are no basic or foundational beliefs and (2) at least the primary basis for empirical justification is the fact that such beliefs fit together and support each other in a variety of complicated ways, thus forming a coherent system of beliefs – or perhaps more than one such system (see below). Such views have been proposed or at least suggested by a variety of authors from the classical absolute idealists (Bradley, Bosanquet, and, especially, Blanshard) through certain of the logical positivists to contemporary figures like Sellars, Rescher, Lehrer, and myself.¹ But despite this fairly impressive list of proponents and a much larger amount of critical discussion of views of this general kind, it is, rather surprisingly, still far from clear that there is a well-defined coherentist position that is even *prima facie* defensible. Accordingly, my first purpose in this chapter is to explore the general outlines of such a position, focusing mainly on those elements and ingredients that seem dialectically essential. On this basis, I will then argue that (contrary to my own earlier view²) all views of this

1 See F. H. Bradley, *Essays on Truth and Reality* (Oxford: Oxford University Press, 1914); Bernard Bosanquet, *Implication and Linear Inference* (London: Macmillan, 1920); Brand Blanshard, *The Nature of Thought* (London: Allen & Unwin, 1939); Wilfrid Sellars, "The structure of knowledge," in Hector-Neri Castañeda (ed.), *Action, Knowledge, and Reality* (Indianapolis: Bobbs-Merrill, 1975), pp. 295–347; Nicholas Rescher, *The Coherence Theory of Truth* (Oxford: Oxford University Press, 1973) and *Methodological Pragmatism* (New York: New York University Press, 1977); Keith Lehrer, *Knowledge* (Oxford: Oxford University Press, 1975) and *The Theory of Knowledge* (Boulder, Colo.: Westview Press, 1990); and *SEK*. Such a view is also frequently ascribed to W. V. O. Quine, but Quine's epistemological views are sufficiently elusive to make it difficult to be very confident about this assessment.

2 In *SEK*.

general kind are faced with obvious, overwhelming, and clearly fatal objections.

In partial contrast to the situation with externalism, where there are at least some other, broadly “naturalistic” motivations at work, the main motivation for coherentism is purely dialectical: simply the avoidance of foundationalism, with no very substantial initial plausibility attaching to coherentism itself. While it is quite plausible that coherence or something like it is *one* important ingredient in empirical justification, it is initially *very* implausible that it is the whole story, or anything at all close to the whole story. The main reasons for this assessment are both familiar and obvious: (1) The idea that epistemic justification depends solely on the internal coherence of a system of beliefs seems to entail that such justification requires no *contact* with or *input* from the world outside that system of beliefs; but how can a relation among beliefs that involves no contact with the world yield a reason for thinking that the beliefs in question correctly describe that world? (2) Since coherentist justification has to do only with the internal relations between the members of a system of beliefs, it seems possible at least in principle to invent indefinitely many alternative and conflicting such systems in a purely arbitrary way, while still making each of them entirely coherent; but it surely cannot be the case that all such systems are thereby justified in the epistemic sense of there being good reason for thinking that their component beliefs are true (and there is obviously no possible way to select among them on purely coherentist grounds). (3) In part because of these two more specific objections, there seems to be no clear reason in general to think that the coherence of a system of beliefs makes it likely that the component beliefs are true (in the realist sense of corresponding with independent reality), thus making it impossible to understand how coherence can be the basis for epistemic justification.

In fact, largely for the reasons just noted, there is probably no one who has ever seriously advocated a *pure* coherence theory of empirical justification, one in which the coherence of a set of beliefs is claimed to be by itself sufficient for justification. The historical coherentist project has rather been, in effect if not very explicitly, to supplement the appeal to coherence in a way that avoids or at least mitigates these objections, while at the same time avoiding a relapse into foundationalism. Given this loosely specified and essentially negative aim, it is perhaps not surprising that the details of the various positive coherentist and quasi-coherentist positions vary quite widely, so much so that it is far from clear that there is very much common

ground between them that can be identified as the core coherentist position.

3.1 The Main Ingredients of Coherentism

I will attempt nonetheless in the following discussion to identify and explain, in a necessarily schematic way, the main elements that are arguably essential to any coherentist position that purports to offer an account of a notion of epistemic justification that is both internalist in character and plausibly conducive to finding truth (understood in the realist way indicated earlier). There are, I will suggest, four main elements that are needed to stave off the most obvious problems and provide even the outline of a reasonably clear and specific position.³

(i) *Nonlinear Justification*

Of the three alternatives with regard to the outcome of the epistemic regress that were outlined in chapter 1, the coherentist clearly must opt for the second, the idea that the chains of justification for particular contingent, empirical beliefs circle or loop back upon themselves. Incautious advocates of coherentism have sometimes seemed to endorse the idea that such a view is acceptable if only the circles are “large enough.” But the obvious objection to circular chains of justification, to which the size of the circle seems entirely irrelevant, is that they involve circular reasoning and hence have no genuine justificatory force.

The only apparent hope for a coherentist response to this objection, stemming originally from Bosanquet,⁴ is the suggestion that the objection depends on the plausible but ultimately mistaken idea that relations of justification fundamentally involve a *linear*, asymmetrical

3 Perhaps not surprisingly, the elements that I regard as essential are realized most fully and explicitly (though not necessarily adequately) in my own former coherentist position, as developed in *SEK*. But it is my belief (which cannot be fully defended here) that any coherentist position that has even a prima facie chance of being tenable will have to involve at least a close approximation to these elements, and that at least all but the fourth can be discerned (with varying degrees of clarity) in the main historical examples of coherentism.

4 See Bosanquet, *Implication and Linear Inference*.

order of epistemic dependence among the beliefs in question. The contrary suggestion is that justification, when properly understood, is ultimately *nonlinear* or *holistic* in character, with all of the beliefs in the relevant system of beliefs standing in relations of *mutual* support, but none being epistemically prior to the others. In this way, it is alleged, any objectionable circularity is avoided. Such a view amounts in effect to making the system of beliefs in question itself the primary unit of justification, with its component beliefs being justified only derivatively, by virtue of their membership in such a system. And the property of the system, in virtue of which it is justified, is of course specified as coherence.⁵

There is an important mistake that needs to be avoided here, however, one that would make coherentism even harder to defend than it otherwise is. Virtually all historical versions of coherentism have in fact construed the idea of holism or nonlinearity in an extreme way according to which the relevant “system” is the believer’s entire body or set of beliefs, all of which would thus stand or fall together. But such a view is in fact highly implausible. In addition to aggravating a number of other problems (see below), it would force the coherentist to say, most implausibly, that the justification of a belief in one area (e.g., nuclear physics) could be undermined by a serious incoherence in the person’s beliefs in a completely unrelated area (e.g., art history). And though historical coherentists have often been willing to advocate the view that all of a person’s beliefs are in fact justificatorily inter-related in this way, their only very serious argument for such a claim has been an appeal to the very holism that is in question.

But on careful reflection, such an extreme holism is both unnecessary and quite unmotivated in relation to the main coherentist view.⁶ What coherentism clearly requires is that the primary units of justification be groups of interrelated beliefs rather than individual beliefs: groups that are large enough to either satisfy or non-trivially fail to satisfy the various aspects of the concept of coherence (see the discussion of these below). Although we have not yet considered the question of just why, according to the coherentist, the coherence of a group of beliefs should be thought to be an indication of the truth of its members, there is no reason apparent so far (and in fact none will emerge) as to why a defensible answer to this question (assuming that

5 For a further elaboration of this idea, see *SEK*, pp. 89–93.

6 Contrary to what was assumed, more or less without discussion and with no very clear rationale, in *SEK*.

there is one) would not apply just as well to a coherent set of beliefs that is much smaller than a particular person's complete body of beliefs. And there is also no clear reason why the size of the relevant groups of beliefs upon which justification depends might not in fact vary fairly widely, depending on just what specific issue of justification is being considered, with some issues being relatively "local" in character and others more "global."⁷

(ii) *The Concept of Coherence*

But what exactly is coherence? The second component of any serious coherence theory must be some relatively specific account of this relation among beliefs. Those who employ the notion agree that coherence is a matter of how the beliefs in a system of beliefs "fit together" or "dovetail" with each other, so as to constitute one unified and tightly structured cognitive whole. And it is also clear that this fitting together or dovetailing will depend in turn on the various more specific logical, inferential, and explanatory relations that exist among the component beliefs of the system in question. But spelling out the details of this idea, particularly in a way that would allow the reasonably precise *comparative* assessments of coherence that a serious account of epistemic justification seems to require, turns out, not really very surprisingly, to be extremely difficult, partly because of the complexity of the overall picture and partly because such an account will depend on more specific accounts of a large number of narrower and still on the whole inadequately understood topics, such as induction, confirmation, probability, explanation, and various issues in logic (particularly those connected with "relevance logic").

Some points are, however, at least relatively clear. First, any conception of coherence that is even *prima facie* adequate as a basis for epistemic justification must, contrary to some suggestions, require more than mere logical consistency among the beliefs of the system. Indeed, in light of both general human logical fallibility and more spe-

7 It is common to use the word "system" to refer to the particular group of beliefs in relation to which the issue of coherence is raised, and I will adopt this practice in what follows. In these terms, the view just argued for in the text is that there is no reason why a given person's total body of beliefs might not fall into many largely or entirely unrelated "systems" of this sort – which seems in fact pretty clearly to always be the case. And of course if the term "system" is used in this way, then it must not also be used to refer to the person's entire body of beliefs.

cific problems such as the paradox of the preface (pertaining to the case in which an author prefaces a complicated discussion by saying that he is sure that some of the claims in it are false), it seems a mistake to view logical consistency as even a *necessary* condition for the degree of coherence required for a significant degree of justification. Second, coherence requires a high degree of inferential interconnectedness between the beliefs in the relevant system beliefs, involving relations of necessitation, both strictly logical and otherwise, and also probabilistic connections of various kinds. One important aspect of this is what might be called *probabilistic consistency*: the minimizing of relations between beliefs in the system in virtue of which some are highly *unlikely* to be true in relation to others.⁸ Third, while some recent positions have emphasized *explanatory* relations as the basis for coherence,⁹ it seems reasonably clear that this cannot be the whole story. The coherence of a system of beliefs is surely enhanced to the extent that some parts of the system are explained by others, thus reducing the degree to which the beliefs of the system portray unexplained anomalies. But not all relevant sorts of inferential connections can be plausibly construed as explanatory in character.¹⁰

As this still very sketchy account suggests, giving a precise account of the nature of coherence remains a largely unsolved or perhaps even unsolvable problem, making it reasonable to ask why this deficiency isn't in itself a sufficient basis for dismissing such theories. A partial response to this objection¹¹ is that difficulties in this area cannot yield anything like a decisive argument against coherence theories and in favor of their foundationalist rivals, because the concept of coherence, or something so similar to it as to be capable of playing essentially the same role and to involve the same problems, is also an indispensable ingredient in virtually all internalist foundationalist theories: an appeal to coherence is seemingly required as at least part of the account of the relation between the basic or foundational beliefs and other non-foundational or "superstructure" beliefs, in virtue of which the latter

8 Though this requirement must not be made too strong, since it seems clear that we are often justified in believing that antecedently improbable events have indeed taken place.

9 This is true, e.g., of Gilbert Harman's coherentist view in *Thought* (Princeton: Princeton University Press, 1973); and is at least strongly suggested by Sellars, in such papers as "Givenness and explanatory coherence," *Journal of Philosophy*, 70 (1973), pp. 612–24.

10 For some further discussion of the concept of coherence, see *SEK*, pp. 93–101.

11 Offered in *SEK*, p. 94.

are justified in relation to the former.¹² For this reason, giving an adequate account of coherence should perhaps not be regarded as exclusively the responsibility of coherentists, despite the more central role that the concept plays in the coherentist position. But while this point is defensible as a dialectical response to attempts by foundationalists to use this problem as an objection to coherentism, it does not alter the fact that without further clarification of its central concept, coherentism remains at best a kind of promissory note, rather than a developed position.

(iii) *A Coherentist Conception of Observation*

Perhaps the most obvious objection to coherentism is posed by the seemingly obvious fact, one that even coherentists have rarely been willing to question or challenge explicitly, that *sense perception* or *sensory observation* plays in some way a central role in empirical justification. The most natural accounts of perception or observation are pretty obviously foundationalist in character (though not necessarily with the perceptual or observational beliefs being themselves basic or foundational). Having rejected foundationalism, any coherence theory that even hopes to be viable must thus attempt to explain how the intuitively central role of perception or observation can be understood and accommodated in a non-foundationalist way.

Here the only real possibility, I think, is the suggestion that while observational beliefs are indeed *arrived at* non-inferentially as a *causal* result of sensory experience, rather than being arrived at via any sort of inference, this does not account for their *justification* (once foundationalist accounts are set aside as untenable). And thus a possible view for the coherentist to hold is that the justification of non-inferential observational beliefs still depends essentially on their coherence with a background system of beliefs. In developing such a view, however, it is crucial that the justification in question still *also* depend in some essential way on the fact that the belief was caused in the way that it was, as a result of something like sense perception, since justification that depended *only* on the coherence of the belief's propositional content with a background system of beliefs would mean that any other way of

12 See, e.g., Lewis, in *An Analysis of Knowledge and Valuation*, ch. 11 (who uses the term "congruence"); and Chisholm, *Theory of Knowledge*, 2nd edn. (Englewood Cliffs, NJ: Prentice-Hall, 1977), ch. 4 (who speaks of "concurrence").

arriving at such beliefs, including simply inventing them arbitrarily, would work just as well. This would make the observational status of the belief justificational and so epistemologically irrelevant – a result that is both deeply implausible from an intuitive standpoint and that would also prevent the appeal to observation from playing any important role in answering the other objections to coherentism (see further below).

My own previous version of this general approach¹³ attempts to provide for observational input by appealing to the idea of a *cognitively spontaneous belief*: one that simply “strikes” the observer in an involuntary, coercive, non-inferential way, rather than arising as a product of any sort of inference or other discursive process, whether explicit or implicit. The suggestion is that ordinary observational beliefs that are a causal result of sensory experience in fact have this sort of status: if I am looking in the right direction with my eyes open in the daytime, I simply find myself thinking willy-nilly, e.g., that there is a large green tree over there.

That a belief is cognitively spontaneous, however, says nothing so far, from a coherentist standpoint, about how or even whether it is justified. Indeed, there is no reason to think that all or necessarily even most cognitively spontaneous beliefs *are* justified, since the category would include hunches and irrational spontaneous convictions, as well as beliefs resulting from perception or observation. The further suggestion is then that certain kinds of cognitively spontaneous beliefs can arguably be justified from within the person’s body of beliefs, by appeal to: (1) the fact of their spontaneous occurrence; (2) the relevant background conditions; and (3) the apparently successful track record with respect to frequency of truth under those conditions of spontaneous beliefs of those specific kinds (identified by such things as their general subject matter, their apparent mode of sensory production (as reflected in the distinctive content of the belief), and concomitant factors of various kinds) – all this being assessed, it is claimed, from within the person’s body of beliefs by appeal to coherence.¹⁴

In this way, it seems at least initially possible for there to be a justifying reason for such a cognitively spontaneous belief that appeals to

13 See *SEK*, 6.

14 Something like this view seems implicit in Blanshard’s talk of “beliefs about the technique of acquiring beliefs” and in Sellars’s talk of “language-entry transitions”; see Blanshard, *The Nature of Thought*, pp. 285–6, and Wilfrid Sellars, “Some reflections on language games,” reprinted in his *Science, Perception and Reality* (London: Routledge & Kegan Paul, 1963), pp. 321–58. It seems also to be hinted at, though less explicitly, in other coherentist views.

its status as cognitively spontaneous and thus non-inferentially arrived at in the way that is characteristic of observation, but still does so in a way that makes the resulting justification still dependent essentially on coherence: specifically, on the fact that the general thesis that a belief of this kind and produced in this way is likely to be true itself coheres with a relevant background system of beliefs. Such a cognitively spontaneous belief would thus be *arrived at* non-inferentially, but still allegedly *justified* by appeal to inference relations and coherence, albeit not by coherence alone. Beliefs that are justified in this distinctive way would then allegedly constitute a kind of observational input from the extra-conceptual world that is still recognizably coherent in its justification.

But this is still not enough for a successful coherentist account of observation to succeed. There are at least two further requirements that such an account would have to satisfy, requirements that will be only briefly listed here (though each of them could be the subject of a much longer discussion). First, the other beliefs needed to give a justifying reason for a particular observational belief must of course *themselves* be justified in some way that does not amount to a relapse into foundationalism. Included here will be: (1) beliefs about the specific conditions under which the cognitively spontaneous belief in question occurred; (2) the belief that cognitively spontaneous beliefs of the specific kind in question are likely, under those specific conditions, to be true; and (3) further beliefs pertaining to that specific belief and its occurrence, including the belief that it was indeed cognitively spontaneous, together with further beliefs that are relevant to specifying the particular kind of belief in question. The justification for (1) will in general presumably have to include other observational beliefs, themselves justified in the same general fashion, so that any case of justified observation will normally or perhaps always involve a set of mutually supporting observations. The justification for (2) will presumably appeal inductively to other cases of correct observation – as judged from within the person's body of beliefs – as well as to more theoretical reasons for thinking that beliefs of the kind in question are generally produced in a reliable way. The justification for (3) will presumably appeal to introspective beliefs, themselves constituting a species of observation, and ultimately to the believer's grasp of the relevant parts of his overall body of beliefs – the status of which is the focus of the fourth of the main elements, to be discussed next.¹⁵ None

15 For more detail, see *SEK*, pp. 124–38.

of this is unproblematic, and attempting to spell it out further seems to generate new problems and difficulties in a way that offers no real basis for hoping that a stable resting place will eventually be reached.

Second, the bare *possibility* of coherentist observation is pretty obviously insufficient to accommodate the role that observation intuitively seems to play in our cognitive lives and to provide for genuine cognitive input from the non-conceptual world. Our intuitive conviction is that sensory observation is not only possible but pervasive and that an appeal to observational evidence, whether direct or indirect, is essential for the justification of at least virtually all contingent beliefs about the world. Thus a coherence theory that does not do violence to these intuitions must *require* and not just allow as a (perhaps unrealized) possibility that a substantial observational element be involved in the justification of any system that includes such contingent beliefs. It must, that is, impose something akin to what I have elsewhere called the “observation requirement,”¹⁶ to the effect that any justified system of putatively empirical beliefs must contain a significant proportion of cognitively spontaneous beliefs that are themselves, on coherentist grounds, likely to be true. Such a requirement seems alarmingly vague, and there is, as far as I can see, no very good way to make it more precise. But this problem, though serious enough, is still rather minor in comparison to the others that a coherentist view must face.

(iv) *The Doxastic Presumption*

Like the other elements already discussed, the final one to be considered here is also in effect a response to an objection that threatens to derail the coherentist position before it even gets off the ground. If it is by appeal to coherence with the believer’s system of beliefs that all issues of empirical justification are to be decided, then an *internalist* coherence theory seemingly requires that the believer have an adequate *and justified* grasp or representation of the relevant system of beliefs, since it is with respect to this system that issues of coherence and so of justification are allegedly to be decided. Such a grasp would presumably have to take the form of a set of meta-beliefs (or perhaps one comprehensive meta-belief) specifying the contents of the relevant system of beliefs. And the glaring difficulty is then that the coheren-

16 See *SEK*, pp. 140–3. But the formulation there is, like much else in that book, distorted by the extreme holism assumed there.

tist view also seems to preclude there being any way in general for such meta-beliefs about the contents of the believer's system of beliefs to themselves be justified.

Meta-beliefs specifying the contents of a person's various systems of belief are themselves obviously contingent and presumably empirical in character, and so must, according to a coherentist view, themselves be justified by appeal to coherence with some appropriate system of beliefs, presumably the system reflecting the person's putative introspective awareness of his own states of mind (for what other system of beliefs would be relevant?). Yet any such account of the justification of these meta-beliefs seems to be inevitably circular or question-begging. What is at issue, after all, putting the issue in the first person, is precisely whether the claim that I actually have certain specific beliefs, including various beliefs about my own states of mind, is correct. To deal with this issue by appeal to the alleged coherence of that claim with the system of beliefs whose members describe my own states of mind would be to take for granted the correctness of a crucial part of the very collection of beliefs whose justification is at issue. In other words, if what is at issue is which beliefs I in fact have, then no appeal to coherence with any system of my beliefs can hope to deal with that issue in a non-question-begging way.¹⁷ Though most coherentists have (somewhat surprisingly) failed to notice this problem or at least to explicitly acknowledge it, it is still pretty obviously one that a viable coherentist position would somehow have to deal with.

In my own earlier version of coherentism, I appeal at this point to what I call the "doxastic presumption." And although this move now strikes me as pretty desperate (a view that many critics no doubt arrived at much more swiftly than I did), I still know of no better way for a coherentist to handle the issue in question. The idea is to mitigate the foregoing objection by treating the comprehensive meta-belief concerning the contents of the person's whole body of beliefs as an *unjustified* hypothesis in relation to which issues of justification are *conditionally* assessed, yielding results of the general form: *if* my representation of my whole body of beliefs is correct, then such-and-such a particular belief is justified in the sense of being likely to be true. The meta-belief in question is still supposed to be a product of introspection. But what this move acknowledges is that there is no apparent way

17 Once excessive holism is abandoned, the formulation of this problem becomes more complicated, but not in any way that affects its ultimate seriousness; compare the discussion in the text with *SEK*, pp. 101–2.

within a coherentist position for this most fundamental result of introspection to itself be justified.¹⁸

This completes my enumeration of the main elements that any coherentist position must arguably include. All of them are responses to deep-seated problems, and this highlights again the dialectically defensive and reactive posture that seems inevitable for a coherence theory. It is, moreover, more than a little uncertain that any of these elements is really adequate to deal with the specific problem that motivates it. But there is worse to come, as we will see in the next section.

3.2 The Standard Objections to Coherentism

The three historically most standard objections to coherentism have already been briefly enumerated above. There is, first, the so-called “isolation problem” or “input objection,” which claims that an account of justification that depends entirely on coherence will have the absurd consequence that contingent, seemingly empirical beliefs might be justified in the absence of any sort of informational input from the extra-conceptual world that they attempt to describe. This would seem to mean in turn that the truth of those beliefs, if they happened to be true, could only be an accident in relation to that world, and thus that there could be no genuine reason to think that they are true and so no epistemic justification.

One of the primary motivations for the coherentist account of observation is to meet this objection by showing how observational beliefs that are causally generated by the world might nonetheless be given a coherentist justification, and how a requirement for such an observational ingredient can be made a necessary condition for empirical justification within a coherentist framework.¹⁹ In this way, the input objection might seem, at least *prima facie*, to be met – though we will have to reconsider below whether this attempted answer really succeeds, even apart from the problems of detail that were noticed briefly above.

The second standard objection is what is usually referred to as the alternative coherent systems objection: Even given a relatively demanding conception of coherence, there will still be indefinitely many different possible systems of beliefs in relation to any given

18 For more extended discussion, see *SEK*, pp. 101–6 and 147–8.

19 See *SEK*, pp. 139–40.

subject area, each as internally coherent as the others. Thus, the members of each of these systems will seemingly be on a par as regards justification according to a coherentist view. And this is surely an absurd result, especially since *any* belief relevant to that subject area that is not somehow internally incoherent in itself will apparently belong to one or more such systems – so that the coherentist must apparently say that any such belief is as justified as any other.

The best attempted response to this objection that seems to be available to the coherentist also depends crucially on the coherentist notion of observation. If the existence of a substantial observational component is made a necessary condition for empirical justification, as suggested above, then, it is claimed, there is no longer any reason to think that such alternative systems can be freely invented in the way that the objection claims, and hence no longer any obvious reason why they should be thought to be genuinely possible. The point here is that there is no reason to think that the cognitively spontaneous beliefs that are judged to be likely to be true in relation to an arbitrarily invented system of beliefs will in fact cohere with that system over time, and thus no reason to think that such a system will *remain* coherent.²⁰ Here too the coherentist perhaps has at least the gist of a *prima facie* adequate response to the objection in question, though this issue too will have to be reconsidered below.

The third and most fundamental of the standard objections is in effect a challenge to the coherentist to give a reason, a *meta-justification*, for thinking that adopting beliefs on the basis of the coherentist view of justification is in fact likely to lead to believing the truth, which is obviously essential if coherence is to be a genuine basis for epistemic justification.

One historically prominent response to this problem is the adoption of a coherence theory of *truth* as well as justification, leading to some version of metaphysical idealism. Here I will simply assume that no such view is acceptable, denying as it does the intuitively obvious fact that a world independent of minds and their beliefs genuinely exists and that it is this world that our beliefs attempt to describe.

The only apparent alternative is for the coherentist to offer an argument of some sort from the empirical premise that a given system of beliefs is coherent (and satisfies the requirement of observation, as roughly formulated earlier) to the conclusion that the component beliefs of the system are likely, to an appropriate degree, to be true. Such an argument would apparently have to be itself *a priori* in char-

20 For more discussion, see *SEK*, pp. 143–6.

acter, since any sort of further empirical ingredient (beyond the initial premise that the system is and remains coherent) would on a coherentist view have to be itself justified by appeal to coherence, thereby rendering the argument viciously circular.²¹ I have elsewhere attempted to sketch the outlines of such an a priori “meta-justificatory argument,”²² centering on the idea that only approximate truth could *explain* the fact of long-run coherence (given the satisfaction of the requirement of observation), but the details of that attempt cannot be considered here. Whether or not it is defeated by other problems, the coherentist version of such an argument now seems to me to be decisively undercut by the more specific objections to coherentism that are discussed in the following section.²³

3.3 Still Further Objections

Coherentism emerges from the foregoing discussion as at best an extremely shaky and problematic position, dialectically on the defensive from the very beginning and afflicted with a multitude of problems and objections that can seemingly at best be only staved off, but rarely if ever decisively answered. Given the apparent strength of the arguments against foundationalism (together with the inadequacy, in

21 It is sometimes suggested that there is nothing objectionably circular about a coherence theory appealing to coherence to justify the claim that coherence is truth conducive. And indeed that *any* theory of justification must make an analogous appeal to its fundamental standard (on pain of abandoning its claim to be a comprehensive account of justification). One way to see that there is something wrong with this response is to note that such a self-invoking justification of the claim of truth conduciveness is equally available for many obviously unsatisfactory views of justification (consider, as a simple example, the view that belief by me is the standard of justification, where I also believe that all of my beliefs are true). The proper conclusion, I think, is that only a view that appeals at some point to a direct insight into or grasp of truth can avoid this problem – which is what paradigmatically foundationalist views, like the one offered in chapters 4 and 5 below, attempt to do. A coherentist view could appeal to such a direct insight into truth only by adopting a coherence theory of truth of the sort that is briefly mentioned in the text. (I am indebted to John Greco for calling my attention to this issue.)

22 *SEK*, ch. 8.

23 Rescher attempts to give a *pragmatic* argument that the practical success that results from the employment of the coherent system makes it likely that the beliefs of the system are at least approximately true. See Rescher, *Methodological Pragmatism*. But the relation of this argument to the coherence of the system is less than clear, and in any case the obviously empirical claim of actual pragmatic success would, for a coherentist, have to be itself justified by appeal to coherence, again making the argument viciously circular.

my judgment, of externalism as a solution of the central epistemological problem that concerns us here), a development and attempted defense of coherentism seemed a project worth pursuing, albeit one that clearly faced pretty long odds from the outset. Now, however, it seems to me clearly time to concede that attempts to develop the coherentist suggestion into a viable epistemology have not succeeded and almost certainly cannot succeed. This result is perhaps already obvious enough in light of the foregoing discussion, but I will support it a bit more by discussing three further objections to coherentism that seem by themselves to be pretty clearly fatal. These are not unrelated to the problems and objections already discussed, but they are somewhat narrower and more focused.

First. Perhaps the most obvious objection pertains to the doxastic presumption. We have already taken note of the problem of access to the beliefs in relation to which coherence is to be assessed, and of the Doxastic Presumption as the only response that is apparently available. But it must be admitted that the result of this move, as many have taken pains to point out, is a very deep and troubling version of skepticism, albeit not perhaps quite the version that historical responses to skepticism have been mainly concerned with: a skepticism according to which no one has *any* justification *simpliciter* for *any* empirical belief, but only at best for the conditional claim that *if* a certain unjustified and unjustifiable presumption is correct, then various empirical beliefs are likely to be true.²⁴ And while it still seems to me that some forms of skepticism are unavoidable and will simply have to be lived with, I find it more and more implausible (partly for reasons implicit in the discussion in the next chapter) to suppose that this is one of them. Nor, as far as I have been able to see, is there any other response to the problem of access to one's own beliefs that is compatible with coherentism. Indeed, it seems clear that would-be coherentists, myself included, have succumbed at this point to a kind of philosophical mistake that is very prevalent and also perhaps the most difficult to avoid: continuing to rely implicitly and inadvertently upon something, in this case the direct, foundational grasp of one's own states of mind, that they have officially repudiated.

24 I have sometimes claimed that *if* this presumption is in fact true, then the various empirical beliefs are after all justified, since there is then a reason why they are likely to be true. This, however, is simply a mistake, since though such a reason would, as it were, exist in the abstract, it would still be inaccessible to the believer, or to anyone else. (It was Richard Fumerton's patient but firm insistence that finally led me to see this pretty obvious point.)

Second. A somewhat less obvious but at least equally serious objection pertains to the coherentist's attempted account of observational input. It still seems to me that something like the account sketched above perhaps succeeds in showing how there could be a kind of input that is justified in at least a partially coherentist way, whether or not it really deserved to be regarded as *sensory* observation. But what is extremely doubtful is that such input can be effectively *recognized* or *identified* as such in an internalistically acceptable way. In particular, the attempt, discussed above, to make the existence of such input a requirement for empirical justification does not seem to succeed, so long as this requirement is construed in such a way that its satisfaction is itself internalistically recognizable.²⁵ The reason for this is that the alternative coherent systems objection, which this account of input is aimed in part to meet, recurs all over again: as long as it is only specified *within the person's body of beliefs* that cognitively spontaneous beliefs occur and that the observation requirement is satisfied, there will be indefinitely many other competing bodies of belief containing analogous specifications but characterizing the world in more or less any arbitrarily chosen way one likes. Of course, such bodies of belief will not in general *genuinely* receive such input, but the fact that they fail to do so will not be discernible on the basis of the contents of the beliefs themselves – which is all that a coherence theory, even given the doxastic presumption, can legitimately appeal to.²⁶

In my earlier discussion of this issue, I attempted to meet this objection by saying that in order to be genuinely justified, such a body of beliefs must be actually believed by someone, as opposed to being a merely possible set of beliefs,²⁷ with the suggestion being, as we saw above, that the systems of belief that are components of an arbitrarily constructed body of beliefs would not remain coherent in actual use, and in particular that the allegedly cognitively spontaneous beliefs specified by such an arbitrarily constructed body of beliefs would not genuinely be found to occur in a cognitively spontaneous way. Unfor-

25 I am assuming here that the satisfaction of this requirement is part of the overall coherentist reason or justification for *any* empirical claim, i.e., that one who has no access to the fact that this requirement is satisfied fails to really possess a reason for thinking that a belief that satisfies the rest of the coherentist account is likely to be true. For only a justification that includes the satisfaction of this requirement can withstand the input and alternative coherent systems objections. (I am indebted to John Greco for pointing out the need to be more explicit on this point.)

26 For a version of this objection, see Susan Haack, *Evidence and Inquiry: Towards Reconstruction in Epistemology* (Oxford: Blackwell, 1993), pp. 57–60.

27 SEK, pp. 149–50.

tunately, however, this sort of response seems to succeed to the extent that it does only because it tacitly appeals to a direct awareness of one's own actual beliefs and their occurrence that is not legitimately available to a coherentist. As long as the occurrence of cognitively spontaneous beliefs and the satisfaction of the requirement of observation is assessed only by appeal to the *coherence* with the rest of the body of beliefs whose *content* stipulates that these conditions are satisfied (which is all that the coherentist legitimately has to go on), the objection stands. (Here again we see the way in which whatever plausibility the coherentist view might seem to have depends on an unacknowledged appeal to a direct or foundational awareness of at least one's own beliefs and their spontaneous occurrence.)

I note in passing that it would of course be possible to avoid this second objection by construing the observation requirement as an *externalist* requirement, one whose satisfaction does not need to be internally assessable. But externalism, for reasons discussed in the previous chapter, seems to me equally unacceptable as the basis for a fundamental epistemological account; and in any case, as already suggested above, such an externalist version of coherentism would have little if any dialectical point, since if externalism were otherwise acceptable, a foundationalist version would be much more straightforward and unproblematic.

Third. A final objection pertains to a topic that has not emerged at all explicitly in our discussion of coherentism so far, but which is nonetheless vitally related to the tenability of the view. As was implicit in the discussion of the alternative coherent systems objection in section 3.2, a coherence theory must appeal, not just to coherence at a moment, but to sustained coherence over a period of time and indeed over at least a relatively long run. It would take time for the coherence of an arbitrarily invented system to be destroyed by new observations; and in attempting to argue for the connection between coherence and truth, it is only long-run or at least relatively sustained coherence that might seem to demand truth as an explanation.²⁸ But then the issue arises of how, according to a coherence theory, the *memory* beliefs upon which any access to the fact of *sustained* coherence would have to rely are themselves justified. Some philosophers have offered coherence theories of the justification of memory beliefs. But, whatever the other merits of such an account might be, it seems clearly to result in vicious circularity if the only reason for thinking that coherentist jus-

28 See *SEK*, ch. 8.

tification is conducive to truth, and hence that the memory beliefs in particular are true, relies on the existence of coherence over time and so on the truth of some of those very memory beliefs themselves. The upshot is that there is no non-circular way for a coherentist to appeal to sustained or long-run coherence, making it even more difficult – or, I think, impossible – to respond to the alternative coherent systems objection or to argue for the connection between coherence and truth.²⁹

Even the foregoing litany does not really exhaust or probably even come very close to exhausting the full range of objections to coherentism, but it surely suffices to make clear beyond any serious doubt the untenability of the central coherentist view.

29 This problem was first called to my attention by the discussion in Alvin Plantinga, *Warrant: The Current Debate* (London: Oxford University Press, 1993), p. 107.

Back to Foundationalism

The conclusion that follows from the previous two chapters, when taken together, is that neither of the two main dialectical alternatives to traditional internalist foundationalism is acceptable as the most basic account of the justification of empirical beliefs. Coherentism is beset with a large number of problems, many of which seem clearly fatal; and externalism, while tenable and even valuable in relation to certain epistemological issues, is incapable in principle of speaking to the fundamental issue with which we began: that of whether we have good reasons for thinking that our beliefs about the world are true. Moreover, while there are still other sorts of epistemological views that have not been explicitly considered here, none of them seems to me to do any better in relation to this central issue or speaks at all clearly to the dialectic that grows out of the regress problem. Thus there seems to be a strong motive to reconsider traditional internalist foundationalism – and more specifically the intuitively compelling but dialectically difficult idea that empirical justification rests finally on a foundation of beliefs about the content of *experience*, especially sensory experience.

Here the initial stumbling block is the dilemma, originally posed by Sellars, that was discussed in chapter 1. To reiterate briefly, the problem concerns the nature of the believer's grasp or apprehension of experience or experiential content. If, on the one hand, this grasp or apprehension is construed as propositional or judgmental in character, as having as its content the conceptual claim or thesis that the experience is of a certain specific sort, then it becomes relatively easy to see how it could, if itself justified, provide justification for a belief with that same content (or perhaps more likely a more abstract content that is suitably related); but quite difficult to see why it does not itself require some sort of justification: some reason for thinking that the

proposition in question is true. If, on the other hand, the grasp or apprehension of experience is construed as non-propositional or non-judgmental in character, as involving no conceptual claim or thesis concerning the character of experience, then it is hard to see why any further justification would be required, because there is apparently nothing to justify; but also difficult to see how a grasp or apprehension of this sort can constitute a reason for thinking that any propositional belief, specifically the belief that the experience is of a certain specific sort, is true. Without some solution to this problem, the initially plausible appeal to experience or experiential content seems to be derailed before it can even really get started.

I now believe that this problem can be solved, indeed that the solution is obvious in that peculiar way in which a philosophical idea that is difficult to see when viewed through the lens of a persuasive though ultimately misconceived dialectical conception can become obvious when brought into proper focus. The core of the present chapter will be an account of this solution and of the view of foundational justification that results. I will begin by considering, in the next two sections, a somewhat tangential but in some ways more easily accessible case: that of the justification of a second-order belief (a *meta-belief*) about the existence and content of an occurrent, conscious first-order belief or thought. The following sections will then apply the lessons thus learned to the main issue of the justification of foundational beliefs about the content of sensory experience.

4.1 Conscious Thought and Constitutive Awareness of Content

It will help to begin with an example. As I work on this chapter, I believe reflectively that I am having various occurrent (as opposed to merely dispositional) beliefs or assertive thoughts about foundationalism and its problems. For example, I believe that I am presently having the occurrent belief or thought that *foundationalism is more defensible than most philosophers think*.¹ This is a meta-belief about the existence of a certain first-order belief; its content is roughly the claim that *I believe that foundationalism is more defensible than most philosophers think*. What then is my justification (if any) for this second-order meta-belief?

1 I use italics to highlight the content of a belief.

The overwhelmingly natural way to answer this question, which is also the one that I want to elaborate and defend here, is to appeal to the conscious experience involved in having the first-order occurrent belief or thought in question. But it is crucial for present purposes that the nature and status of this experience be understood in the right way.

Clearly one sort of experience that I may have of such a belief is the sort that is standardly referred to as *apperceptive*: I may be aware of that specific belief by virtue of having a *higher-order* state of awareness that has the first-order belief as its *object*, but which is metaphysically separate and distinct from it. The content of this higher-order, in this case second-order, awareness would be the proposition that *I am presently believing or thinking that foundationalism is more defensible than most philosophers think*, a proposition clearly distinct from the proposition that is the content of the original belief and which is essentially the same as the content of the meta-belief. Indeed, this second-order state of awareness seems simply to be an occurrent version of the meta-belief with whose justification we are concerned. To appeal to this sort of experience to justify the meta-belief would then be entirely circular and unhelpful.

But is this second-order awareness the only relevant sort of experience that is available? In fact, for reasons that will emerge more fully below, it seems to me a fundamental mistake to think that my primary conscious awareness of my occurrent belief or thought and its content is of this apperceptive sort, which essentially depends on the existence of a second state of mind. My suggestion is instead that an essential and intrinsic aspect of having any occurrent belief just *is* being consciously aware of the two correlative aspects of its content: first, its propositional content, in this case the proposition that *foundationalism is much more defensible than most philosophers think*; and, second, the assertory rather than, e.g., questioning or doubting character of one's entertaining of that content. These two awarenesses (or rather, more plausibly, two aspects of one awareness) are, I am suggesting, not in any way apperceptive or reflective in character: they do not require or involve a distinct second-order mental act with the propositional content that I have the belief in question. Instead, they are (at least partly) *constitutive* of the *first-level* state of occurrent belief or thought itself in that they are what make it the very occurrent belief that it is, rather than some other occurrent belief or a different sort of state altogether. The key points here are simply that occurrent belief or thought is, after all, *itself* a conscious state, not merely a state of which one can be conscious via a second, independent state; and that what

one is primarily conscious of in having such a belief is precisely its propositional and assertive content. Not to be consciously aware of that specific content would be not to have that specific conscious, occurrent belief at all.

It is this account of the primary experiential aspect of occurrent belief that seems to me to allow an escape between the horns of the anti-foundationalist dilemma. The crucial point is that the most fundamental experience involved in having an occurrent belief is *neither* a second-order apperceptive or reflective awareness having the judgmental content that a certain sort of belief state has occurred *nor* a purely non-cognitive awareness that involves no specific awareness of the character of the belief and its content. Instead it is an intrinsic and constitutive awareness of the specific propositional and assertive content of the belief: an awareness that is in part an awareness *of* a certain propositional content, but is not an assertive, higher-order awareness *that* such a proposition is believed.²

Because of its non-apperceptive, constituent character, this “built-in” awareness of content, as it might be described, neither requires any justification itself, nor for that matter even admits of any. The first-order belief itself does, of course, require justification: I need a reason, one that I am presently in fact in the process of explaining, for thinking that it is true that *foundationalism is more defensible than most philosophers think*. But there is no comparable issue of justification that arises for the intrinsic awareness of this content that I have simply by virtue of having the first-order belief in an occurrent, conscious way (as long as this awareness is not confused with the second-order, apper-

2 I interpret this as “going between the horns” of the dilemma, because I am construing the horns as embodying the development and elaboration indicated in the earlier discussion and just summarized in the text, according to which the propositional horn involves a higher-order propositional and judgmental awareness that a state of the specified sort occurs. Though the constitutive awareness of the content of the belief of course involves the proposition that figures in that content, it is not a judgmental awareness *that* I have a belief with the content in question, that a certain second-order proposition is true, and so, as explained further in the text, does not raise any issue of justification. But one could instead interpret the present argument as showing that the conceptual and propositional side of the propositional/non-propositional dichotomy is not necessarily incompatible with foundationalism after all, because it need not involve such a judgmental awareness; this would amount to “grasping one of the horns of the dilemma” rather than going between them. (On this latter interpretation, the analogous possibility for sensory experience, discussed below, would show that the non-propositional horn also includes a possibility that is compatible with foundationalism.) I am grateful to Matthias Steup for helping me to see this alternative way of viewing the relation of the constitutive awareness of content to the Sellarsian dilemma.

ceptive one). Indeed, such a non-apperceptive, constituent awareness of content might be said to be strictly *infallible* in something like the way that foundationalist views have traditionally claimed (but which most have long since abandoned). Since it is in virtue of this constitutive or “built-in” awareness of content that the belief is the particular belief that it is with the specific content that it has, rather than some other belief or some other sort of state, there is apparently no way in which this awareness of content could be mistaken – simply because there is no independent fact or situation for it to be mistaken about.

Given this result, the idea is then that it is by appeal to this intrinsic, constitutive awareness of the propositional and assertive content of the first-level belief or thought that the second-level meta-belief can be justified. Such a constitutive awareness of content seems obviously enough to constitute in and by itself, at least if other things are equal, a reason for thinking that the second-level belief that I have an occurrent belief with that very content is true (or, perhaps more realistically, for thinking that a meta-belief that gives a less detailed, more abstract description of the first-level content, but one that the actual, more specific content falls under, is true). The way in which this works, elaborated further below, is that the meta-belief is a *description* of the very content involved in the constitutive awareness of content, so that by consciously having that constitutive awareness, I am in an ideal position to judge whether or not this description is true.

In this way, such a meta-belief can have the epistemic status required by foundationalism: it can be justified in the sense of there being a clear and internally accessible reason for thinking that it is true, but where the reason in question does not itself depend on any appeal to a further belief that would itself be in need of justification – though we now see that it is the first-level constitutive or “built-in” awareness of content, rather than the meta-belief that it justifies, that turns out to be the ultimate source of justification.

The infallibility that pertains to the “built-in” awareness does not, however, extend to the apperceptive meta-belief that it justifies. It would still be at least possible to apperceptively misapprehend one’s own belief, i.e., to have a second-level belief that does not accurately reflect the content contained in the constitutive or “built-in” awareness constitutive of the first-level belief. Such a mistake might result from mere inattention, from the complexity or obscurity of the belief content itself, or from some further problem or distraction. But unless there is, in a particular case, some special reason to think that such a misapprehension has occurred or that the chances of one having occurred

are substantial, this possibility of error does not seem to prevent the second-level meta-belief from being adequately justifiable by appeal to the first-level constituent awareness. This is just to say that while such justification is defeasible in various ways, it is *prima facie* adequate until and unless it is defeated, rather than requiring an independent and prior showing of reliability.

But don't I still have to *judge* that the content reflected in the "built-in" awareness and that claimed in the second-order belief are the same, and doesn't this higher-order judgment itself require some independent justification, thus continuing the regress? Internalist epistemologists are sometimes charged with "overintellectualizing" the situation of justification, an accusation that generally seems to me to have little force, being based as it is on little more than the internalist demand that justification involve a genuine and intelligible reason for thinking that the belief to be justified is true. Here, however, I am inclined to regard the suggestion that a direct comparison of two conscious states need involve an independent judgment that must in turn be justified by something other than the conscious contents of the states themselves as a clear case of objectionable overintellectualization. If any intellectual comparison or assessment can ever be direct and unmediated by a further judgment, surely this one can. And to deny that this is ever possible is to guarantee vicious regresses in all directions, rendering the operation of the intellect inherently futile.

4.2 Constitutive Awareness versus Higher-order Thoughts

The foregoing account of the foundational status of meta-beliefs about one's own occurrent beliefs or thoughts relies on a more general view concerning the nature and status of the property of being conscious as it pertains to conscious mental states: that for a mental state to be conscious is for it to involve as an intrinsic, constituent feature, one that is a part of its own internal character and that depends not at all on any further reflective or apperceptive state, a conscious awareness of its distinctive sort of content (and possibly of other things as well, such as the attitude toward that content that it embodies).³ This view seems to me quite plausible, even obvious, from a sheerly intuitive stand-

3 The presence of such an intrinsic awareness is claimed to be a feature of *conscious* mental states, not necessarily of all mental states; that is a further issue.

point. But since some may not find it so, it will be useful to explain and defend it further by contrasting it with the only very clear alternative: a view that has become known as the “higher-order thought” theory of consciousness, first explicitly advanced and defended by David Rosenthal.⁴ According to Rosenthal’s view, consciousness is not an *intrinsic* property of any mental state. Instead, one mental state becomes conscious only by being the object of a *second* mental state, a “higher-order thought” that one has or is in the first mental state.

The issue between these two views of the nature of consciousness is obviously crucial for the defensibility of the version of foundationalism that was suggested in the previous section. If Rosenthal is right, the conscious dimension of my first-level belief or thought that *foundationalism is more defensible than most philosophers think* depends on the existence of an appropriate second-order thought, viz. the thought that *I have the occurrent belief that foundationalism is more defensible than most philosophers think*. In consequence (as he never quite says explicitly but seems clearly to intend), there would be no consciousness at all of the content of the first-level occurrent belief were that content not apprehended in the second-order thought. And if this were so, then conscious awareness of the first-level content would apparently occur only as a part of the content of the second-order thought, so that there would be no built-in or constitutive awareness of the content of the first-order belief or thought to appeal to for the justification of a meta-belief, as discussed in the previous section. And since such a second-order thought appears to be just as much in need of justification as the original meta-belief – indeed it might apparently just be that meta-belief in an occurrent form – appealing to its content would also yield no ultimately foundational justification. Thus if the higher-order thought theory of consciousness is correct, the foundationalist view suggested above will not work.

Fortunately, however, there is a clear and decisive reason why the higher-order thought theory cannot be correct. (A second, somewhat less decisive but still weighty reason will emerge later.) This can be seen by first noticing that the higher-order thoughts whose occurrence supposedly confers consciousness on lower-order thoughts cannot all themselves be conscious, as this is understood by Rosenthal’s account.

4 David Rosenthal, “Two concepts of consciousness,” *Philosophical Studies*, 94 (1986), pp. 329–59; reprinted in Rosenthal (ed.), *The Nature of Mind* (New York: Oxford, University Press, 1991), pp. 462–77. (Page references in the text are to this reprint.)

One such higher-order thought may of course be conscious by virtue of being the object of a still higher-order thought, but since it is reasonably clear that an infinite hierarchy of such higher-order thoughts does not and probably cannot exist, there must in any sequence of such thoughts, each having the previous one as its object, be a highest-order thought in that sequence that is not in turn the object of a still higher-order thought and that thus, according to the higher-order thought theory, is not itself conscious. All this Rosenthal accepts and indeed seems to advocate (p. 467), though, if I am right, without fully appreciating its significance.

Now consider again my earlier example of my first-order conscious belief that *foundationalism is more defensible than most philosophers think*. On Rosenthal's account, as we have seen, the status of this thought as conscious must result from a second-level thought that I have the first-level thought, rather than from any built-in or constitutive conscious awareness. Let us suppose, for the sake of simplicity, that the second-order thought is in this case not itself the object of any higher-order thought, and so, according to his account, is not itself conscious at all. Rosenthal seems to regard this as the most typical case (pp. 465–6), and considering more complicated possibilities would yield the same ultimate result, albeit in a somewhat more complicated way.

The problem is now to understand how and why according to this picture I am conscious of the content of my first-level belief or thought at all. It is clear that I am not conscious of that content merely by virtue of having the first-order thought. And though the first-level content is reflected in the content of the second-level thought, *I am not conscious of that second-level content either*, on Rosenthal's view, since there is no higher-order thought about *it*. Thus it is entirely obscure where the consciousness of the first-level content is supposed to come from or to reside. If the first-order thought is not in itself conscious and the second-order thought is not in itself conscious, and if there are in this case no higher-order thoughts (which would only yield a longer sequence of non-conscious thoughts), then there seems to be no consciousness of the first-level content present at all – contrary to the stipulation that we are dealing with a first-level thought that is (somehow) conscious. This seems to me to constitute a *reductio ad absurdum* of the higher-order thought theory of consciousness.

My diagnosis is that Rosenthal (along with many others including, perhaps, even Descartes himself) has confused two subtly but crucially

different things: first, the consciousness of the *content* of a conscious mental state, which is, I have suggested, intrinsic to the occurrence of that state itself; and, second, the reflective or apperceptive consciousness of that state itself, i.e., the consciousness *that* such a state has occurred, which I agree requires a second-level or apperceptive state. This confusion is plainly reflected in Rosenthal's statement that "conscious states are simply mental states we are conscious of being in" (p. 462): to be conscious of *being in a state* requires a consciousness of the state itself as a kind of object of thought and not merely a consciousness of its content. Conflating these two things thus leads quite inevitably to the view that a mental state could be intrinsically conscious only by somehow, paradoxically, having *both* its ordinary content and the further, self-referential content that it itself occurs. Rosenthal is surely right to reject such a view (pp. 469–70), but wrong that it is the only alternative to his higher-order thought theory.

One important point is worth adding. If a particular mental state does not involve an intrinsic awareness of its own content (which I have not argued here to be impossible), one may still come in some way to have a higher-order thought that it (the lower-order state) exists. But such a higher-order thought, even if it were itself intrinsically conscious in the way indicated, would not, contrary to Rosenthal's claim, somehow transmute the state that is its object into a conscious state, even though the subject would be conscious *of* that state as an object. In such a situation, the subject would be conscious of the content of the lower-order state only indirectly and only as described or characterized in the content of the higher-order state – which description might of course be incomplete or less than fully accurate and would in any case be in conceptual terms. Thus the lower-order state would still not be itself a conscious state in any genuine way. Especially where the state in question is a qualitative or sensory state with a content that is not itself conceptual (see further below), there is all the difference in the world between an external and conceptual awareness that it occurs and an actual conscious awareness of the qualitative or sensory content itself.

My conclusion here is that the higher-order thought theory is untenable and indeed obviously so, with the only apparent alternative being the view advocated here that an awareness of the appropriate sort of content is an intrinsic, constitutive feature of those mental states that are conscious – thus making it possible, as discussed above, to appeal to that awareness to justify a foundational belief.

4.3 The Justification of Beliefs about Sensory Experience

The account given in the last two sections seems to me to make it reasonably clear how one relatively specific sort of belief, namely, an apperceptive meta-belief about the existence and content of an occurrent first-order belief, can be basic or foundational in the sense of there being an internally available reason why it is likely to be true that does not depend on any further belief or other cognitive state that is itself in need of justification (though, as we have seen, it is really the constitutive awareness of content rather than the meta-belief that turns out to be the ultimate source of justification). Apart from possible objections based on externalist theories of belief content, which I find extremely implausible but have no time to go into here,⁵ the account in question also seems to me to be extremely obvious – almost too obvious to warrant discussing, were it not that so many, my own earlier self included, have managed to miss it.

Where does this leave us? Even this much of a foundationalist ingredient would be a valuable, indeed essential addition to erstwhile coherence theories like those discussed in chapter 3. Though there would still be a serious problem of what to say about non-occurrent beliefs, to which the account sketched so far is not applicable in any very straightforward way, at least a good deal of the problem of access to one's own beliefs would be solved, thus perhaps avoiding the need for anything as dubious as the doxastic presumption. But it is very doubtful at best that foundational beliefs whose content is restricted in this way to the existence and character of one's own beliefs are enough by themselves, even if worries stemming from the occurrent-dispositional distinction are set aside, to provide an adequate basis for the justification of beliefs about the objective physical world in general. What seems needed for this purpose is a foundation that also includes beliefs that reflect the content of other kinds of experience, especially the sensory or perceptual experience on which knowledge of the physical world seems so obviously to depend. Can the foregoing account be extended to beliefs of this sort?

Consider a state of, e.g., visual experience, such as the one that I am presently having as I sit at my computer table. Like an occurrent belief,

5 For some relevant discussion, see my paper "Is thought a symbolic process?," *Synthèse*, 89 (1991), pp. 331–52.

such an experience is a conscious state. What this means, I suggest, is that, in a way that parallels the account of occurrent belief or thought offered above, it essentially involves a constitutive, or “built-in,” non-apperceptive awareness of its own distinctive sort of content, namely sensory content.⁶ And, again in parallel fashion, such a constitutive awareness of sensory content is in no need of justification and is indeed infallible in the sense that there is no sort of mistake that is even relevant to it. Since it is this awareness of sensory content that gives my experiential state the specific character that it has and thus constitutes it as the specific experiential state that it is, there is simply no logical room for this awareness to be mistaken about the content in question, no independent fact that it could somehow get wrong. Such a built-in awareness of sensory content is thus also apparently available to justify genuinely foundational beliefs in a way parallel to the case already discussed. Thus if I have the belief that, e.g., my present visual experience includes a red, square patch in the approximate middle of my visual field, and if my constitutive or built-in awareness of the content of my actual conscious visual experience includes an awareness of such an element, then this latter awareness seems to provide a clear and compelling reason for thinking that the belief is true. And if I am aware of this reason, as of course I might not be, then, at least as long as other things are equal, the belief is seemingly justified in a way that does not rely on any further belief, thus making it basic in the sense required by foundationalism.

There is, however, a recently popular objection to such a view that needs to be addressed. This objection, which is present with various degrees of explicitness in the thought of philosophers as different as Popper, Sellars, Davidson, and Rorty,⁷ is obviously related to the Sellarsian dilemma discussed earlier, but still different enough on the surface to warrant independent consideration. It begins with the idea that the distinctive content of a sensory or perceptual experience, the

6 Such content is not, as we will see, propositional or conceptual in character, and this may seem to some to make the very word “content” inappropriate. But while agreeing that there is a certain potential for confusion here, I know of no better term for what one is conscious of in having sensory or phenomenal states of consciousness, and so will continue to employ it.

7 See Karl Popper, *The Logic of Scientific Discovery* (New York: Harper, 1959), sections 25–30; Wilfrid Sellars, “Empiricism and the philosophy of mind,” reprinted in his *Science, Perception and Reality* (London: Routledge and Kegan Paul, 1963); Richard Rorty, *Philosophy and the Mirror of Nature* (Princeton, NJ: Princeton University Press, 1979), chs 3 and 4; and Donald Davidson, “A coherence theory of truth and knowledge,” in Dieter Henrich (ed.), *Kant oder Hegel* (Stuttgart: Klett-Cotta, 1983), pp. 423–38.

content that makes the experience the very experience that it is, is *non-conceptual* (and non-propositional) in character – where what this means is at least that the most basic awareness of this content is not couched in general or classificatory terms, is not an awareness of the experience as falling under general categories or universals. And from this the conclusion is drawn that an awareness having content of this sort cannot stand in any intelligible justificatory relation to a belief formulated in conceptual and propositional terms, and hence that the relation between the two must be merely *causal*. As Davidson puts it:

The relation between a sensation and a belief cannot be logical, since sensations are not beliefs or other propositional attitudes. What then is the relation? The answer is, I think, obvious: the relation is causal. Sensations cause some beliefs and in *this* sense are the basis or ground of those beliefs. But a causal explanation of a belief does not show how or why the belief is justified.⁸

And if this were correct, what I have been calling the constitutive or built-in awareness of sensory content, even though it undeniably exists, would be incapable of playing any justificatory role and thus would apparently have no real epistemological significance.

The premise of this argument, namely, the claim that the basic content of sensory experience is essentially non-conceptual in character, seems to me both true and important, even if somewhat elusive. One way to get at the point is to notice that the content of, e.g., the visual experience that I am having as I look out over my presently rather wild yard seems far too specific, detailed, and variegated to be adequately captured in any conceptual or propositional formulation – or at least in any that I am presently able to formulate or perhaps even understand. I can enumerate the large-scale features: a number of maple and red cedar trees, a deck, various sorts of planters, bushes and shrubs of various kinds, etc., and could further characterize each of these in conceptual terms. But to describe in conceptual terms the subtle patterns of light and shading and color that are present in even a small portion of my visual field is plainly beyond my present conceptual powers, even though all of this is nonetheless plainly there in my conscious experience. Moreover, and even more basically, if we can somehow imagine having an ideally complete and fine-grained conceptual description of such an experience, it seems clear that thinking

8 Davidson, "Coherence theory," p. 428.

in conceptual terms of very specific shades of color in some complicated pattern is not at all the same thing as actually *experiencing* the pattern of colors itself. (Here we see the second objection, alluded to above, to the higher-order thought theory of consciousness: a higher-order conceptual thought could not account for the distinctive sort of consciousness that a conscious sensory state involves.)

But despite the wide influence of this argument (to which I myself was once susceptible), careful reflection will reveal that the conclusion simply does not follow from the premise. Even if we grant and indeed insist that the specific content of a sensory experience is itself non-conceptual in the way claimed, this provides no reason at all to deny that such non-conceptual content, like the various other kinds of non-conceptual phenomena, can still be conceptually *described* with various degrees of detail and precision. The relation between this non-conceptual content and such a conceptual description thereof may not be strictly *logical*, as Davidson uses the term, i.e., may involve nothing like the relations of inference, consistency or inconsistency, etc., that may exist between two propositions, but it is also obviously not merely a causal relation. Rather it is a *descriptive* relation, having to do with the accuracy or inaccuracy of fit between a conceptual description and a non-conceptual object that the description purports to describe. And while the assessment or evaluation of the accuracy of a description is not quite the same thing as the logical assessment or evaluation of an inference, it is nonetheless normative and even logical in a broader sense that would have no application to a merely causal relation.

But the important point for our purposes is that where such a relation of description exists, the character of the non-conceptual object is what determines whether the conceptual description is correct or true. And thus an awareness of that non-conceptual character can seemingly constitute a kind of *reason* for thinking that the description is true or correct (or equally, of course, untrue or incorrect) – thus apparently providing a basis for the justification of the conceptual claim.⁹ Such a reason or basis is, of course, only available to one who has some sort of *independent* awareness of the character of the non-conceptual item, i.e., an awareness that does not rely on the conceptual description itself. In the most familiar sorts of cases, e.g., where it is some physical object or situation that is being described, one could apparently

9 This seems to me a perfectly acceptable use of the term “reason,” albeit one that is not quite parallel to other uses of the term. But the use of this particular term is not essential to the view I am presenting here.

have an awareness that is independent of the description in question and otherwise unproblematic from an epistemological standpoint only via some second conceptual state that embodies a second, perhaps more specific description, and this second description would of course itself equally require justification, so that no foundational justification would result.¹⁰ But in the very special case we are concerned with, where the non-conceptual item being described is itself a *conscious* state, one can be aware of its character via the constitutive or “built-in” awareness of content without the need for a further conceptual description and thereby be in a position to recognize directly that a conceptually formulated belief about that state is correct, that the relation of description that is implicitly claimed by the belief to hold really does hold.

Thus when I have a conscious state of sensory experience, I am, as already argued, aware of the specific sensory content of that state simply by virtue of having that experience. And hence if an apperceptive belief that I entertain purports to describe or conceptually characterize that perceptual content, albeit no doubt incompletely and abstractly, and if I understand the descriptive content of that belief, i.e., understand what an experience would have to be like in order to satisfy the conceptual description, then I seem to be in a good, indeed an ideal, position to judge directly whether the conceptual description is accurate as far as it goes, and if so, to be thereby justified in accepting the belief. Here again there is no reason to think that mistake is impossible and thus no reason to think that such an apperceptive belief is infallible or indubitable. But as long as there is no special reason for thinking that a mistake is likely to have occurred, the fact that such a belief seems via direct comparison to accurately characterize the con-

10 There is an important assumption being made here that should not pass unnoticed. From a commonsense standpoint, it often seems as though we are in a situation in which a description of a physical object is directly justified in this way. Thus if I believe that there is a pencil on my computer table, it is natural to think that I can judge the correctness of this descriptive claim and thereby be justified in accepting it simply by looking to see if there is indeed a pencil there. At one level, there is no doubt that I can do this and that we are often justified in this way. But the epistemological problem, of course, is that the justification of the perceptual belief (or judgment or awareness or whatever it should be called) being appealed to in such a case is itself anything but unproblematic. In the absence of some account of how we have direct or immediate cognitive access to the physical world, an account that I at least see no way to give, this perceptual belief must itself apparently be justified by appeal to the underlying sensory experience in something like the way that will eventually be outlined here – thus preventing the perceptual awareness of physical objects from constituting in itself a basis for foundational justification.

conscious experience that it purports to describe apparently provides an entirely adequate basis for thinking that the description is correct and hence an adequate basis for justification. It is the fact that in this case and this case alone, the person is able to judge the accuracy of a conceptual description in a way that raises no further issues of justification – not some dubious claim of infallibility or incorrigibility – that provides the basic reason for thinking that it is beliefs about conscious experience that constitute the foundation of empirical knowledge.

4.4 “Direct Apprehension” Revisited

Here we seem in fact to have found a solution to a very difficult and long-standing philosophical problem, one that is implicit in the epistemic regress problem discussed in chapter 1 and is also related to the third of the standard objections to coherence theories discussed in chapter 3 (though much more general in its application): Assuming, as I am here, that truth is correspondence to or agreement with the relevant chunk of independent reality, how is it even possible for us to have good reasons to think that such correspondence has actually been achieved in a particular case? Clearly the reason for thinking that correspondence has been achieved in one case might appeal to other, prior cases in which correspondence has allegedly been achieved. But it has seemed obvious to many philosophers that not all assessments of correspondence can be thus indirect, that there must be at least *some* cases where the achievement of correspondence can be recognized or judged directly if there is to be any adequate basis for indirect assessments in other cases. And at the same time, however, many philosophers, myself again alas included, have rejected any such direct confrontation between conceptual thought and reality as inherently impossible, though without ever really making clear why such a result does not lead immediately to skepticism. How, as it is sometimes put, can we step outside our minds to judge that our own conceptual states agree with something external to them and to us?

But if the foregoing account is correct, we seem to have found a case of exactly the sort of direct comparison or “confrontation” between a conceptual description and the non-conceptual element or chunk of reality that it purports to describe which seems intuitively to be essential if our conceptual descriptions are ever to capture reality in an ascertainable way. Such a comparison can only take place, to be sure, where the reality in question is itself a conscious state and where the

description in question pertains to the conscious content of that very state, but in that very specific case it seems to be entirely unproblematic and perfectly genuine. Thus contrary to many recent critics of foundationalism, the idea that reality is in some circumstances simply *given* to the mind in a way that makes the truth of claims about it directly and unproblematically apparent is, after all, not a myth!

I believe in fact that it is this sort of non-apperceptive, intrinsic awareness of the content of a conscious state that epistemologists such as those mentioned in chapter 1 had at least primarily in mind in their use of the notion of “direct apprehension” or “immediate acquaintance.” But if this is right, then discussions of direct acquaintance were often needlessly obscure, suggesting as they did some sort of mysteriously authoritative or infallible apprehension of an independent cognitive object, rather than an awareness that is simply constitutive of a conscious state itself. Moreover, and more importantly, the claim of some proponents of direct or immediate acquaintance that one might possibly be immediately acquainted in this sense with other sorts of entities, perhaps even with physical objects or their surfaces, simply makes no sense on the present account of what immediate or direct acquaintance really amounts to, since neither a physical object nor its surface can be literally part of the content of a conscious state. This seems to vindicate the frequent, but usually unargued claim of other proponents of direct acquaintance that one can be immediately acquainted only with one’s own mental states.¹¹ I also believe that it is this sort of constitutive or “built-in” awareness of the content of a conscious state that Chisholm had in mind in speaking of states that are “self-presenting,”¹² a terminology that seems rather more appropriate to the phenomenon in question than “acquaintance.”

The foregoing discussion seems to me to show pretty clearly that a potential foundation for empirical justification genuinely exists, consisting, more or less as traditional foundationalists thought, of beliefs about the content of sensory experience, together with beliefs about the contents of other conscious states. It also seems to me to suggest

11 I am limiting my attention here to claims of direct acquaintance with matters of concrete and contingent fact. The idea of direct acquaintance is sometimes also applied to the a priori awareness of necessary truths and abstract entities, but it is clear that the present account of what direct acquaintance involves will not work in that area. Whether there is some other account that is applicable there is an issue that lies beyond the scope of the present discussion.

12 See, e.g., Chisholm, *Theory of Knowledge*, 3rd edn. (Englewood Cliffs, NJ): Prentice-Hall, 1989), pp. 18–19.

strongly that no other sort of foundation for empirical justification is available. But while this result is a clear and essential prerequisite for the viability of a traditional foundationalism, it still leaves a host of further questions and difficulties to be dealt with. In the final chapter, I will try to say a little about some of these, leading up to a necessarily brief and schematic discussion of the venerable problem of the external world.

The Conceptualization of Sensory Experience and the Problem of the External World

Despite much recent skepticism, we seem, if the argument of the previous chapter is correct, to have located a secure foundation for empirical justification more or less where traditional foundationalists always thought that it was to be found: in the form of beliefs about the content of conscious states and particularly of states of sensory experience. Having located such a foundation, the question is now what can be done with it. Can the various non-foundational beliefs that seem intuitively to be justified, especially beliefs about physical objects and their properties, be plausibly justified by appeal to the foundation thus identified? This is the issue that will be explored in this final chapter, though I should add that I regard the present discussion as substantially more tentative and provisional than that of the earlier chapters. Despite the long history of this topic, there are a number of thorny questions that seem to me to be inadequately explored and quite possibly inadequately understood. Thus my aim in this chapter is more to bring the main issues into clearer focus than to attempt any very definitive resolution.

5.1 The Conceptual Formulation of Sensory Experience

Perhaps the most immediately urgent question is what specific form the allegedly foundational beliefs about sensory content might take. In asking this question, I am taking it for granted that for the content of sensory experience to play any epistemic role, it is necessary that such content be conceptually formulable in beliefs that are explicitly about

it. Things would be far easier if it were plausible to hold, as some have,¹ that non-conceptual content could somehow directly justify beliefs that are not immediately about it, e.g., beliefs whose content has to do with physical objects, without the experiential content needing to be itself formulated or formulable in conceptual terms. I believe, however, that any such view is pretty obviously untenable, that Davidson and the others are right in thinking that no intelligible relation of justification can hold between non-conceptual sensory content and conceptual beliefs in general. I have argued that a special sort of justificatory relation can exist in the specific case where the conceptual belief is a purported description of the conscious, non-conceptual experiential content itself, but the clear implication of that discussion is that this is the *only* sort of case in which such a justificatory relation between non-conceptual experience and belief is possible.²

Thus it is impossible, in my view, for a foundationalist to avoid the issue of how non-conceptual experience is or should be described by the foundationally justified beliefs. It may be, of course, that different sorts of descriptions are possible, but the main question is what sort of description is optimal for epistemological purposes. Here there are two main views, which it would be possible to combine in various proportions:

One view is that the content of sensory experience should be described in *phenomenological* terms, terms that are as close as possible to the apparent character of the given experience itself – that is, in terms of something like the pure sense-datum concepts envisaged by various philosophers in the first half of the twentieth century.³ The

1 See, e.g., Paul Moser, *Knowledge and Evidence* (London: Cambridge University Press, 1989). Moser's view is that non-conceptual contents can directly justify physical-object claims in virtue of the fact that the latter *explain* the former, but he says almost nothing about how the explanatory relation in question is supposed to work – or, more importantly, about how any such relation of explanation could be intelligible apart from some conceptual description of the phenomena to be explained.

2 It is reasonably clear that ordinary people do not in fact formulate such conceptual descriptions of their experiences in any very explicit way, so that justification that depends on such a formulation will be at best available to them on reflection and perhaps only very much in principle. Whether and to what extent this is a problem will be briefly considered further on.

3 In saying this, I do not mean to be committing myself to an ontology of sense-data. In fact, I am inclined to think that the so-called “adverbial” account of the contents of experience is almost certainly correct, the main reason being that there is no acceptable account to be given of the relation that would have to exist between ontologically independent sense-data and the mind that apprehends them that is independent of the adverbial theory: if the sense-data somehow affect the mind in a way that reflects

advocates of such views have usually assumed that the resulting description of, e.g., visual experience would be in terms of patches of color arranged in visual space,⁴ and I am inclined tentatively to accept such a picture.

The second main view is that sensory content should be described primarily in terms of the physical objects and situations that we would be inclined on the basis of that experience, other things being equal, to think we are perceiving. Thus on this view, for example, the conceptual description of my present visual experience would characterize it as the sort of experience that in the absence of countervailing considerations would lead me to think that I am sitting about a foot away from a large wooden table, on which a computer, speakers, various books and pieces of paper, etc., are located (all of which could be spelled out at great length). The usual way of putting this is to say that what I am conceptually aware of is certain physical-object *appearances* – or, in adverbial terms, of ways of being “appeared to” that are characterized in physical-object terms. Where the appearance in question is a visual appearance, we may say alternatively that it *looks* as though there are objects of the sorts indicated, and analogously for other sensory modalities.⁵

As already suggested, these two views are not entirely incompatible: it would be possible to hold that sensory content is or at least

their character, then the resulting adverbially characterizable states of mind are really all that matter, making the sense-data themselves superfluous; and if they do not affect the mind in such a way, then their apprehension by that mind is difficult or impossible to make sense of. But the most important point for present purposes is that any characterization of sensory experience that can be given in sense-datum terms can equally well be adopted by an adverbial theorist, simply by construing a comprehensive sense-datum description of one’s sensory experience as characterizing the specific manner in which one is adverbially “appeared to.” This is why the issue between these two views, while important in other respects, makes no difference at all, in my judgment, to the epistemological issues that we are mainly concerned with here.

4 Such views are divided on the issue of whether the visual space in question is two- or three-dimensional. But while ultimately quite important, this is an issue of detail that is beyond the scope of the present discussion.

5 Some philosophers have objected pretty strenuously to the idea that, e.g., ordinary “looks” statements can be construed as descriptions of non-conceptual sensory content. See, e.g., Wilfrid Sellars, “Empiricism and the philosophy of mind,” in his *Science, Perception and Reality* (London: Routledge and Kegan Paul, 1963), sections 10–23. I do not have space to enter into this controversy here and must content myself with saying that the objections in question seem to me to show at most that there are other senses of “looks” besides the one that is relevant here (e.g., one that indicates a tentative or guarded opinion about what is actually there), but have no serious tendency to show that the sense of “looks” presently at issue does not exist.

could be described partly or even entirely in each of these two ways. But the main question is which is optimal, both practically and epistemologically, and here there are a number of different issues that are relevant.

The idea that a purely phenomenological or sense-datum description could accurately capture the content of experience would nowadays be rejected by many as wrong-headed in principle,⁶ but this seems to me to be a mistake. I can see no reason why it would not be possible for us to have the conceptual resources to provide such a phenomenological description of experience to any level of precision and accuracy desired, even though it seems obvious that we would always fall short of an ideally complete description – as seems to be the case with any conceptual description.

But even if an account of experience in phenomenological terms represents a theoretical possibility, it seems clear that most people do not in fact possess the conceptual resources that would be needed for anything approaching a full description. Ordinary people are indeed capable of giving reasonably precise and accurate phenomenological or at least quasi-phenomenological descriptions of some aspects of their experience, and a person, such as an artist or a wine taster, who cultivates this ability can often do a good deal better in a particular area. But it is doubtful whether even those whose abilities of this sort are the best developed are in a position to conceptually formulate a strictly phenomenological characterization of their overall sense experience that is sufficiently detailed and precise to capture all or even most of its justificatory significance for claims about the physical world (assuming for the moment that it has such significance). And in any case, it is exceedingly clear that most of us do not even begin to approach such a capacity. Moreover, even if we did possess the needed conceptual resources, it seems also clear that the time and effort required to formulate a justificatorily adequate description in such terms, whether overtly in language or internally to oneself, would be prohibitive.

These difficulties may seem sufficient to show that the alternative characterization of experience in terms of physical-object appearances must be the primary one (assuming that these are the only two alternatives). I propose to tentatively accept this conclusion and see where it leads us.

6 See, e.g., P. F. Strawson, "Perception and its objects," in G. F. Macdonald (ed.), *Perception and Identity* (Ithaca, NY: Cornell University Press, 1979), pp. 41–60.

5.2 The Significance of Physical-object Characterizations of Experience

Many philosophers have questioned whether ordinary people in ordinary perceptual situations normally or standardly have beliefs of *any* sort about the non-conceptual content of their sensory experience, even ones that are couched in physical-object appearance terms. But it is hard to see how a person who has, e.g., a visual belief about a certain sort of physical object can fail to also have at least an implicit grasp of the character of his visual experience in physical-object appearance terms. Such a person is after all surely aware that the perceptual claim in question is a result of vision, i.e., that he *sees* the object; and this seems to bring with it at least an implicit realization that his visual experience is such as to make it look as though an object of that specific sort is there. That most ordinary persons would not couch matters in such explicitly philosophical terms must, of course, be granted, but this does not seem to me to show that they are not aware in a more inchoate way of what the philosophical account more explicitly formulates.

But while such conceptual characterizations of experience in physical-object appearance terms do seem to me to be generally available even to ordinary people, it is easy to misunderstand their significance, and many philosophers seem to have done so. In particular, it is crucially important to distinguish a description of experience that merely indicates what sort of physical objects and situations seem to be presented on the basis of the experiential content itself from one that embodies some further causal or relational claim about the connection between that experience and the physical realm. For since the justification of such a further claim would have to appeal to something beyond the experienced content itself, no belief involving such a claim could be foundational in the way indicated in the previous chapter.

A useful example of the sort of danger that I am warning against is provided by Susan Haack's book *Evidence and Inquiry*, in the course of which she attempts to give a specification of the evidential force of a state of perceptual experience. Her suggestion is that this can be captured by a set of propositions ascribing the perceptual states to the subject in question. Thus, for example, such an ascription might say that the subject "is in the sort of perceptual state a person would be in, in normal circumstances, when looking at a rabbit three feet away and in good light" or "is in the sort of perceptual state a normal subject

would be in, in normal circumstances, when getting a brief glimpse of a fast-moving rabbit at dusk.”⁷ Haack’s discussion of this point is not as clear or full as one might like, but the specific formulations offered make it reasonable to suppose that these characterizations are intended to describe the experience in terms of the physical situations that are *causally* or *lawfully* connected with it, rather than in terms of its intrinsic experiential content. This, however, is precisely the sort of description that cannot be justified by appeal to the experienced content alone. The content of my experience may no doubt incline me to think that a rabbit is present, but that content obviously cannot by itself reveal that it is in fact of the sort that is normally (or indeed ever) *caused* by rabbits. A useful way of putting the point is to say that the claims about physical appearances or ways of being appeared to that constitute our conceptual formulations of the intrinsic content of sensory experience must be understood in what Chisholm has called the “descriptive, non-comparative” sense of the terms or concepts in question,⁸ for only in that sense can the claim to be “appeared to” in a certain way be adequately justified simply by appeal to our constitutive or “built-in” awareness of the non-conceptual sensory content alone.

As already suggested, it seems reasonably plausible that we do in general have the ability to grasp or represent the character of our perceptual experience fairly accurately, albeit somewhat obliquely, in terms of such physical-object appearances. But once illegitimate construals like Haack’s are set aside, it seems to me far from obvious exactly what such characterizations of experience really amount to. If in describing an experience as an appearance of a rabbit (or as “being appeared to rabbit-ly”), we are saying nothing about its relation, causal or otherwise, to actual rabbits, what then are we saying? It seems clear on reflection that there is a correlation of some sort between experiential content and physical situations that we are relying on here: one of which we have a detailed and confident albeit merely tacit grasp, and that we also assume without much question to be shared by others. Whether this correlation is learned or whether it is at least largely innate is difficult to decide, but there is no doubt that we are confidently guided by it in the vast majority of cases, even though we are unable to even begin to formulate it explicitly.

In speaking here of a “correlation,” I do not mean to suggest that it is a *mere* correlation, that the experiential content and the

7 Susan Haack, *Evidence and Inquiry: Towards Reconstruction in Epistemology* (Oxford: Blackwell, 1993), p. 80.

8 See, e.g., *Theory of Knowledge*, 3rd edn., p. 23.

corresponding propositional claim about physical objects are only externally coordinated, without being connected with each other in any more significant way. Indeed, this seems quite obviously wrong. From an intuitive standpoint, it seems plausible to think that the experiential content is in itself *somehow* strongly suggestive of or perhaps even in some interesting way structurally isomorphic to the correlated physical situation.

Ultimately I will want to suggest that this seeming isomorphism needs to be spelled out in something like phenomenological or sense-datum terms in order for the justificatory force of sensory experience to be adequately captured. But the point for the moment is that once the nature of this correlation is understood, there is no immediately apparent basis for thinking that it is in fact dependable, that beliefs adopted on the basis of it are likely to reflect in an accurate way what is really going on in the physical world. And plainly this is not something that can simply be *assumed* in an epistemological context where it is the very justification of physical-object beliefs that is in question.

5.3 The Inference from Sensory Experience to the Physical World: Four Alternatives

This brings us then finally to the main issue: how is the inference from sensory experience (which we are assuming for the moment to be characterized in physical-object appearance terms) to the physical world to be justified? Here there are four main alternatives that are worthy of mention (though the last really amounts to a repudiation of this way of construing the issue).

First. Perhaps the most historically standard solution is the reductive phenomenalist attempt to define concepts pertaining to physical objects and their properties in terms of sensory appearance concepts.⁹ Though there are difficult issues of conceptual priority involved,¹⁰ it seems to me that this approach may well succeed for the specific case of secondary qualities. But the problems afflicting a more global phenomenalist approach are both well known and, in my judgment, clearly fatal.

9 Perhaps the most thoroughly developed version of this view is that offered by C. I. Lewis, *An Analysis of Knowledge and Valuation* (La Salle, Ill.: Open Court, 1946).

10 See, e.g., the discussion of this issue in the first of Sellars's Carus lectures, "The lever of Archimedes," *The Monist*, 64 (1981), pp. 3–36.

Second. A quite different solution is advocated by H. H. Price¹¹ and, in what seems to be a rather seriously qualified form, by Chisholm,¹² among others. The core idea of this view is that the mere occurrence of a physical-object appearance or state of being appeared to confers *prima facie* justification on the corresponding physical claim. Chisholm's somewhat more explicit version of this solution appeals to a supposed logical relation of "tending to make evident" that is alleged to exist between claims or beliefs about sensory appearances and the corresponding claims or beliefs about the actual perception of physical objects, where the idea seems to be that this is a primitive relation of which no further explication or analysis can be given. Thus, for example, the claim would be that my belief that my present visual experience involves appearances of a computer on a table, or my belief that I am being appeared to in the corresponding way, *tends to make evident* my belief that I am actually perceiving such a computer and such a table and thus that they really exist in the physical world. Such an evidential tendency is capable of being defeated by countervailing evidence, but where no such defeater is present, it is claimed, the claim of genuine perception and so of corresponding physical reality is justified.

The difficulty with this sort of view is that it seems very implausible to suppose that such a primitive and unanalyzable logical relation of "tending to make evident" or "tending to justify" genuinely exists between an individual belief about physical-object appearances and the corresponding belief about physical reality. To be sure, Chisholm's claim is not that any such relation is discernible *a priori* in itself, but only that its existence is an *a priori* consequence of the "general presupposition" or "faith," roughly, that epistemological success is possible,¹³ a view that already concedes a good deal to skepticism. But over and above concerns of this sort, the problem is that if a belief about a physical-object appearance is construed, as I have argued and Chisholm seems to agree that it must be construed, as merely a useful though somewhat oblique way of describing the non-conceptual content of sensory experience, then it is hard to see how such a belief could by itself have any direct or immediate evidential or justificatory significance of the sort claimed. Why should the occurrence in my mind of a conscious state with a particular sort of non-conceptual

11 See H. H. Price, *Perception*, 2nd edn. (London: Methuen, 1950), ch. 7.

12 *Theory of Knowledge*, 3rd edn., pp. 46–54, 64–8, 71–4.

13 *Theory of Knowledge*, 3rd edn., pp. 4–6, 72–3.

content be thought to have any bearing at all on whether a specific sort of physical object exists in the mind-independent world? If there is any answer to this question, it seemingly must depend in some describable and explicable way on the specific features of the experience and perhaps also the way in which that experience relates to other experiences of the same general sort. It thus seems hard to make sense of any such primitive and inexplicable relation of “tending to make evident” or “tending to justify,” whether a priori justifiable or not.

One way to appreciate this point is to notice that if descriptions of the given content of experience in terms of physical-object appearances are understood in the way just indicated, rather than as embodying some further claim for which additional justification would be required, then it would be a mistake to think that they have any epistemological, as opposed to practical, advantage over descriptions of such experience in purely phenomenological terms. The experiential content being described is the same in either case, and its justificatory capacity is not somehow enhanced by failing to conceptualize it in the terms that would give the most explicit and detailed description of it. But there seems to be no plausibility at all to the idea that a purely phenomenological description of the same experience that is in fact conceptualized as a particular physical-object appearance (or state of being appeared to) would by itself, in isolation from other such experiences, have any tendency to justify or render evident the corresponding claim about the physical world.

Third. It might be suggested that the basis for the needed inference from sensory appearance to physical reality is to be found in two fundamental facts about such physical-object appearances, facts that were noticed by Locke and Berkeley, among others: first, the involuntary, spontaneous character of their occurrence; and second, the fact that the physical descriptions they embody fit together and reinforce each other in a *coherent* fashion, presenting a relatively seamless and immensely complicated picture of an ongoing physical world.¹⁴ These two fundamental facts are, of course, the central ones appealed to by Locke, in justifying his inference from sensory ideas to the external world; and also by Berkeley, in justifying his inference to the God who is supposed to produce our ideas. In both cases, the underlying idea, rather more explicit in Berkeley, is that some *explanation* is needed for the combination of involuntariness and coherence, since it can be ascribed

14 See John Locke, *An Essay Concerning Human Understanding*, Book IV, ch. xi; and George Berkeley, *The Principles of Human Knowledge*, paragraphs 28–30.

neither to anything like voluntary selection nor, because of the high degree of coherence, to chance. And each of them argues on this basis that the explanatory hypothesis that he advocates (the physical world hypothesis for Locke and the God hypothesis for Berkeley) is thereby justified as the best explanation of the facts in question.¹⁵

It indeed seems clear that the combination of spontaneity and intricate coherence requires *some* explanation beyond mere chance. But *why* an explanation in physical-object terms, rather than any of the other possibilities (including Berkeley's) that so obviously exist? What makes the physical explanation so obviously salient is our ingrained inclination, based on the tacitly understood correlation already discussed, to describe the experiential content in physical terms (or, indeed, to leap directly to a physical claim with no explicit acknowledgment of the experiential premise). Since it is, however, this very correlation between experience and physical-object claims whose justification is ultimately at issue, no appeal to that correlation can as such have any justificatory weight. Thus some further, independent reason needs to be offered for thinking that the physical-object explanation is indeed the best one.

Fourth. The final alternative appeals to the idea, already briefly mentioned, that the correlation between experiential content and physical objects is not a *mere* correlation, that there are features of the experiential content itself that are strongly and systematically isomorphic or structurally similar to the correlated physical situations. If this is so, it might provide a reason to prefer the physical explanation of experience to the various others that might be given, and it is hard to see what other sort of reason there might be. But this seems to indicate in turn that the conceptual characterization of experience in physical-object terms cannot be taken as epistemologically primary after all, that something like a sense-datum or phenomenological characterization is needed if an adequate basis for the inference to the physical world is to be found. In the next section, relying on the pioneering work of Broad and especially Price,¹⁶ I will attempt to sketch, in a necessarily brief and schematic way, some of the features of experience that seem to be relevant here. Then, in the final section, I will present and defend an equally schematic account of how the account of experience thus

15 For a useful discussion and elaboration of Locke's argument, see J. L. Mackie, *Problems from Locke* (Oxford: Oxford University Press, 1976), ch. 2. See also *SEK*, ch. 8, for a somewhat different version of the same underlying idea, couched there in terms of a coherence theory.

16 C. D. Broad, *Scientific Thought* (London: Routledge & Kegan Paul, 1923), Part II; and H. H. Price, *Perception*, 2nd edn. (London: Methuen, 1950).

given can be used to justify an inference to the existence and properties of physical objects.

5.4 A Preliminary Phenomenological Characterization of Sensory Experience

We are now in a position to see that the characterization of the non-conceptual content of sensory experience in terms of physical-object appearances is epistemologically unsatisfactory in two closely related ways. First, it depends on the very correlation between non-conceptual experience and physical objects whose reliability is ultimately at issue, thus making it difficult or impossible to avoid begging the central question as to the accuracy or reliability of this correlation. Second, and even more importantly, it tends to obscure the various fine details and nuances of experience upon which, I am suggesting, both the correlation itself and the justification that is being sought must ultimately depend. Think, for example, of the extremely large number of different visual presentations that could be lumped together as appearances of a table or even as appearances of a relatively specific sort of table at close range. A crucial part of the overall issue is precisely what it is that warrants viewing these qualitatively distinguishable experiences as all appearances of one and the same specific sort of physical object or physical situation, something that can be adequately dealt with only by considering that qualitative character in its own right, rather than lumping experiences that are intrinsically very different under the same physical-object appearance description.

Such a qualitative characterization of sensory experience in its own right is something that very few philosophers have even attempted to give: partly because of the extreme difficulty of doing so in a clear way while operating with a public language whose conceptual resources have been shaped by more narrowly practical concerns; and even more because of the great length that would seemingly be required to do an even approximately adequate job. Given the space limitations of the present essay, I propose to simplify the task in three ways. First, I will largely confine myself to a sketch of some of the features of experience pointed out by Price in his much more extensive account in *Perception*,¹⁷ an account that while surely debatable on many points of detail, still

¹⁷ Mainly in chs. 4 and 8 of *Perception*. Since I am severely summarizing and condensing a much longer discussion, I will not attempt to give specific page references.

seems to me fundamentally correct in its main outlines. Second, I will, for the sake of brevity, make use of a tool that although perhaps acceptable in a preliminary sketch, would be obviously objectionable in a full account: namely, the use of physical-object descriptions to direct the reader's attention to the features of non-conceptual experience that are relevant. While the justificatory force of such experiential features in relation to physical-object claims ultimately depends, I believe, on their being describable in ways that are independent of the physical object claims that they seem to support, the plausibility of the general sort of justificatory argument that I want to consider can, I think, be adequately appreciated for present purposes without actually giving such an independent description. Third, I will employ sense-datum terminology as a convenient way of referring to various relatively specific aspects and features of sensory experience, even though I think that the view that sensory experience is literally an acquaintance with entities of the sort that such terminology suggests is very likely mistaken.¹⁸

In these terms, the suggestion I eventually want to make is that the specific characteristics of our sense-data, and especially their spatial characteristics, are such as to be easily and naturally explainable by supposing that they are systematically caused by a relatively definite world of mostly solid objects arranged in three-dimensional space, and by no other hypothesis that is not, in a way to be further explained below, essentially parasitic on that one. What then are these characteristics?

First. Think of the visual sense-data experienced while doing what we commonsensically think of as moving around a medium-sized physical object, perhaps Price's favorite example of a matchbox, at close range, observing its various sides from various distances and angles under relatively uniform and "normal" conditions of lighting. In relation to each side of the box, there will be a collection of spatial sense-data, varying more or less continuously in shape, intuitively as the angle from which the object is being perceived is altered. Within each such collection, there will be a much smaller set of what might be called central sense-data, having geometrically similar two-dimensional shapes and also satisfying the following two conditions: (1) The two-dimensional shapes of the other data in the collection can be

18 See the discussion in n. 3 to this chapter. Restating the following discussion in adverbial terms would greatly complicate the formulation, while not altering the essential content.

regarded as perspectival distortions in various directions of the shapes of these data (which are in this way central to the collection). (2) At a constant apparent distance, the two-dimensional shapes of the central data in the collection corresponding to each side of the object are such as could be fitted together in three dimensions to form a closed three-dimensional shape, one which is in fact the same as the commonsensical shape of the box; and perspectively distorted versions of two or three such two-dimensional shapes can be experienced as adjoined to each other in ways that can be viewed as perspectively distorted images of that three-dimensional shape. These central sense-data are roughly what Price calls “nuclear sense-data.”¹⁹

Second. Think now of the sense-data corresponding to intuitively less adequate visual perceptions of the matchbox: perceptions at greater distances, through distorting media of various sorts (e.g., water or wavy glass), and under varied conditions of lighting. Here too the sense-data experienced can be regarded as related to the intuitively more adequate “nuclear” data already discussed via various sorts of distortions that are still perspectival in a somewhat broader sense. And it is in general possible, at least in principle, to experience what Price calls a “gradual transition series” leading from one of these intuitively less adequate sense-data to one of the more adequate “nuclear” ones via a series of intermediate sense-data in which the changes between any two members in the series can be made as gradual as one likes.

Third. Consider now the tactual sense-data that would be experienced while touching the matchbox, running one’s hands over its various sides and edges. Here too, there will be “nuclear” sense-data having shapes that correspond intuitively to the various faces of the box, and that can be experienced much more directly than the visual data as fitting together to form a three-dimensional shape. There will also be other sense-data that depart from these nuclear data via various sorts of milder distortion and incomplete perception. Price’s view, which I am inclined to accept, is roughly that the collection of tactual sense-data can be coordinated with the larger and more complicated collection of visual sense-data by virtue of the geometrical properties that sense-data in each group, especially the “nuclear” ones, jointly

19 Though his account differs by claiming that at relatively close range, the shape and size of the sense-data are constant, that is, unaffected by perceptual distortion. This goes along with his view that sense data are located and literally face different directions in a three-dimensional visual space.

instantiate.²⁰ Taking the visual and tactual sense-data together, we have, corresponding to our commonsensical matchbox, roughly what Price calls a “family” of sense-data, where the central structure of such a “family” is constituted by the relations in which its members stand to the “nuclear” sense-data and thereby to the three-dimensional shape to which those data are related. His most important claim about all this, which I believe to be correct (though it surely has not been adequately established by the foregoing sketch), is that all of this structure among the sense-data (as well as the further points to follow) is in principle discernible in its own right, without reference to common sense or scientific conceptions of the physical objects to which we intuitively think the sense-data in question are related.

Fourth. Sense-data pertaining intuitively to other senses, mainly aural, olfactory, and thermal sense-data, can also be viewed as members of such “families,” with the main connection between them and the visual and actual sense-data being the way in which their intensity is coordinated with the visual or tactual data that are simultaneously experienced, increasing in general as one “moves” through a “gradual transition series” leading to an experience of larger visual sense-data of the “family” in question. (Thus, for example, if the matchbox is an old one with a distinctive musty smell, the olfactory data in question can be connected to the visual and tactual sense-data of the matchbox “family” by appeal to the fact that as the visual sense-data of that family grow larger, i.e., occupy larger and larger regions of one’s visual field, the musty smell becomes correspondingly stronger.)

Fifth. Consider now the regular and repeatable sequences of sense-data from different “families” that intuitively correspond to moving around in space, experiencing first one object, then an object adjacent to the first (where it is often possible to simultaneously experience members of the two families corresponding to the two objects), then a further object adjacent to the second one, etc., etc., in various directions, perhaps returning ultimately to the original object, but

20 What is mainly at issue here is, of course, Molyneux’s problem. Price takes it as just obvious that the very same geometrical properties are exemplified by both visual and tactual data, so that there is no need for an empirical correlation between the properties of one and the properties of the other. (He also seems to think that the coordination between the visual and tactual sense-data pertaining intuitively to a single object can be done by considering those sense-data alone; whereas I believe, because of the problem posed by multiple objects with approximately the same geometrical properties, that it will require considering the sense-data corresponding intuitively to a relatively large group of spatially related objects.)

now viewed from a different perspective. (Think of walking across a university campus, experiencing one building, then an adjacent one, then one adjacent to the second, and so on, in various directions.)

Sixth. Consider the various patterns of sense-data that intuitively reflect the ways in which one object can block, either wholly or partially, our view of another and the way in which these patterns relate to those connected with movement, as just discussed.

Seventh. Consider the extremely varied sense-data that intuitively reflect the causal actions of commonsense objects on each other, e.g., such things as the effects of a source of heat on objects that are subject to melting or burning, together with the ways in which the sense-data corresponding to such effects are correlated with those intuitively reflecting movement from one such object to another and the time required for such movements. Here the rough idea is that the sense-data corresponding to causal effects in general often vary in intensity in a way that correlates with the intuitive distance between the two objects as reflected in the sense-data intuitively correlated with movement. (Think of a fire or perhaps just a source of light.) There are also patterns of sense-data that reflect what we intuitively think of as various blocking or shielding effects.

Eighth. Think, finally, of the ways in which the families of sense-data and their relations to each other as thus far specified change over time in ways that intuitively reflect both changes in the objects in question and the movement of the observer.

Though the foregoing is obviously only the merest sketch of matters that would require a book much longer than Price's to be described in full detail, I hope that it is adequate to indicate in a very approximate way the facts about our experience upon which the inference to the physical world must, I am suggesting, finally rest. One important point, which I have so far ignored, should, however, be added. The features of experience that I have described are not realized in any very complete way in the sense-data actually experienced by a given person or even, for that matter, by all persons taken together. Many of them pertain in large part to sense-data that are merely possible or (better) *obtainable*: sense-data that *would* be experienced if, from an intuitive standpoint, the observer's point of view were to change in ways that it does not in fact change. This means that there are really two distinguishable inferences (or attempted inferences) involved in the attempted justification of belief in the physical or material world on the basis of sense experience: one from actual experience to obtainable experience, which seems to be broadly inductive in character; and a

second abductive or explanatory inference from actual and obtainable experience to physical objects. While the first of these two is by no means entirely unproblematic, I will focus my attention here almost entirely on the second, where the problems and difficulties are obviously much greater.²¹

5.5 The Inference to the Physical World

On the view that I want to tentatively defend, one main premise for the inference from experience to the physical world is the detailed account of experience of which a preliminary sketch has just been given. A second main premise is that some explanation is clearly needed for this complicated pattern of experience, that it cannot be plausibly viewed as either just a matter of chance or as somehow an ultimate and not further explicable brute fact (the latter being essentially the view of the phenomenalist). The intuitive credentials of this second premise seem pretty obvious, and I will not offer any more explicit defense of it here.

It seems quite clear that at least one possible explanation for the experiential patterns in question is the approximately commonsensical idea that my sensory experiences are systematically caused by a realm of three-dimensional objects, (1) having at least approximately the shapes corresponding to those reflected in the “nuclear” sense-data, (2) through which I move in such a way as to change my point of view, (3) which are spatially related to each other in the ways reflected in the sequences produced by my apparent movement, and (4) which have causal properties and change over time in the ways corresponding to the relevant further aspects of the experiential patterns. Obviously a full account of such an explanation would have also to involve at least a general account of the various human senses and the way they are affected by different kinds of circumstances, with vision being pretty clearly the most important sense of all. All of this could and ultimately should be spelled out in vastly greater detail, but the general idea should be clear enough for present purposes.

Despite this lack of detail, it seems to me reasonably safe to assume that the explanation just roughly sketched, which I will henceforth refer to as the *quasi-commonsensical* hypothesis (with the reason for

21 I am also ignoring here the problem of other minds, which would also have to be considered in a fuller account.

the qualification to be explained later), is at least one relatively adequate explanation of the details of our sensory experience. The main issue then is what other explanations are available. Here it will be useful to draw a distinction between two general kinds of explanations. On the one hand there are explanations, such as the one involving the quasi-commonsensical hypothesis just indicated, that explain the features of experience more or less directly by appeal to basic features of objects in the hypothesized world: objects that are, in a way that is difficult to characterize precisely but still reasonably intelligible, directly reflected or represented in those experiences. And on the other hand, there are explanations that explain experience by appeal to the combination of something like a *representation* of the sort of world that figures in the quasi-commonsensical hypothesis, together with some agent or mechanism that produces experience in perceivers like us in a way that, as it were, mimics the experience that we would have if the represented world were actual and we were located in it, even though neither of these things is in fact the case. Here Berkeley's view, if understood in the most natural way, provides an example: in systematically causing ideas in us, God is guided by his own more complete set of ideas, which constitute a much more adequate representation of the world we seem to inhabit.²² Adopting a useful technological metaphor, I will refer to the former sort of explanation as an *analog* explanation and the latter as a *digital* explanation.

In these terms, it is in fact, I submit, far from obvious that there is any alternative *analog* explanation available that is even approximately as good as that provided by the quasi-commonsensical hypothesis. The main points that make the prospects for such an alternative seem quite dim are the following. (1) It is very hard to see how the spatial features of experience could be explained in analog fashion by anything other than a spatial world: certainly a multi-dimensional world of some sort seems needed to account for the various sorts of experiential paths that return to the same experiential starting point. (2) For essentially the same reason, a two-dimensional world does not seem to have enough

22 Berkeley prefers, of course, to regard all of the relevant ideas – or perhaps rather just God's ideas – as *constituting* the reality of the world, rather than as representing a world that does not actually exist. I doubt, however, whether such a view is really intelligible, whether even God could make sense of the order reflected in his causation of ideas without viewing them as representations of a physical realm, even if no such realm actually exists. Part of the point here is that it cannot be supposed to be merely an accident or coincidence that our experience conforms so well to what would be expected from such a material world.

internal complexity to account for all of the possible experiential sequences and variations. (3) Thus an alternative analog hypothesis would arguably have to involve a world of at least three dimensions containing objects whose shapes and relations differ systematically from those that are actually reflected in our experience, with the character of experience thus involving something like a systematic distortion. But if the differences in question are supposed to be large enough to be interesting, then the very features that make the quasi-commonsensual hypothesis work so well also make it hard to see how such an alternative could work at all.²³ At least it seems fair to say that no one has ever actually described such a possible alternative in even an approximate way, and here it seems to me quite reasonable to place the burden of proof on those, if there are any, who seriously believe that such analog alternatives genuinely exist.

But even if I am right that there are no serious contenders for an alternative analog explanation of sensory experience, there are still many, probably indefinitely many, possible *digital* explanations. Indeed, it seems obvious that for any analog explanation of anything, there are *always* indefinitely many possible corresponding digital explanations, since anything that can be explained by appeal to the features of a particular object or structured set of objects can also be explained by appeal to a representation of those objects (perhaps in conceptual terms or perhaps in a form analogous to a map) together with an appropriate mechanism of some sort that “translates” from features of the representation to the output that the represented objects would produce (given relevant background conditions). In the present case, as already suggested, such explanations will involve some sort of representation of the world depicted by the quasi-commonsensual hypothesis, together with some agent or mechanism (Berkeley’s God, Descartes’s demon, or the computer that feeds electrical impulses to a brain-in-a-vat) that generates experiences of the sort that we would have if we existed in the quasi-commonsensual world thus represented.

Is there then any rational basis for preferring the analog explanation offered by the quasi-commonsensual hypothesis to the various digital alternatives? Consider again the basic *modus operandi* of a digital explanation: Rather than supposing that the features of the world that explain experience are as close as possible to those actually reflected

23 Some of the relevant points here are contained in Jonathan Bennett’s discussion of “size blindness” in his “Substance, reality, and primary qualities,” *American Philosophical Quarterly*, 2 (1965).

in experience, such an explanation claims instead that the true features of the world are utterly different from those that are reflected in experience. It seems clear that there is something rather arbitrary about such a view, as partially reflected in the fact that there is and apparently could be no basis at all for preferring one such digital mechanism to another. It seems intuitively unreasonable to adopt such a further, arbitrary mechanism until and unless there is some specific feature of the data to be explained that requires it.

There is also a second reason for preferring an analog explanation like the quasi-commonsensical hypothesis to the corresponding digital explanations. The explanatory success of a digital explanation of experience depends in effect on the truth of two claims: first, that the corresponding analog explanation could indeed account for the experience in question (that a material world could produce experience of the sort that we in fact have), since the digital explanation works by emulating the action of the cause or causes postulated by the analog explanation; and, second, that the specific translating mechanism postulated by the digital explanation in question can indeed successfully do the job of emulation (that God or the computer can indeed systematically produce the sort of experience that would be produced by the represented material world). But the explanatory success of the analog, quasi-commonsensical hypothesis depends on the truth of only the first of these claims. My tentative suggestion is that this makes the quasi-commonsensical hypothesis less vulnerable to problems and so more likely to be true than the various digital alternatives.²⁴

My claim is thus that there may in this way be a good reason to think that the quasi-commonsensical hypothesis provides the best explanation for the detailed contours of our sensory experience, so that we would thereby be justified in accepting it. I will conclude this necessarily sketchy discussion with three further comments about this argument.

First. Even if the reasons advanced here for preferring the quasi-commonsensical hypothesis succeed in showing that it is more likely to be true than any of the digital alternatives, they do not, at least not without substantial further discussion, show either: (a) that it is more likely to be true than the disjunction of all the alternatives (including the

24 It is tempting to put this point in terms of the fact that the probability of a conjunction of two non-necessary claims is lower than that of either claim separately, but this would not work if the second claim were a necessary truth, as might be claimed for at least the God hypothesis. Since I have no space to discuss this issue here, I will settle for the fuzzier formulation in the text.

hypothesis of chance or sheer coincidence) – and so more likely to be true than not; or (b) that it is likely enough to be true to satisfy our commonsense intuitions in this area or to satisfy the requirement for knowledge (whatever this turns out to be; see the discussion at the end of chapter 1). But there is no room here to explore these further, obviously important issues.

Second. At best this explanatory argument justifies attributing to physical objects only those features that are clearly required to explain the character of our experience. As already noted, these features will arguably include spatial, temporal, and causal properties. But it is pretty clear that they will not include secondary qualities like color, but only the causal power to produce experiences of such qualities in us, thus leading to the familiar Lockean view of the status of such qualities. (This is the main reason that the hypothesis in question is only “quasi-commonsensical.”)²⁵

Third. Even if everything else about this argument turns out to work, it does not seem plausible to claim that anything very closely approximating it is in the minds of ordinary people when they make claims about the physical world (even though they are arguably aware in a less explicit and unified way of the relevant features of experience). Thus if this is the best justification available for such claims, it will follow that most ordinary people are not fully and explicitly justified in making them (and so do not possess “knowledge” of such matters, if “knowledge” requires such justification). The most that can be said is that the essential elements for such an argument are at least roughly within their purview, so that the argument is in principle available to them. I do not find this to be a terribly implausible result, but others may disagree. There is, however, once again no room here to pursue the issue further.

25 The results of modern science seem to show that it is also possible to infer in the same basic way to other “theoretical” properties of such objects, i.e., such things as mass, electric charge, etc. (Though it is at least doubtful whether any of the things that can be arrived at in this way really amount to intrinsic, as opposed to relational, properties of the objects in question.)

II

**Beyond Internal
Foundations to
External Virtues**

Ernest Sosa

Knowledge and Justification

6.1 Non-propositional Knowledge

A witness may know *how many* people are in a lineup, *which* of the members of the lineup is the culprit, *why* he did the deed, *how much* he stole, *when* he did it, and *who* the victim was. What constitutes such knowledge might be, respectively, her knowledge *that* there are five people in the lineup, that the second from the left is the culprit, that his motive was robbery, that he stole two hundred dollars, that he did it at 10 p.m. the previous night, and that the victim was her husband. All such knowledge is propositional in that a proposition, *that such and such is the case*, forms its content. Things we know can often be formulated explicitly. Our witness's knowledge, for example, is in every case a proposition that we, and the witness herself, can easily formulate as a declarative sentence of English: "There are five people in the lineup," etc. But much of our propositional knowledge is not thus formulable, not easily. The witness knows what the culprit looks like, for example, what his face looks like: that is, the culprit has a certain facial appearance such that the witness knows *that* the culprit has that very appearance. Yet the witness could not well capture much of what she thereby knows in any words or symbols at her command, nor, let us suppose, could she even do so by drawing a picture. She may be able to manifest her knowledge only through images in the privacy of her own mind, or through her external conduct, perhaps through her public identification of the culprit, by picking him out of the lineup. Such knowledge can apparently play the roles in our thought, practical and theoretical, and in the justification and explanation of our conduct, normally played by linguistically formulable beliefs. The only discernible difference appears to be the difference in formulability, which seems less important than the similarities.

All the knowledge considered so far is thus propositional, whether or not it can be put, by us or by the subject, as a declarative sentence. Some knowledge can be put in symbols not part of a natural or public language, as when someone knows the opening bars of Berlioz's *Symphonie Fantastique*, knows that they go a certain way, and puts his knowledge in musical notation, without being able to put it in sentences of English or of any other natural language. Most of us know tunes that we can whistle but cannot put into words or into any symbols at all – well, unless we invent some code on the spot. But even then, we still had the knowledge *before* inventing the code. All the knowledge considered so far seems reducible to propositional knowledge in the ways indicated.

Other knowledge is not reducible. The witness, for example, knows her husband, knows him very well indeed. What constitutes this knowledge? Involved in it is propositional knowledge about the husband: about his appearance, age, biography, virtues, foibles, favorite dishes, and so on, for an enormous collection of extremely various facts about him. However, if the husband is a spy, then the CIA may know a lot of what the wife knows and, in addition, a great deal more that she does not know, and yet no one in the CIA may be able to claim correctly that they know the man himself. That is to say, no one there may have ever met him or even seen him in the flesh. Knowing someone or something, knowing some “object” in the broadest sense of this term, seems at least sometimes to require having had some special causal interaction with that “object.” This is plausibly a requirement for knowing a person, and for knowing an experience or a sight, say the sight of the Boston skyline two miles from the south on highway 93, or the experience of a cold shower after a hard run. Other “objects” can be known, however, independently of any such causal interaction requirement. Thus one can know not only a person, or a sight, or an experience, but also a theory, a field or discipline, a historical period, a novel, and so on for an enormous variety of possible objects of knowledge. In all of these propositional knowledge plays a crucial role, sometimes even an exclusive and exhaustive role, although in certain cases a certain causal interaction is required as well.

So much for *objectual knowledge*. A second sort of knowledge in apparent contrast to propositional knowledge is *procedural knowledge*. Propositional knowledge is “knowledge that,” procedural knowledge “knowledge how.” But first appearances are here deceiving; we need a closer look. Someone may know how to swim the crawl stroke while unable to capture this in words. But, again, such inability is not enough

to show that one's knowledge is not propositional. About a certain way of swimming someone may know *that* the crawl is swum in that very way. Of course there is more than one way to know that someone is swimming the crawl. One may know it as a poolside spectator, for example, or alternatively one may know it in the way one does when one is oneself the active swimmer. And these are very different ways of knowing it. The spectator knows it through the visual appearance of the swimmer in the act of swimming. The spectator can pick out those swimming the crawl from those who are doing the butterfly or the breaststroke or the backstroke. But the swimmer knows himself to be doing the crawl in a very different way. The swimmer knows not by sight, but perhaps through the kinesthetic sensations experienced in the act of swimming. Compatibly with this, however, procedural knowledge may still be seen in every case as a rather special sort of propositional knowledge, of knowledge, with respect to a certain way of doing something, that one does it *that* way.

Is knowing how something is done tantamount to having a skill? It is a commonplace, first of all, that the best practitioners are not always the best coaches. The coach excels in his knowledge of how it is done, the practitioner in his skill at doing it. One might retain one's knowledge of how the crawl is swum, moreover, even in the absence of any opportunity for manifesting that knowledge. But the skill too can be retained even absent the opportunity: one does not lose one's ability to swim the crawl when there is no water nearby. Perhaps then the swimmer's skill *is* identical with his knowledge of how it is done? If so, it is identical at most with the *agent's* knowledge of how it is done, not the spectator's. Not even the agent's knowledge seems identical with possessing a skill, however, since a defect in the efferent mechanisms may not affect the stored knowledge of how to proceed, a knowledge that *would* be triggered if only the efferent mechanisms were in good working order. Someone recently afflicted with rheumatoid arthritis may still know how one ties a certain nautical knot despite the ravages of the disease on his finger joints, which removes his ability actually to do it. So we can plausibly distinguish three things: the spectator's knowledge of how it is done, the agent's knowledge of how it is done (the knowledge of how actually to do it), and the agent's skill or ability at doing it, all three of which must be distinguished in turn from actual manifestations of the knowledge how or of the ability. A competence of any of these three sorts may be present despite failures of performance due to some lack in the circumstances or in the relatively peripheral equipment required. But defects in peripheral equipment

affect knowledge how and ability rather differently. Again, an ability may be lost through stiffening joints while the agent retains knowledge of how the thing is done. A successful operation may restore the ability immediately, suggesting that the less peripheral state plausibly describable as “knowledge of how it is done,” or even “knowledge how to do it” had remained.

6.2 Propositional Knowledge

Propositional knowledge is hence deeply involved across the gamut of knowledge attributions, even if occasionally a causal transaction is required for knowledge of some “object.” Actually, even in the knowledge of an object, much propositional knowledge is often required. In focusing on propositional knowledge, therefore, we are not artificially isolating some idiosyncratic sector of the field of human knowledge. All sorts of human knowledge involve the propositional knowledge that we shall take as our main subject, in which respect we join a philosophical tradition stretching back to Plato’s *Theaetetus*.

That early Platonic dialogue already contains suggestions about the nature of propositional knowledge, a subject taken up more recently by A. J. Ayer, C. I. Lewis, R. M. Chisholm, and many others. The tripartite account of propositional knowledge – as justified, true belief – was long the received view. Beliefs based on prejudice, superstition, and ignorance can still for all that turn out to be right, if only by luck. But no such belief amounts to knowledge, surely, which means that knowledge cannot be just true belief, but is at most true belief with a proper rationale or basis or source or status. (This is questionable, and questioned, in ways we shall consider below.) One’s belief is knowledge only given a proper rationale or grounding or basis, only with a proper status of that sort. Philosophers long thought of this status as that of a justified belief: a belief constitutes knowledge, according to that view, only if it is not only true but also “justified.” However, it is not enough that the belief be justified through securing some practical benefit. An athlete may be helped to win by her strong and steady confidence that she *will* win, which may provide her with practical justification for somehow acquiring and sustaining that confidence even in the teeth of contrary evidence. But such practical justification does not bear on whether she knows what she believes, unlike the evidence against her belief. The tripartite account hence takes that into account as follows: Subject *S* knows *P* (knows that *p*), if and only if, (1) *P* is

true, (2) *S* believes *P*, and (3) *S* is (epistemically) justified in believing *P*.

In a paper as incisive as it is brief, Edmund Gettier refutes the tripartite analysis: First, it is possible that a subject *S* be justified in believing that *p* even if it is *false* that *p* – where, for convenience, we assume that if you are justified in believing something, you must at a minimum indeed believe that thing; so we need to distinguish between someone who *would be* justified in believing that *p*, presumably so long as he believed it on a certain basis or in a certain way and someone who *is actually* justified in so believing. Second, someone justified in believing that *p*, who deduces that *q* from his belief that *p*, and believes that *q* as a conclusion of that deduction, would seem justified thereby in believing that *q*. For example, suppose you believe that a friend of yours, Mary, owns a country cottage. You just sold it to her an hour ago, and your relevant evidence points unequivocally in one direction: she still owns that cottage. Since you see that Mary is in the room, therefore, you conclude that someone in the room owns a cottage (from your belief that Mary both is in the room and owns one). Actually Mary's cottage has just burned down. But as luck would have it, someone else in the room, Jane, has just inherited a cottage, so someone there now *does* have a cottage, after all. Mary does not, but Jane does. So you have true justified belief that someone there owns a cottage, without knowing what you believe. This fits Gettier's recipe and shows how you can have a true justified belief without yet knowing what you believe.

Many are the ways one can respond to Gettier with some plausibility, and by now all of these (plus some others) have been tried and defended with considerable sophistication.¹ There are at least two simple but widely divergent responses that one might try. One can respond by focusing on the notion of epistemic justification, while rejecting either or both of Gettier's assumptions about such justification, thus undercutting the counterexamples. Alternatively one can accept the Gettier reasoning and try to minimize the damage by downgrading its importance, either through some substitute analysis, or by rejecting that whole project of analysis. Next we consider some examples of these responses, starting with the second, more dismissive, strategy. According to this, although Gettier's reasoning may be applicable against the traditional, tripartite analysis, it is clearly powerless

1 See Robert Shope, *The Analysis of Knowing* (Princeton, NJ: Princeton University Press, 1983).

once we have in focus the correct understanding of the correct conditions for knowledge, which are in fact exceedingly simple, amounting to nothing more than truth and belief.

Against the objection that a lucky guess cannot amount to knowledge, it is replied that a lucky guess does not amount to true *belief* either. No belief is an island, it is argued, and once we see that a belief must be surrounded by a whole continent of belief, it is not so easy to find simple examples of beliefs that are just true by luck and do not amount to knowledge. And there is moreover a positive argument, as follows. Our aim in inquiry is surely knowledge: knowledge is thus the *telos*, the ultimate aim of inquiry. Now, either we want justification in our beliefs because it is a means to *truth*, or justification is desired along with truth as part of the *telos* of inquiry, namely knowledge. But if justification is of merely instrumental value, because justified beliefs tend to be true, then justification is not itself part of inquiry's *telos*, which is then exhausted by true belief. In that case, true belief is what knowledge must amount to, since knowledge is by definition the *telos* of inquiry, and that's what true belief turns out to be. However, if justification is added to true belief as *also* part of inquiry's *telos*, then our concept of that *telos*, our concept of knowledge, is just *incoherent*. Why so? Because it presents to us *two* goals that our beliefs aim to fulfill, two goals of inquiry – truth and justification – which cannot always be realized simultaneously, which can come apart. In order to avoid such incoherence, therefore, we should opt for a concept of knowledge as, merely, true belief.²

Why, however, should we think that knowledge must be identified with the ultimate *telos* of inquiry? It is true that in inquiry we aim for truth, we aim to attain knowledge of the truth. But why should we think that what we aim for in aiming both for truth and for knowledge is the *ultimate end* of our inquiry? Within the sport of archery we aim to hit the target as close to the bullseye as possible, an end intrinsic to that sort of activity. When engaged in the activity, don't we also prefer to hit the bullseye by means of skill and not just by luck? A gust of wind might come along and guide our arrow to the bullseye, but this will be less sweet than a hit unaided by the lucky gust. Of course a hit that through skill compensates for the wind might be sweeter yet. So I see nothing unacceptable in a notion of a good, skillful shot that goes

2 Compare Crispin Sartwell, "Why knowledge is merely true belief," *Journal of Philosophy*, 89 (1992), pp. 167–80; and "Knowledge is merely true belief," *American Philosophical Quarterly*, 28 (1991), pp. 157–65.

beyond that of a mere winning or accurate shot. A winning, accurate shot may have been just lucky and not at all skillful, and not in that sense a good shot. In archery we want accurate, winning shots, but we also want shots that are good and skillful. Are the goodness and skill that we want in our shots qualities that we want merely as means? Maybe so, but it seems unlikely given that we would not be fully satisfied even with many accurate, winning shots, so long as they derived from sheer luck, and manifested no real skill, even if we would gain *some* satisfaction through hitting the mark (not to speak of prizes, fame, etc.). In any case, whether or not the goodness and skill that we want are desired as means only, it remains that there is a perfectly understandable concept of a good, skillful shot, which we are free to understand as a concept that includes *both* hitting the mark *and* doing so through skill appropriate to the circumstances. Can there be any doubt that we have such a concept concerning archery? Surely we do, along with many analogous concepts in other sports, *mutatis mutandis*. What precludes our conceiving of knowledge in a similar way, as a desideratum that includes an intrinsic success component as a proper part, a hitting of the mark of truth, while containing also a component that evaluates *how* one accomplishes that, how one succeeds in hitting the mark of truth? On this conception, knowledge is not just hitting the mark but hitting the mark somehow through means proper and skillful enough. There seems nothing “incoherent” in any pejorative sense in such a desideratum of “knowledge,” and plenty of analogous desiderata grace the wide gamut of human endeavors.³

3 And it should also be noted how unusual is the concept of incoherence in play, and potentially how misleading. That a goal is thus “incoherent” does not mean that in no conceivable circumstances could it possibly be attained, but at most that some conceivable circumstances would preclude its attainment. Thus the “incoherence” of a goal of true justified belief arises only from the fact that in conceivable circumstances – those of the victim of the Cartesian evil demon, for example – the justification of one’s beliefs may come at the cost of allowing much falsehood as well into one’s view of things. But this does not mean that our goal of attaining true justified beliefs could not possibly be attained, ever. Clearly there is reason to avoid goals that are truly, strongly incoherent, since we shall be inevitably frustrated in pursuing them. (Actually one might still adopt some such goal as an “ideal,” if its *pursuit* is known to yield substantial enough benefits.) *In any case*, that it is “incoherent” in the weak sense is no sufficient reason for avoiding a goal. Anyone thirsty and hungry may have as an ultimate goal both eating and drinking. The fact that in conceivable circumstances the two goals might not be jointly attainable is no reason to *avoid* that complex goal, and opt for either eating or drinking *to the exclusion of the other goal*. This would seem absurd in a situation where both food and drink are at hand for the taking.

Again, what is not so cannot be known to be so: if anyone is to know that p then at a minimum it must be so, it must be true, that p . Truth is thus necessary for knowledge; but it is not sufficient. Many truths go unremarked by anyone at all; and any one of us misses truths without number. If one is to know that p , not only must it be true that p : one must also believe it, at least implicitly. However, even truth combined with belief still does not necessarily constitute knowledge. A gambler may be convinced that the dice will come up seven, and may predict correctly, and profitably, without thereby manifesting knowledge. In some sense one's belief must be well enough justified, which the gambler's belief is not. If one is to know that p , one must believe (presume, take it for granted, etc.) that p with adequate justification. But even that is not quite enough. I may read from a clock that it is five o'clock, and thus arrive at a correct and justified belief, without knowing what I believe, since unbeknownst to me the clock is *stopped* at that reading. My justification apparently depends on my assumption that the clock is working properly, and I may be perfectly well justified in so assuming. (For years I have depended regularly upon that clock.) But if, unbeknownst to me, the clock is stopped, then my justified true belief that it is five o'clock is not knowledge. The Gettier problem derives from such counterexamples to the thesis that any true justified belief amounts to knowledge. In one sort of Gettier counterexample, as we have seen, a subject flawlessly infers a true belief from a false justified belief. The new belief thereby becomes a true justified belief, but if the subject's only basis for accepting it is the inference from the falsehood, then despite being true and justified, the new belief is no knowledge.

Note the reliance of this Gettier recipe on a false justified belief, without which there can be no such Gettier problem. The problem hence will not touch the rationalist, who allows only such fail-safe faculties as intuition and deduction. This commits him to rejecting the very possibility of a false justified belief, which makes him immune to such counterexamples. However, our rationalist epistemology goes with a conception of knowledge as superknowledge, and of justification as superjustification, rarities hardly ever present when one attributes "knowledge" or "justification" to oneself or others. More realistic, less demanding epistemologies do entail the possibility of a belief at once false and justified, and thus give rise to the Gettier problem.

Justificationalist solutions to the Gettier problem strengthen the justification condition: for example, by requiring that the underlying argument or inference, if any, not depend on any false premise or

lemma or presupposition, or by requiring that the justification not be “defeasible” (in a technical sense) by any falsehood. This sort of solution in any case retains the requirement of justification for knowledge. Let us next consider such justification and its sources.

6.3 Sources of Justification

We focus on the empirical justification of knowledge that someone might enjoy at a given time. There are many things one might mean by such “justification,” but here we shall focus on the epistemic quality of how a given belief is formed or sustained. Is the belief “well formed?” This is a matter of how the believer comes to have or sustain that belief. In calling it *empirically* well formed or justified we are passing judgment on the *epistemic* quality of whatever then leads to it for that subject. Supposing you are then empirically justified in holding a certain belief, what might that fact derive from? What sorts of conditions can make you thus justified in so believing? It is useful to distinguish next among three things:

- B* The subject’s beliefs
- E* The subject’s sensory experiences
- W* The world external to the subject’s psychology

(All for a given subject *S* at a time *t*.)

How might someone come to be empirically justified? Coherentism tries to explain this by appeal simply to relations of coherence among beliefs in *B*. But this unacceptably detaches empirical, epistemic justification from truth. For a belief, i.e. an instance of believing by a given subject at a specific time, would then seem empirically justified if accompanied by a comprehensive and coherent enough body of beliefs, regardless of how that belief was related to experience or the world beyond.

It might be argued that the conditions required for a belief to have its content $\langle p \rangle$ already guarantee that it is likely to be true. Hence the coherentist has not after all detached justification from truth, since no belief can be detached wholly from truth: simply in virtue of being a contentful empirical belief it must enjoy some truth connection. But this also shows why a problem remains. We want a way to distinguish beliefs that are justified epistemically from those that are not. Surely

it is *possible* to believe something without adequate epistemic justification. What distinguishes beliefs that *are* epistemically justified from those that are not?

It might be thought, with the coherentist, that the difference resides in the greater coherence of the justified beliefs within *B*. The paranoid can be coherent without being epistemically justified, however, and part of the problem lies at the belief/experience or the experience/world interface. The paranoid conveniently “see” things that are not really there and that others cannot “see.” Whether this is a defect in how they respond to their actual sensory experience or whether it is a defect in how their experience reflects the world beyond, it is a defect either way, and no amount of coherence will remove it or prevent it. Such a defect can remain despite the most brilliant coherence (paranoia being compatible with logical brilliance and imaginativeness). The beliefs of the paranoid need not be epistemically justified despite being comprehensively coherent.

When a belief is *epistemically* justified, something renders it justified, and distinguishes it from beliefs that derive from paranoia, or wishful thinking, or superstition, and the like. What is thus epistemically justified, epistemically effective, must be truth-connected. So, when the belief is about the external world, the feature that renders it justified must somehow involve that world, and must do so in a way that goes beyond the world-involvement required even for it to have its worldly content. After all, the unjustified beliefs of the paranoid and the superstitious are also about the world and must have the world-involvement required for them to have their proper content.

Many are persuaded that coherentism cannot provide any adequate truth connection. Some then turn to foundationalism, in one or another of its two main varieties: classical foundationalism, and externalist foundationalism. These agree that epistemic justification requires more than just belief coherence. But they disagree in what they propose as a further source of justification. Classical foundationalism proposes what is given in sensory experience and more generally in states of consciousness. Externalist foundationalism proposes some external relation to the world beyond.

Against externalist foundationalism one might object that what justifies one in holding a given belief must be *accessible to one's own reflection*. Externalist foundationalism violates this intuition in proposing, as a source of empirical justification, an external relation between the justified belief and the world beyond, perhaps a relation of causation, or of subjunctive tracking, or of reliable belief formation.

No such relation to the external world is accessible through mere reflection.

Classical foundationalism can more easily satisfy the accessibility requirement, since the further source of justification that it invokes, namely experience, still lies within the mind of the subject, as does her body of beliefs. That further source is hence accessible to a “reflection” broadly enough understood. Indeed the classical foundationalist supposes that we enjoy a kind of privileged access to our own states of consciousness. These states are *given* to us or *present* to our minds in ways that account for their foundational status and for how well we know them.

6.4 Foundationalism versus Coherentism

We have glimpsed alternatives as to how a belief gets the sort of epistemic justification (or aptness or warrant, or positive epistemic status) that it needs in order to constitute knowledge. These include:

Coherentism. When a belief is epistemically justified, it is so in virtue of being part of a coherent body of beliefs, one sufficiently coherent and appropriately comprehensive.

Classical Foundationalism. When a belief is epistemically justified, it is so in virtue of being either the taking of the given, the mere recording of what is present to the mind of the believer, or else based ultimately on such foundations.

Externalism. When a belief is epistemically justified, that is because it comes from an epistemically, truth-conducively reliable process or faculty or intellectual virtue.

The beliefs supporting a belief *B* might be thought to form one of at least three possible structures. First, a tree that branches infinitely upwards from the “root” belief *B*. Second, a pyramid with foundational beliefs on which rest level after level of beliefs supported by reasoning, with belief *B* at the tip of the pyramid. Third, a raft one of whose planks is belief *B*, with each plank held in place by its ties of coherence with other planks.

A belief may be both true and justified without being knowledge. This is shown by our Gettier examples in which a true conclusion is deduced through a false lemma that the believer is internally, rationally, justified in believing. Her belief of that conclusion is then both

true and justified without being knowledge. Also Gettieresque is the evil-demon victim who is allowed a *true* belief that he faces a fire, along with a set of necessarily associated truths, a miniscule set by comparison with the falsehood that massively surrounds it in our victim's total body of beliefs. As regards what is internal to the experiences, memories, introspections, and reasonings of the victim, there is not the slightest flaw. He is, internally, as cognitively worthy as the best of us. In that sense, therefore, the victim is perfectly well justified in believing that he faces a nearby fire. And a fire does burn nearby, before that victim. But that is not something he can be said to know, surely, despite his belief's being true and well justified, if his supporting web of beliefs, his broader belief system, is mostly a tissue of falsehoods: if, for example, his visual, auditory, and other experiences as of a fire have nothing to do with the fire there.

What is involved in such internal justification? Classical coherentism and foundationalism are best understood as accounts of such justification. According to coherentism a belief *B* is thus justified if and only if it coheres well enough within the subject's system of beliefs. Let us now explore this in greater depth.

Is a set of beliefs coherent simply because its contents are logically and probabilistically interrelated? That is insufficient: such a set might easily lack any claim to acceptability or justification. Thus any set including *p* and *q* can simply be expanded to include $p \supset q$ (which follows deductively from *q*) and $q \supset p$ (which follows deductively from *p*).⁴ And now *p* follows deductively from other beliefs in the set, namely *q* and $q \supset p$; while *q* follows similarly from *p* and $p \supset q$. And we can continue thus at will until we acquire a very large and deductively interlocking set. Take arbitrary beliefs *p* (that the moon is made of cheese) and *q* (that clouds are cotton candy) and surround those beliefs with the likes of $(p \supset q)$, $(q \supset (p \supset q))$, $(q \supset p)$, $(p \supset (q \supset p))$, etc. The result will be a set obviously lacking desirable comprehensiveness though tightly coherent and as large as one may like. Such comprehensiveness is hence not just a matter of cardinality; scope and content of subject matter also count.

For another example, take any location *L* and any kind *K*, and consider the series: (P1) There is at least 1 *K* at *L*, (P2) There are at least 2 *K*'s at *L*, (P3) There are at least 3 *K*'s at *L*, . . . , (P*n*) There are at least *n* *K*'s at *L*, etc. For each *P_i*, *P_{i+1}* could be adduced as a supporting

4 The letters *p*, *q*, etc., will function variously; the context will be expected to reveal their specific function in that context.

reason, to infinity. A logically interrelated set is not necessarily an epistemically coherent set, therefore, and in fact the most epistemically relevant notion of coherence is not just a matter of abstract logical or probabilistic relations among propositions, but involves rather *beliefs* that are appropriately interrelated logically *and* causally. Moreover, even if the set $\{p, q, p \supset q, q \supset p\}$ is accepted in its entirety by some believer, that does not by itself suffice for the believer to have beliefs that are epistemically justified or apt for constituting knowledge, even should they be true. And the same goes for the beliefs about the number of K's at location L. It seems also necessary that the body of beliefs in question be comprehensive enough. But the relevant comprehensiveness is not just a matter of the set's containing lots and lots of beliefs. For the set about the number of K's at L might expand infinitely without that mattering much on its own for whether the believer is epistemically justified in believing any of its members.

What kind of comprehensiveness *is* required, then, if it is not just a matter of the set's cardinality? Two sorts of requirements suggest themselves. First, we might require that one's relevant beliefs cohere with each other and *also* cohere with one's relevant experiences (which makes better sense if we think of the relevant experiences as having their own propositional content, as in "it was as if I were seeing something white and round at arm's length before me"). And, secondly, it will help also to include some account of how one acquires one's beliefs. Thus one might face something red and round, and enhance one's knowledge of this fact through awareness of the source of one's belief in a good-enough faculty of color vision (along with such supportive beliefs as that it's a ripe tomato one sees, etc.).

Coherence involves the logical, explanatory, and probabilistic relations among one's beliefs. However, it would not do to attain a tightly interrelated system by lopping off whatever beliefs may refuse to fit. Theoretically, it would be possible (aided by futuristic cognitive technology) to perform such surgery on one's belief system. But that need not yield epistemic justification.

A body of beliefs may moreover be comprehensively coherent (tightly interlocking, impressively diverse, and vastly wide-ranging) and *still* fall short of justification. Take one's I/now perspective. Suppose one replaces one's concepts of (a) oneself, and of (b) the present, wherever they may figure in one's vast system of object-level beliefs, by corresponding concepts of (a') the holder of SS no. *n*, and (b') May 18, 2002. This system will nearly match one's original system in true comprehensiveness and in interlocking coherence. Yet it has

little semblance of justification. One would be wildly unjustified in attributing to the holder of SS no. n , as of 5/18/02, the vast set of things that one attributes to oneself now. The coherentist needs to require interlocking comprehensiveness not only about the object level of beliefs, but also in a way that ascends to meta-levels, where one takes note, at least implicitly, of the sources of one's beliefs and of how reliable these are, and so on. It is this requirement that blocks the transformation from the I/now system to the one involving SS no. n and 5/18/02.

Even that seems insufficient, however, if it is conceivable that such a comprehensively coherent system of beliefs could still fail to mesh properly with the subject's sensory experience. This failure of mesh might occur in either or both of two ways. S might experience as if p and might have no reason to question or resist such prompting to believe that p , but might nonetheless believe something, that q , incompatible with the proposition that p , and might even believe that not- p for good measure. In addition S might have a splitting headache while believing that he does *not* have any headache. Such failure of mesh, which seems logically possible, would preclude even a highly coherent and comprehensive system of beliefs from rendering its member beliefs justified.

Coherentism is hence well advised to adopt corresponding requirements of (1) comprehensiveness, (2) perspectival content, and (3) mesh with experience.

With such improvements, coherentism is on a rapprochement course with classical foundationalism as accounts of (internal) justification. This is especially evident given that foundationalism, for its part, must allow that the internal coherence of a body of beliefs can help give justification to its member beliefs. The potential for a meeting of the minds is evident if we put it this way:

The coherence needed for epistemic justification in a system of beliefs requires that the system be appropriately comprehensive, which means that it must include an epistemic (meta) perspective, *and* a suitable complement of *foundational* beliefs!

Once it is put thus the game is obviously up. Both sides score points. Coherentism scores for its emphasis on appropriately comprehensive (and perspectival) coherence, an indispensable component of our concluding view. And foundationalism scores for its emphasis on the extra-belief components needed in an appropriately justified system: for example, appropriate mesh with experience.

6.5 More on Epistemic Coherence

Some sort of awareness of the coherence of one's beliefs is required for justification. But this awareness may be constituted by the sensitivity to such coherence that one manifests by accepting one's system of beliefs (or a large enough fragment) *in virtue of its coherence (at least in part)*, and by adjusting one's degree of assurance to the degree of that coherence.

A further example shows how a rich enough concept of epistemic coherence requires more than just the bare logical coherence of a set of propositions. Take someone who arbitrarily believes every fifth proposition he reads and thus acquires a beautifully coherent system of beliefs, all interrelated by a sophisticated proof that only a handful of logicians would be able to grasp. The moral: epistemic coherence requires not only a logically coherent pattern in our beliefs, but also our awareness of it, in some way, to some extent. And here comes a problem for the appeal to coherence. Consider the body of beliefs in one's epistemic perspective; it's the coherence of that perspective that supposedly yields one's reflective justification and, in turn, one's reflective knowledge. The mere logical coherence of those beliefs falls short, however, so one might wish to require also awareness by the believer that the beliefs do cohere. Just here is the problem: What about this further awareness? Must it not itself be more than just a lucky guess? Must it not also qualify as a *justified* awareness with the sort of status required for knowledge? A vicious regress looms.

In considering this we must recognize the varieties of belief, and the various ways in which belief can be manifested. We need to distinguish, among beliefs, first the fully conscious from the subconscious; also, secondly, those that *are* from those that *are not* formulable by the believer; and, finally, we also need to distinguish belief manifested through episodic acceptance of a proposition somehow present to one's mind, even if not symbolically present, from belief that is manifested only in other ways. Thus a belief that $2 + 2 = 4$ can remain, subconsciously, even in sleep, even below consciousness; one can believe that someone looks or sounds a certain way, even if one is unable to capture symbolically the full content of one's belief; and one can know that in order to avoid disaster on a bicycle one must now shift one's weight or turn the handlebars a certain way, even when in fact one is prevented from doing so. One's full knowledge of how to ride a bicycle could not be spelled out in symbols, nor could one even manifest it

fully through offline imagery; yet one may still manifest such knowledge, surely, by actually riding a bicycle and successfully putting one's knowledge to the test. At various junctures in that episode of bike-riding something will need to be done and one somehow knows what it is, i.e., one's behavior manifests embodied knowledge of what then needs doing. Of course if one is deliberately trying to look awkward, for some ulterior end, one may know well enough what needs to be done, while now manifesting one's knowledge precisely by *not* doing it. So how one manifests knowledge and belief turns out to be quite complicated, as with propositional attitudes generally.

Something similar may be said of more intellectual procedure. One may acquire and sustain beliefs, manifesting thus, at least in part, sensitivity to the place of such beliefs in one's coherent networks of belief. What is more, the beliefs constitutive of one's epistemic perspective are no exception. Jungle guides, farming peasants, and experienced sailors embody much practical lore that they cannot articulate. Given a certain gestalt look of the environment, they draw practically appropriate inferences: that a storm is coming, say, despite their inability to formulate the knowledge that they embody, the knowledge that when the sea and the heavens look a certain way, a storm is likely brewing. What reveals their belief *that such an outcome is expectable when things look that way* is their repeated expectation of the outcome in particular situation after particular situation when things in fact do look that way. So they accept something of the form "When things look *F* they are likely to turn *G*" even though their acceptance is neither fully conscious nor manifestable through explicit formulation, but only through a pattern of "inferences."

And we come finally to the question "What might justify commitment to such an inference pattern?" Sundry things could do so: the fact that the pattern fits one's experience, for example, and is accepted or sustained in part at least because of that. This answer would provide a way of stopping the threatening regress. Of course an innate pattern might also do the job, especially one evolved with the species.

6.6 Lucky Knowledge?

It may be objected against externalism that a set of beliefs might be comprehensively coherent without manifesting any external virtue or any causal sensitivity to the environment. Suppose we can conceive of such a set of beliefs, beliefs thoroughly independent of the surround-

ing world that they are about. More specifically, consider the case of a believer – call him Lucky Strikes – who has as extensively true a belief system as it is humanly possible to have, but one acquired accidentally: perhaps he was struck by lightning and that is how he acquired his beliefs. For that matter, the beliefs might have arrived not in a neat package, courtesy of the lightning, but rather one by one, or anyhow a few at a time, each by its own lucky route. The evil demon in charge might for example have crews of randomizing lesser demons who instill beliefs in their victims without regard to what beliefs other lesser demons are instilling. By a huge accident, Lucky Strikes comes out with an enormous and wholly correct belief set. Shall we say that Lucky *knows* the various things that he believes? Well, recall that the victim is to have as extensive a belief system as it is humanly possible to have, and his beliefs are to be wholly true. So Lucky must believe with regard to nearly all his first-order beliefs about the things around him, his past, etc., that the reason why he believes these things has nothing to do with their being as he believes them to be. In fact, he accepts that he believes them only because the demons drew those particular beliefs at random as ones to be instilled in him; all in complete disregard of whether the propositions thus fed were true or not. In this case it is hard to find any sense in which Lucky “knows” the things that he believes.

Anyhow that is how it seems to me, even though I am quite flexible on epistemic terminology, and am even willing to grant that a super-market door can “know” that someone is approaching. In any case, there is *perhaps* some stretched sense in which there is a lot of knowledge in Lucky. Thus compare the “knowledge” that we attribute to Lucky Contestant when he guesses the right answer. Somewhat incautiously we might say that he won because he “knew” more right answers than his rivals (although more strictly we should say that he “got” more answers right). In any case, whether they can count as library knowledge or not, Lucky’s true beliefs do not qualify as the sort of knowledge to which humans can aspire, as reflective beings or even as animals.

Two further requirements for epistemic coherence plausibly flow from the case of Lucky Strikes. First, the beliefs in an epistemically coherent set cannot just sit there loose from each other. Second, for at least many an object-level belief *B*, it would be epistemically incoherent to combine *B* with a meta-belief that *B* is believed in complete causal independence of whether or not it is true. Perceptual beliefs, in particular, epistemically repel any such meta-belief.

The most comprehensive coherence accompanied by the truth of what one believes will not yet amount to knowledge. The New Evil Demon problem already shows this. Suppose we are now victims of Descartes's evil demon. Could that affect whether or not we are *epistemically justified* in believing what we believe? If we are justified as we are, we would *seem* to be equally justified, in some appropriate sense, so long as nothing changed within our whole framework of experiences and beliefs. However, if by sheer luck one happened to be right in the belief that one faces a fire, one's being *both* thus justified *and* right still would fall short of one's knowing about the fire.⁵ So whatever is to be said for coherence, or even for comprehensive coherence, one thing seems clear: none of that will explain fully what a true belief needs in order to be knowledge. One's beliefs can be comprehensively coherent without amounting to knowledge, and the same goes for one's beliefs and experiences together.

That being so, the sense of "epistemic justification" in play here will not capture fully the epistemic status that a true belief must have in order to constitute knowledge. Compatibly with this, there may still be a kind of justification, however, a kind of rational, internal justification, that is unaffected by anything external to the subject's mind. On this view, knowledge derives from a belief that both (1) is true and (rationally, internally) justified, and (2) *also* satisfies some further requirements concerning how reliably that belief is produced, or how well it tracks the truth, etc. Recall our classical foundationalism, which, again, may be viewed as merging with a kind of coherentism, especially a coherentism of both beliefs and propositionally-contentful experiences. Such internalist foundationalism does provide a good account of what is involved in a belief's being internally justified (or reasonable, or rational). Nevertheless, even this more moderate internalism is untenable, or so I will argue below.

6.7 Foundationalism Redux

Foundationalism postulates foundations for knowledge. On this its two branches – the rationalist and the empiricist – agree, even while they disagree in their preferred foundations, and on how a superstructure might rise from their favored foundations.

5 *Our* victim must *have* a body long enough during her travail to permit there being a fire "before" her, so this must depart in various ways from the strict Cartesian case. Nevertheless, surely a coherent example could be tailored to suit.

Only rational intuition will give the rationalist a secure foundation, and only deduction will build further knowledge on that foundation. The axiomatic system serves here as model, with its self-evident axioms and its theorems derived through logical deduction. In trying to reduce even arithmetic to self-evident logical axioms, therefore, logicism is in sympathy with rationalism. In his *Meditations*, Descartes sketches a more ambitious strategy for rationally founding all knowledge, not only mathematical knowledge. But his strategy required substantive commitments that proved problematic. The failures of rationalism are evident both in Descartes and in logicism.

Empiricists by contrast accept not only foundations by rational intuition but also foundations by present or recalled sensory experience. Equally unsuccessful, however, was their project of reducing all physical reality to sensory experience, whose most radical and fully realized embodiment is Carnap's phenomenalism. Besides, as Hume showed, the future cannot be predicted deductively: the reasoning required outstrips logical deduction. In response empiricism becomes doubly more liberal than rationalism: it accepts, in the first place, a broader foundation, one given not only by rational intuition but also by present or recalled sensory experience; and it admits, secondly, not only deductive but also inductive reasoning.

Such liberalization proves insufficient, however, since our knowledge is largely unsupported by sensory experience, present or even recalled. Almost nothing that one knows of history or geography or science, for example, has such adequate sensory support. And the same goes for the names of friends and relatives, and a great diversity of knowledge about artifacts, about dishes and how they taste, about how people react, and so forth. Little of that can be defended solely by induction on the basis of sensory experience, present or recalled. Take, finally, observational knowledge of immediate surroundings perceived without instruments. Not even this can easily be explained merely by induction or deduction from what one knows by introspection of one's own sensory experience. Enumerative induction evidently falls short. Nor has it been shown how we might reason validly to the external world through some abductive inference to the best explanation.

An even more liberal empiricism is now tempting, one with a broader foundation that includes not only what we intuit rationally and what we know by introspection of our own sensory experience but also what we know by direct observation of our surroundings. Consider this broader foundation, however, in its three parts: the intuitive, the introspective, and the observational. What is a rational intuition?

Is it a true belief, without inference, in something logically necessary? No, a belief in something necessary might derive from guessing or superstition or brainwashing; and no such belief would amount to knowledge. The question remains: What is a rational intuition?

Similar questions arise for other empiricist foundations: What is introspection? What is observation? Suppose a well-lit, white, triangular surface against a black background. From a favorable angle and distance, the observer sees the white triangle and knows two things. He knows, first, that his visual experience has a certain character: that of being visual experience *as if* he had a white triangle before his eyes. And he knows also that in fact he does have before him at a certain distance a white triangular surface. These are indeed paradigms of knowledge by introspection of one's own experience, and by observation of one's immediate surroundings.

Once again, suppose an observation of a white surface, well lit and ideally situated, against a black background, this time with twelve equal sides rather than three. The observer sees the white dodecagon and has two thoughts. He thinks, first, that his visual experience has a certain character, that of being visual experience as if he saw a white dodecagon. And he thinks, further, that in fact he sees a white dodecagon a certain distance away. Although he is twice right, however, he is right only by chance, for he lacks the capacity to distinguish dodecagons with a high probability of success; indeed he often confuses dodecagons with decagons. Therefore, not every introspective or observational belief constitutes foundational knowledge.

Foundational empiricism proposes three ways in which a belief might constitute foundational knowledge: intuition, introspection, and observation. These are the sources associated with the traditional doctrine of the given. The classical, internalist version of the view focuses more narrowly on the former two choices: rational intuition and introspection. Let us now explore in greater depth this traditional approach: classical, internalist foundationalism. We shall then consider its background and content, and its emblematic doctrine of *the given*.

Does Knowledge Have Foundations?

7.1 Acquaintance and Awareness

Consider first Leibniz's commitment to direct awareness:

Our direct awareness of our own existence and of our thoughts provides us with the primary truths *a posteriori*, the primary truths of fact . . . [They can be called "immediate"] because there is no mediation between the understanding and its objects.¹

Compare Russell:

We shall say that we have *acquaintance* with anything of which we are directly aware, without the intermediary of any process of inference or any knowledge of truths.²

An ambiguity is shared here by Leibniz and Russell. One's consciousness contains experiences that go unremarked: unnoticed altogether, or at least unnoticed as experiences with an intrinsic, experiential character that they nevertheless do have. Just as one automatically jumps one's jumps, smiles one's smiles, and dances one's dances, however, so one experiences one's experiences. And since experiencing is a form of awareness, one is thus in one sense automatically aware of one's experiences, precisely in experiencing them. In the same way one is aware even of experiences that escape one's notice and of which one is hence *unaware*, in another sense. What is more, it is not only her smile that the Mona Lisa smiles; she smiles her specifically enigmatic

1 *New Essays Concerning Human Understanding*, Book IV, ch. 9.

2 *The Problems of Philosophy* (Oxford University Press, 1997), p. 46 (first published in 1912).

smile. Similarly, one experiences not just one's experiencing but also one's experiencing in the specific ways in which one does experience.

Which kind of awareness do Leibniz and Russell intend: (1) noticing, intellectual awareness, whereby one occurrently believes or judges the thing noticed to be present, as characterized a certain way; or (2) experiential awareness, whereby one is "aware" directly of an experience of a certain specific sort simply in virtue of undergoing it?

That distinction – between n(oticing)-awareness, and e(xperiencing)-awareness – is important as follows. From the fact that one is e-aware of something it does not follow that one is n-aware of it. To notice a fact about one's experience at a given time is to believe correctly that it is so, but just a guess will not do: the correct belief must also be at a minimum justified, or reasonable, or epistemically appropriate, or some such thing. So what sort(s) of experience can be discerned with epistemic justification through believing it (them) to be present to one's consciousness at the time?

Foundationalists through the ages have tried to understand how we can be justified foundationally in a certain belief. And they have appealed crucially to what is "given" in one's experience, or to what is "present" to one's consciousness. So they have appealed to what we are "directly" aware of. But this requires that we be clear on the kind of awareness invoked: In particular, is it e-awareness or is it n-awareness? The latter will *not* enable the desired explanation, since the concept of "noticing" is itself epistemic in a way that unsuits it for the explanatory work that it is being asked to do. What we want is an explanation in non-epistemic terms of how a non-inferential, foundational, belief can acquire epistemic status in the first place, so that holding it is not just arbitrary, so that conclusions drawn from it can inherit epistemic status. Our explanation hence cannot rest with "noticings" that are supposed to have epistemic status *already*. The question will remain as to how *these* beliefs constitutive of the "noticings" have acquired their status.³

Thus are foundationalists led to mental phenomena epistemically more primitive than any "noticings" or beliefs, to conscious states "given in" or "present to" consciousness. In our terminology, what

3 Of course one can be justified in believing *that* one believes such and such, and the mere presence of this embedded occurrent belief might largely account for one's justification for self-attributing it, *even without the embedded belief being itself justified*. The present discussion in the main text abstracts from such cases.

foundationalists are thus led to is e-awareness: that is, to states constitutive of the subject's total consciousness at the time, including both *noticed* states and also ones that escape their notice. But we now face the "problem of the speckled hen," which worried Roderick Chisholm since his 1942 publication in *Mind* of a paper so-titled. The problem concerns the gap between e-awareness and n-awareness. Much in the intricate character of our experience can, again, escape our notice, and can even be mischaracterized, as when one takes oneself to be able to tell at a glance that an image has ten speckles although in actual fact it has eleven rather than ten. If the classical foundationalist wishes to have a theory and not just a promissory note, he needs to tell us *which* sorts of features of our states of consciousness are the epistemically effective ones, the ones *by corresponding to which specifically* do our basic beliefs acquire epistemically foundational status. Having a visual image with forty-eight speckles seems not to qualify, whereas having a visual image with three speckles may (at least when they are large and separate enough). What is the relevant difference?⁴ The full dimensions of this problem for foundationalist epistemology have not yet been properly appreciated, or so I will argue.

The distinction between e-awareness and n-awareness is also relevant to Gilbert Ryle's definition of privileged access, whose first clause reads as follows:⁵

- (1) that a mind cannot help being constantly aware of all the supposed occupants of its private stage.

What kind of awareness is involved in this clause? Is it e-awareness or is it n-awareness? If the former, then the clause is trivially empty, amounting only to the claim that a mind cannot help undergoing whatever experiences it undergoes. If the latter, then the claim is absurdly strong, and not something that any sensible form of privileged access can require: one is not required actually to notice every aspect, nor even every constitutive or internal or intrinsic aspect, of every experi-

4 When he introduces the problem on p. 122 of his *Foundations of Empirical Knowledge* (London: Macmillan, 1940), Ayer acknowledges that he owes it to another philosopher. Ironically, that philosopher was Gilbert Ryle. The objections pressed by Ryle against the possibility of privileged access are weak, whereas he also spots a problem potentially much more serious, one that, ironically, he fails to press.

5 *The Concept of Mind* (London and New York: Hutchinson, 1949), ch. 5, p. 154.

ence one undergoes. Much of the intricacy of our visual experience escapes our notice.

How then is privileged access to restrict its claim (1) within more defensible bounds? What distinguishes those occupants of one's consciousness one can know to be present "by direct inspection"? Or, put another way: What is the kind of state whose being given in one's consciousness is automatically a source of foundational status for a corresponding belief?

7.2 Knowledge of the Given: Through Thick and Thin

If by affirming a declarative, meaningful, indexical-containing sentence, one produces an utterance all of whose indexicals refer, one's utterance will have a content and thus will be true or false. By affirming such an indexical affirmable, one endows one's affirmation with content in such a way that it will be true or false, relative to its context of affirmation, depending on the truth or falsity of its content. Thus I might say "This pen is black" as I show a pen, thereby giving my utterance a content composed of the pen and the property of being black, such that my affirmation will be true if and only if the pen then has that property. If I mumble those same words while tossing and turning half-asleep, or if I affirm them as an actor in a play, however, my utterance may lack both determinate content and, therefore, truth value.

Consider the great variety of ways of being or acting: ways of being shaped, ways of being colored, ways of thinking, ways of tying one's shoes, ways of waving good bye, and so on. Such ways are among the things one can refer to by means of indexicals. Thus: "This is the way I tie my shoe, or I tie my shoe this way, or I tie my shoe thus." "This is how I wave good bye, or I wave good bye like this, or I wave good bye thus." "This is how a pencil looks when immersed in water." And so on.

Such ways of being or acting are made salient when one points to an exemplification, or perhaps when one saliently *produces* an exemplification. The way thus made salient is then the referent of the appropriate utterance of "this way" or of "thus."

Indexical affirmations occur not only in speech but also in thought. And our points remain true, *mutatis mutandis*, for indexical thought as for indexical speech. Thus one may think "This itch has been bothering me longer than that itch," while attending to the respective itches, which takes the place of physical pointing as when, pointing, one says:

“This book has been here longer than that book.” Selective attention is the index finger of the mind.⁶

Consider now the progression from ... to to, etc., to At some point in this progression from 3-membered linear arrays to 4-membered ones, etc., one will hit arrays that one can no longer classify numerically at a glance. Let us focus on a well-lit 11-membered array:

A1

Here we can distinguish:

- (a) One’s “seeing” an 11-membered linear array (as when Macbeth in his hallucination could “see” a dagger).
- (b) One’s believing that one is “seeing” an 11-membered linear array.

This second is the sort of “conceptual” belief whose content, ⟨This is an 11-membered linear array⟩, permits such inferences as: ⟨This is a linear array with less than 12 members⟩, and ⟨This is an array with half as many members as a 22-membered linear array⟩, etc.

- (c) One’s believing that one is “seeing” *thus*, where one selectively attends to the way in which that array (A1) now presents itself in one’s visual experience.

Compare A1 with the following:

A2 ...
 ..
 .
 ..
 ...

Here again we can make analogous distinctions. Thus:

- (d) One’s “seeing” an 11-membered hourglass array.
- (e) One’s believing that one is “seeing” an 11-membered hourglass array.
- (f) One’s believing that one is “seeing” *thus*.

6 I assume here that the conditions of contextual reference through the use of “thus” enable its use, aided by some mechanism of attention, to pick out a color or shape as present in that very image. Of course the requirement of such a mechanism of attention would seem to import a need for some presupposed *nondemonstrative* concepts to give content to the attention, which means that such demonstrative reference could not be conceptually fundamental.

Let's focus on (c) and (f), while assuming for simplicity that in visual experience one "sees" particular images (whatever their ontological status may be). Thus my images when I dream or hallucinate are different from yours. And let's assume visual images to have certain definite features: e.g., of shape and cardinality. Thus the images one has while focusing on A1 and A2 have certain definite features of shape and cardinality.

Consider now beliefs expressible in terms like those of (c) and (f). But let's now put it as follows:

One's believing that *this is thus*.

One such belief simply picks out an image and attributes to it the properties of shape and cardinality exemplified in it, where one picks out these properties *as* exemplified in that image. In that case, assuming I have picked out an image, I cannot go wrong in the shape and cardinality that I attribute to it. However, that is quite compatible with my failing to know that the image is *11-membered* in either case. The "phenomenal" concepts used in such a belief that "this is thus" are more primitive than even the simple arithmetical ones involving cardinality of 11. But we need to take account of a difference between A1 and A2.

So far the kind of statement or thought we are considering is just like "I exist" or "This exists" in that the conditions of reference guarantee truth. However, the guarantee of truth comes with a corresponding conceptual lightness.

From the *cogito* not much follows logically about one's nature. One could be a body, a soul, the World Spirit, whatever. And similarly little follows logically from the fact that "this is thus" as we are understanding this statement or thought *vis-à-vis*, e.g., A1. Not much can be derived about how this is from the fact that it is "thus". The concept involved here is a thin indexical one, with minimal conceptual content.

The thinness of such beliefs may be appreciated by comparing A1 and A2. Thus consider your discriminatory and recognitional capacities *vis-à-vis* these patterns. Or compare your recognition of a familiar face with your knowledge that you see a façade with a certain specific look, though one you could not distinguish in all its specificity from many similar façades; or, better yet, consider a highly irregular pattern that you would be unable to discriminate a minute later from many similar ones.

There is a difference between having just an *indexical concept* which one can apply to a perceptible characteristic, and having a

thicker *perceptual concept* of that characteristic. What is the difference? The latter involves discriminatory and recognitional capacities. It is not enough that one be able to believe correctly that this is thus, that one be able to use indexical concepts such as *this* or *thus* in order to capture a fact that one has an image of a certain specific sort, or the like. Such indexical concepts are highly portable all-purpose conceptual tools by means of which we can capture a great variety of facts by situating ourselves appropriately for use of the indexical involved. But they are also, again, very thin; not much follows from their mere content.

Thicker perceptual concepts go beyond thin indexical ones at least in requiring some ability to recognize the commonality in a diversity of items that co-exemplify some feature. Possession of such a perceptual concept would involve sensitivity, when appropriately situated, to the presence or absence of that feature. It may be thought that full grasp of the concept would require that one be able to recognize the feature as the same again when it reappears, but this implies, absurdly, that concepts cannot be lost. For if the concept is lost then one may of course lose the ability to recognize the feature as the same again. What we may require for possession of the concept *at a given time* is rather this: that within a certain set of possible alternative settings one would *at that same time* have recognized the pattern if *then* presented it; one *would* have classified the instances of the features appropriately, and would not have mischaracterized instances of other, distinct, patterns as instances of this one. This is what seems missing once we reach the level of eleven linearly arrayed dots, as above. If we had *right now* been on the next page, where there are ten similar dots, or twelve, we would not have been sensitive to these differences and might too easily have responded intellectually and linguistically as we do now when we see eleven dots. This is in contrast to the eleven-membered hourglass pattern that again we may have no word or symbol for, but that we can think to be present not just as the figure being *thus* patterned, but also as a specific phenomenally grasped pattern that we can go on to recognize on another page as the same, that we would have recognized now had we been looking at that page.

Grasping such a phenomenal concept comes with a certain guarantee of reliability, then, since it is defined in part by sensitivity to the relevant feature of which it is a concept. It is defined in part by the ability to tell when that feature is present and when absent in our experience. So we must be sufficiently reliable in the application of the concept in order to so much as grasp it. Again, therefore, the condi-

tions for giving our thought such “thicker” conceptual content automatically require that we be reliable in forming and sustaining beliefs by use of such concepts, so long as the conditions of application are undefective for their use. Thus our thick perceptual concepts enjoy a certain guarantee of reliability, as do those that are thinly indexical. What makes *introspective* beliefs with both varieties of concept so reliable is that in both cases we rely so little on conditions of application that might be defective. In these cases our relevant faculties or mechanisms operate with little or no benefit of medium or channel. We do not depend on light or air or the quality of these. We are maximally reliable because our very grasp of those concepts requires reliability in the right circumstances, and the circumstances are nearly always right, leaving little scope for possible failure.

We move beyond such concepts already with the theoretically richer concepts of arithmetic and geometry. When we apply *these* concepts to our present experience, we can easily go wrong, in the way of a claim ventured about the number of dots in linear array *A1*. So the question now remains as to which of *such* beliefs we can be justified in holding foundationally, unsupported by reasons or inferences, or anyhow independently of any such support.

Classical foundationalism needs some such beliefs with arithmetical or geometrical content, since from purely indexical or phenomenal concepts very little could be inferred, even allowing explanatory induction from the given to the external. But we still lack any adequate explanation of how we are justified through taking what is given in sheer sensory experience. More specifically, foundationalism has yet to vindicate our justified application of the thicker concepts required if we are to move adequately from the given to what lies beyond consciousness itself. For example, we still wonder how we might be justified foundationally in applying arithmetical and geometric concepts to our experience. How might one explain such classical foundational justification? This is not just a peripheral issue for classical foundationalism. Without at least a sketch of an account, only a vacuum is left where one would expect to find the solid core of the position. Let us next consider one possible sketch of the required account.

7.3 Recourse to Attention

Let us first distinguish among three sorts of concepts: indexical ones; phenomenal ones; and simple geometric and arithmetical ones: “SGA

concepts,” for short. All three can be applied to our experience. The former two – indexical concepts and phenomenal ones – come with a certain guarantee of reliability. To grasp them is at least in part to be able to apply them correctly to one’s experience, in ways sketched above. SGA concepts differ from phenomenal ones in this respect: no guarantee of reliability in applying them to experience derives simply from understanding them. SGA concepts differ from indexical ones because their conditions of reference fail to guarantee their correct application. However, the mere application to experience of indexical or phenomenal concepts will not provide a rich enough foundation for the empirical knowledge enjoyed by a normal human. Our problem for classical foundationalism is, more specifically, that it seems unable to account for how more contentful concepts, such as SGA concepts, might be applied with foundational justification. Here for example is one attempt:

An SGA belief that one’s experience has feature F (an SGA feature) is foundationally justified so long as (a) one’s experience does have feature F, and (b) one believes that one’s experience has feature F.

But the speckled hen case (see section 7.2) shows the inadequacy of this attempt.

It may be argued that classical foundationalism *can* provide an account of how we are foundationally justified in applying indexical and phenomenal concepts. Far from denying this, I have myself suggested it above. Regarding these concepts my doubt is mainly this: that they seem too thin, not thickly contentful enough to provide what is needed in a foundation of empirical knowledge.⁷ But there is also another point: the classical foundationalist explanation of the justification even of indexical and phenomenal thought (which applies indexical and phenomenal concepts respectively) may still be relatively superficial, since it may in the end turn out that the most satis-

7 And if one prefers an observationalist foundation in perceptual judgments about one’s surroundings, a similar problematic will attach to the modified classical foundationalist view that such judgments derive their justification from the mere fact that they are judgments that attribute an observational characteristic to a perceived object on the basis of the presentation of that characteristic in one’s sensory experience. The problem will be that this seems wrong for the cases where one’s judgment does fit the observational characteristic of the perceived object *and* its presentation in one’s sensory experience, while nevertheless *one is generally inept in making such judgments*, and rarely gets them right.

fyingly coherent and deeply explanatory account of the justification even of these thoughts will invoke the exercise of reliable cognitive virtues. After all, the use of such indexical and phenomenal concepts comes with a built-in guarantee of reliability. The use of such concepts may hence be seen as the exercise of corresponding reliable cognitive virtues. If in other departments of our thought we need appeal to such virtues in order to explain epistemic justification and knowledge, therefore, it may be that the deepest and most general account of the classical foundationalist justification pertaining to indexical and phenomenal thought will invoke the more generally powerful idea of reliable cognitive virtue, with indexical and phenomenal virtues falling into place as special cases.

It may be argued, however, that classical foundationalism *can* after all sketch an account of the foundational status not only of some introspective indexical and phenomenal beliefs, but also of some introspective SGA beliefs. We need only add to our account a further clause requiring that the believer “attend” to the SGA feature F. An intriguing idea,⁸ this recalls our invoking attention in the conditions required if one is to pick out an aspect of one’s experience by saying or thinking that one is experiencing “thus.” Somehow one must zero in on the target aspect of one’s experience, to be picked out as “experiencing thus” – which is to “attend” to that aspect.

Are foundationally justified beliefs perhaps those that result from attending to our experience and to features of it or in it?

Consider again the hourglass eleven-membered pattern of dots. Suppose you have a grasp of a corresponding phenomenal concept (which requires that you would have been able to recognize this pattern if it had occurred suitably within your experience), and that you would have been able to detect its absence.⁹ This is compatible with your not having counted the dots, however, while by assumption only counting can make you justified in believing that there are eleven. Your phenomenal concept of that eleven-membered array is then *not* an arithmetical concept, and its logical content will not yield that the dots in it do number eleven. (If the hourglass pattern with eleven dots does

8 This was suggested as worth considering by Richard Feldman in his commentary on an earlier version of this material, which was delivered at the APA Eastern Division meetings of 1999.

9 This second requirement is by the way particularly stringent and difficult for those who appeal to graspings of features *as featured* in one’s experience; when a feature is absent it cannot be grasped thus, and hence our correct denials of the concept in those cases will not be amenable to the same kind of explanation of our reliability.

not establish the point, we can substitute an hourglass pattern with nineteen dots, or one with twenty-nine, etc.)

What is worse, one *can* “attend” to a pattern in one’s experience without having a phenomenal concept of it. Suppose you see a well-lit white decagon against a black background in an otherwise darkened room. You can attend to that shape, you can focus on it and see it stand out clearly from its background. All of that you *can* do despite lacking a phenomenal concept of that property, which would require at a minimum that you would have been able to spot other exemplifications of it in appropriate conditions, and would have been able to spot its lack as well. You may simply lack this sort of ability; thus, you may be unable to attach a label to the pattern in such a way that you would have been able to apply that label with systematic correctness in enough otherwise dissimilar cases where the pattern was also present.

Consider now a *modified classical foundationalism* as follows:

An SGA belief that one’s experience has feature F (an SGA feature) is foundationally justified so long as (a) one’s experience does have feature F, (b) one believes that one’s experience has feature F, and (c) one attends to feature F.

Our example of the decagon shows this to be inadequate. One would not be foundationally justified in believing the white figure in one’s visual field to be a decagon, despite it’s being one, and despite one’s then “attending” to that shape as it appears in one’s visual experience.

7.4 Fumerton’s Acquaintance Theory

In his insightful defense of classical foundationalism, Richard Fumerton makes a further proposal. He proposes a special relation of acquaintance which he explicitly distinguishes from any intentional propositional attitude. He then writes:

If my being acquainted with the fact that P is part of what justifies me in believing P and if acquaintance is a genuine relation that requires the existence of its relata, then when I am acquainted with the fact that P, P is true. The fact I am acquainted with is the very fact that makes P true. The very source of justification includes that which makes true the belief. In a way it is this idea that makes an acquaintance foundation theory so attractive. I have no need to turn to other beliefs to justify my belief that I am in pain because the very fact that makes the belief true is unprob-

lematically before consciousness, as is the correspondence that holds between my thought and the fact. Again, everything one could possibly want or need by way of justification is there in consciousness.¹⁰

A page earlier he had required for the non-inferential justification of one's belief *B* with content *P*, that one be acquainted with three things: first, one's belief *B*; second, the fact that *P*; and, third, the correspondence holding between that belief and that fact. Now we must ask what Fumerton means by acquaintance. Is it enough for one to be acquainted with an item that it figure in one's consciousness, perhaps at its surface? If so then it would seem that one would be automatically acquainted with one's occurrent beliefs and with one's conscious sensory experiences.

Take a case where one sees a black figure against a white background. One sees an image, a triangular image *I*, say. So there is one's experience *E* of seeing that image just then. In addition suppose one also has an occurrent belief *B* that one sees that triangular image. Both *E* and *B* figure in one's consciousness, then, at its surface. In addition, there is the fact that *B* corresponds to *E*, at least in the sense that *E* makes *B* true. Does this also figure in one's consciousness? Well, if one had two images *I* and *I'*, both triangular, would it not figure in one's consciousness that the two are isomorphic, or at least would not their *shape-sameness* be given? That would seem to be also constitutive of one's consciousness at the time. And, if so, it could also plausibly be held that the correspondence of *B* to *E*, both items in one's consciousness at a given time, would also figure in one's consciousness at that time. And so, if it is enough, for one to be "acquainted" with an item, that it figure in one's consciousness (at its surface, perhaps) at the time, then all that is required for one to satisfy Fumerton's three conditions is that *E* and *B* figure in one's consciousness and that *B* correspond to *E*. But this runs against the problem of the speckled hen.

7.5 Experience, Concepts, and Intentionality

Might things look different to two people in otherwise similar perceptual circumstances simply because one has a phenomenal concept that the other lacks? This is trivially true in intellectual senses of "looks," as for example in a sense involving simply an inclination to believe (an inclination requiring possession of the constitutive concepts). The

10 *Metaepistemology and Skepticism* (Lanham, Md.: Rowman and Littlefield, 1995), p. 76.

more interesting question is whether, *in a sensory sense of “looks,”* possession of a phenomenal concept could make a difference to how things “look.” If so, it may be claimed, accordingly, that if one lacks a certain phenomenal concept then one’s experience will *not* in fact have the feature corresponding to that phenomenal concept. Anyone who *does* have a given phenomenal concept, whose experience has moreover the corresponding feature, plausibly *will* believe, and believe reliably, that it does, since part of what is involved in possession of such a phenomenal concept is the ability to tell when the corresponding feature *is* present in one’s experience.

If that is so, then phenomenal concepts admit no gap of the sort that gives rise to the speckled hen problem. No such phenomenal feature could possibly characterize the experience of a subject lacking the relevant phenomenal concept. But once in possession of the phenomenal concept while having an experience with that feature, one could not possibly miss that feature, given the conditions that one must satisfy in order to so much as possess such a concept.

How plausible is it that those who have and those who lack a given phenomenal concept differ in the way specified? Here now are some arguments against that dubious idea.

Suppose I am presented with an appropriately-sized, well-lit regular decagon that stands out from a black background, and suppose I focus on this image at the center of my visual field. As it happens, I myself *lack* any such phenomenal concept of a decagon, so that, for example, I am unable to tell decagons by simple inspection. Yet that image surely has a definite shape in my visual field. But what shape could that be other than that of a decagon? If that *is* the shape of the image in my visual field, finally, does it not follow that the image looks sensorily to me like a decagon.

That is a first argument; here is a second. Consider the following arrays:

```

• •
• • •
• • • •
• • • • •
• • • • • •
• • • • • • •
• • • • • • • •
• • • • • • • • •
• • • • • • • • • •
• • • • • • • • • • •

```

Your cardinality judgments for the top half or so of these horizontal arrays are quite reliable even before you count, unlike your judgment for the bottom half (that's how it is for me, anyhow, and we could surely modify the example so that it is also true of just about anyone – except maybe one of Oliver Sacks's idiot-savants).

Consider now the images of one-line horizontal arrays that you have as you stare at that large pattern. One might think that the top half or so of these imaged arrays have a determinate cardinality, unlike those towards the bottom. But is it not determinate that each array has one more dot than its predecessor as we move down from the top to the bottom array? As we consider any pair of proximate imaged arrays, can we not see clearly that the successor has one more member than its predecessor? And, if so, does it not follow that if (a) the first array has, determinately, two dots, and if (b) each array has, determinately, one more than the preceding array, then (c) the tenth array must have, determinately, eleven dots?

We could also reason similarly with the imaged arrays corresponding to the following fragment of our pattern:

```

. . . . .
. . . . .
. . . . .
. . . . .
. . . . .

```

Here the reasoning would be as follows:

- 1 The top imaged array has six dots.
- 2 There are five imaged arrays from top to bottom.
- 3 Each imaged array has one more dot than its predecessor.
- 4 Therefore, the fifth array must have ten dots.

It may be replied that experience is intentional in a way that makes our argument questionable. Thus one might experience that there are lots of speckles, more than a dozen, without experiencing that there are specifically 48, or any other nearby number of dots. Here again the aim is to preclude the possibility of the sort of gap that gives rise to the speckled hen problem. The aim is to block the possibility that one have an experience of a certain specific sort while mischaracterizing that experience in one's beliefs about it. Armed with the claim that experience is intentional, one could defuse the speckled hen problem

by holding that the hen image has many speckles *without* having any particular number of speckles. Even if experience is indeed intentional, however, my argument still goes through in some form, or so it seems to me. For someone could *still* believe that the hen image *does* have 47 speckles. Now our reckless subject would not be wrong because the image actually has 48 speckles or any other specific number of speckles. But he would be wrong in any case, simply because what he believes about the image fails to be true about it (even if it is not quite false, but only indeterminate, or whatever).

In any case, I will now try to turn the argument from intentionality on its head, by showing that, even at the heart of the intentional, the argument against classical foundationalism still goes through.

Obviously one might think that there is a triangular item in a box B, without thinking that there is an isosceles one, or an equi-angular one, or an equilateral one, etc. Nevertheless, intentional facts can still be metaphysically determinate. What is determinately so in the case before us is that I am “predicating” or “invoking” the property of triangularity. Such determinacy suffices for my main point, as suggested by the following, which moves beyond sensory experience to occurrent thought. My main point about the fallibility and unreliability of introspection can be made about as plausibly in terms of occurrent thought as in terms of experience.

Suppose you are having right now the occurrent thought:

(T1) That if squares have more sides than triangles, and pentagons more than squares, and hexagons more than pentagons, and heptagons more than hexagons, then heptagons have more sides than triangles.

Presumably one can affirm that complex conditional occurrently, which requires that it be the content of one’s affirmation at the time one affirms it. Compare this occurrent thought:

(T2) That if squares have more sides than triangles, then squares have more sides than triangles.

We can know about each of these that it is a conditional thought, and how many atomic conditions make up its antecedent, and how many geometric figures are specified in each. Take just the last of these: how

many geometric figures are specified. Most of us can know this about T2 immediately, but to know it about T1 we would need to count. Consider now a nonchalant judgment about one's T1 occurrent, conscious thought, unaided by counting. Is there not a way in which this judgment is unreliable and therefore unjustified while the corresponding judgment about T2 is not?

Something similar could be done with any thought that involves a multiplicity of properties (or concepts, depending on your view of thought). Thus you might right now occurrently think this:

(T3) That if everything blue is spherical, and if everything white is cubical, and if everything red is ellipsoid, and if nothing is spherical, cubical, and ellipsoid, then nothing is red, white, and blue.

Now quick: How many (nondisjunctive) color or shape properties are you thinking about? Can you tell right off, without counting? Can you make an honest mistake? But isn't there a determinate number of color or shape properties that your (intentional) thought is about, even before you count? Otherwise how to attribute to your thought a given logical form, with a certain content involving the predication of certain properties, etc.?

Alternatively, we might consider the sort of thinking that is involved in simply consciously *considering*, without necessarily assenting to, a certain proposition. And we might now wonder how many properties (or concepts) of a certain sort figure in that proposition. The reasoning above would then be applicable here again. In some cases we could tell readily, without counting, but in other cases we would need to count or to use some other inferential process.

Our anti-givenist argument is not restricted to experiences, then, but may be framed with similar plausibility in terms of thoughts. Some intrinsic features of our thoughts are attributable to them directly, or foundationally, while others are attributable only based on counting or inference. How will the classical foundationalist specify which features belong on which side of that divide? It is hard to see how this could be done without appeal to cognitive virtues seated in the subject. For example, an attribution of a feature to an experience or thought is perhaps foundationally justified only when it derives from the operation of a reliable virtue or faculty. This may then yield the important difference between the claim about our thought T2 that it is about two

shapes, on one hand, and on the other the claim about our thought T1 that it is about seven.¹¹

7.6 The Evans/Peacocke Proposal

Classical foundationalism focuses on how conscious states themselves rationally ground introspective beliefs. On this point Christopher Peacocke agrees, as he draws on his broader view about conditions for concept-possession, and in particular conditions for the possession of the concept of belief.¹² Appeal to such conditions can illuminate how the occurrent belief that *p* can so appropriately prompt the self-attribution of that very belief. The possession of the concept of belief itself requires such self-attribution. However, don't we need some reason at least to believe that in self-ascribing beliefs we are not likely to go wrong? Michael Martin raises a related worry:

[We] . . . can see Peacocke's account as implicitly offering to delimit in part the range of . . . beliefs about which we have authority: namely, those which we can ascribe on the basis of conscious thoughts. But . . . this is too broad a basis for authoritative self-ascription. . . . A too familiar phenomenon in philosophical enquiry is the realization that the proposition that one felt that one had a real conviction was so, turns out to be not the proposition that one identified initially, but really another proposition, which is a close cousin.¹³

Martin goes on to appeal, with Peacocke, to insights due to Gareth Evans, who writes: "The crucial point is [this] . . . : in making a self-ascription of belief, one's eyes are, so to speak, or occasionally literally, directed outwards – upon the world."¹⁴ According to Martin, there is perhaps a moral to be drawn from Evans's approach:

11 At least that seems to be so except perhaps for savants who can tell directly the number of figures involved in thoughts like T1 or even in those involving many more figures than that.

12 Christopher Peacocke, "Conscious attitudes, attention, and self-knowledge," in C. Wright, B. C. Smith, and C. MacDonald (eds.), *Knowing Our Own Minds* (Oxford: Oxford University Press, 1998).

13 Michael Martin, "An eye directed outward," in C. Wright, B. C. Smith, and C. MacDonald (eds.), *Knowing Our Own Minds* (Oxford: Oxford University Press, 1998), p. 116.

14 Gareth Evans, *The Varieties of Reference* (Oxford: Oxford University Press, 1982), p. 225.

When the subject has her eyes directed outwards, they are directed on the world, as it is for her. Now, we should not assume that all of a subject's conscious states form part of the world as it is for her, not even all of her conscious beliefs. So, we might say, what we expect is that a subject should at least have authority with respect to those mental states which comprise in part her point of view on the world. Peacocke himself endorses a close link here between consciousness and the subjective point of view when he claims that "The requirement that the reason-giving state is one which is, or could become, conscious is intimately related to our conception of an agent as someone with a point of view, one whose rational actions make sense to the subject himself . . . given that point of view".¹⁵

Despite the insights contained in this approach, it cannot be a complete account, not even of the restricted field of states that make up one's subjective point of view. A gaping deficiency in the account suggests a different approach, one that will not try to explain the authority of beliefs about one's current mental states merely by noting that they are rationalized through the support of those very states.

The gap concerns our knowledge that we do *not* consciously believe such and such, and our knowledge that we do *not* consciously intend such and such, and our knowledge that we do *not* seem to see anything red or any triangle. In no such case is our knowledge to be arrived at by means of first-order inquiry or deliberation of any sort; nor is our knowledge based on some conscious state that rationalizes it. On the contrary, it is precisely the *absence* of a relevant belief or intention or experience that makes it a case of knowledge. Moreover, it is doubtful that our belief in the *absence* of the state must be based on any conscious state to which it is appropriately responsive.

We must distinguish between believing that not-*p* and not believing that *p*, between desiring that not-*p* and not desiring that *p*, between experiencing as if not-*p* and not experiencing as if *p*. Nevertheless, we enjoy privileged access to both sides of each such pair. The more general problem for the approach in terms of an intermediate conscious state is this: We enjoy privileged access both to what is present in our consciousness and to what is absent therefrom. What explains the one is unlikely to differ dramatically and fundamentally from what explains the other. So, even if the account of what is present to consciousness must invoke some relevantly distinguished conscious states that help rationalize our pertinent beliefs, something deeper must

15 Martin, "An eye directed outward," p. 117.

explain what is held in common between that case of privileged access and our equally privileged access to what is absent from consciousness. This is a problem for classical foundationalism, since, again, it seems unlikely that the explanation in terms of taking the given goes deep enough if it is powerless to explain our foundationally justified beliefs about what is *absent* from our consciousness at the time, and yet it seems unlikely that the explanation of that foundational justification would be entirely other than the explanation of our foundational justification for foundational beliefs about what is then *present* in our consciousness.

7.7 Foundational Knowledge

We have been offered various accounts of how beliefs can have *foundationally* justified status through sticking to the character of the subject's own conscious experience at the time. None of these has been successful, or so I have argued. Although discussing earlier failures can be illuminating, here I would like to sketch instead a positive view that seems more promising.

What distinguishes the case of 48 speckles where one guesses right, and does not know, from the case of 3 speckles, where one does know foundationally? We need to appeal not just to

- 1 the specific property of the experience, its containing however many speckles, say, and
- 2 the propositional content of the occurrent thought as one judges the image to contain that many speckles, and
- 3 the match between that property of the experience and that propositional content – such that the propositional content predicates that very character of the experience.

For in the case of the 48-speckled image, where one guesses right in taking the image to contain that many speckles, all of those conditions are met, and the judgment does fit the character of the image. Yet one fails to know by acquaintance, and one fails even to be justified.

It is hard to resist the conclusion that what we require, in addition to 1–3, is a further appeal to the following:

- 4 some causal or counterfactual connection between the character of the experience and the propositional content of the judgment.

This is abetted by the thought that if the judgment (with its content) is to be rationalized by the experience (with its relevant character), then the former must be appropriately responsive to the latter, in such a way that variations in the latter would have led to corresponding variations in the former. Or perhaps it will suffice for appropriate responsiveness that one might easily not have believed this without that belief's being then accompanied by one's experiencing in a corresponding way.

Accordingly, what seems required is that one's judgment about the phenomenal character of one's experience be appropriately causally or counterfactually (and reliably) related to the character of the experience. If this is right, it is fascinating to find at the heart of givenist, internalist, classical foundationalism a need for the sort of relation so often used by its externalist opponents, over the course of recent decades, for their various externalist alternatives. Some have invoked straight causation of a belief by the fact that is its content, others a requirement of non-accidentality, others a counterfactual tracking requirement, and yet others a requirement of reliable generation and sustainment of the belief. These have been proposed mostly as requirements that a belief must satisfy in order to qualify as knowledge, whereas our main focus here has been on how a belief gets to be epistemically justified. But there is a close connection between the two concerns, though this has yet to be spelled out in satisfactory detail, given the unclarity and ambiguity in the relevant terminology of "epistemic justification."

I will conclude with a sketch of the sort of requirements that seem to me most promising. How then would one distinguish

- 1 an *unjustified* "introspective" judgment, say that one's image has 48 speckles, when it is a true judgment, and one issued in full view of the image with that specific character,

from

- 2 a *justified* "introspective" judgment, say that one's image has 3 speckles?

The relevant distinction is that the latter judgment is both *safe* and *virtuous*, or so I wish to suggest. It is "safe" because in the circumstances not easily *would* one believe what one now does in fact believe, without being right. It is "virtuous" because one's belief derives from a way of forming beliefs that is an intellectual virtue, one that in our

normal situation for forming such beliefs would tend strongly enough to give us beliefs that are safe.

One does not know nor is one so much as justified foundationally in taking one's image to contain 48 speckles even if one's image *does* in fact contain 48 speckles, so that one's belief corresponds precisely to what is then given in one's consciousness. One falls short in that case because of how easily one might have believed that one's image had 48 speckles while it had one more speckle or one less. But that is not so for the belief that one's image has 3 speckles.

It is not sufficient, however, that one's belief be thus *safe*. Consider the belief that $(2^2)^2 = 16$. Not easily would anyone believe this without being right, since not possibly would anyone believe it without being right. Nevertheless, if one derives and sustains that belief only by means of a procedure which assumes that $(x^n)^n = x^{n+n}$, then one fails to know, despite one's belief's being perfectly "safe," in the sense defined.

For that reason, knowledge requires one's belief to be not only safe but also virtuously sustained, through the use of a reliable ability or faculty, through an "intellectual virtue."¹⁶

An opponent could now take a different tack, by arguing against what he perceives as a presupposition of our reasoning. Thus it might be argued that we are missing the point, that *justification* is not just whatever must be added to true belief in order to attain knowledge. Justification may be said to answer to its own requirements, and to have its own separate intuitive basis, one more closely allied to concepts of reasonableness and rationality. Thus what one is assessing in calling a belief justified is rather the worth of the mind of the believer in respect of holding that belief with the basis that it has and in the internal circumstances in which it is held. Thus the relevant focus of evaluation is rather the relevant coherence of that mind, either at that moment or over the stretch of its history that is relevant to the acquisition and sustenance of that belief.

Nevertheless, my basic point remains. It is not enough that one's beliefs at that time agree in logical respects with the experiences one is then undergoing. For the belief that one has an image with 48 speckles could hardly agree better with that image's indeed having exactly that many speckles, while yet the belief remains unjustified despite such impressive coherence. What is missing here but present in the case where one's image has 3 speckles, that being exactly how many

16 In a fuller treatment these notions need to be clarified, and their virtues displayed more fully in dialectical interplay with main competitors.

speckles one believes that it has? Once again what matters is plausibly one's virtuous ability to discern cases of 3 speckles from those involving fewer speckles, or more, by contrast with one's *inability* to discriminate comparably as to 48 speckles.

In evaluating that proposal, recall, we seek to understand a source of *epistemic justification* (in a sense allied with rationality and reasonableness) that will be *foundational*, i.e., that will not derive from any inference, implicit or explicit, or at least one that will not derive wholly from that. It is a mistake, therefore, to oppose our proposal by arguing that someone *could* be justified in self-attributing an experience of 48 speckles despite lacking a virtuous ability to discriminate such experiences directly: by arguing, for example, that someone *could* be thus justified simply by inferring from a directly introspected inclination to so believe, along with a well-justified belief that such inclinations almost always turn out correct. This *would* explain a source of epistemic justification for that belief all right, even absent any virtuous ability to discriminate such experiences directly, but that justification would be inferential, not foundational.

We shall return to a fuller defense of this sketched proposal in chapter 9. But first we turn to the problematic of philosophical skepticism, to the relation between this problematic and the internalist stance in epistemology, which will lead to more hard questions for epistemic internalism.

Skepticism and the Internal/External Divide

8.1 Descartes's Paradox

“A belief is knowledge only when proof against all doubt, even the most hyperbolic” – so premises Descartes. If unable to rule out the possibility that we are demon victims (or envatted brains), therefore, we know neither what we ostensibly see, nor whatever conclusions we may inductively, and fallibly, infer from such “data.”

This skeptical argument is naturally dismissed for imposing an unreasonably high requirement. Ordinarily we require nothing nearly so stringent, but only well-justified beliefs, based perhaps on less-than-conclusive reasoning. We do not require apodictic reasoning beyond any reasonable doubt, from axioms infallibly known.¹

How superficial that dismissal seems, however, in light of the following examples. Newly-arrived in Alaska, I spot a husky as such, even though from my distance and angle no feature distinguishes it from the many wolves in that vicinity, nor am I able to rule out circumstantially that I see a wolf rather than a husky. You believe Tom stole a book because Dick tells you so, and in fact you are right, but when Dick is accused of slander you have no basis to rule that out. Thirst-crazed by the desert sun, you hallucinate an oasis where by coincidence there happens to be one just as pictured in your feverish imagination. Such cases plausibly suggest this “principle of exclusion”:

(PE) Take any hypothesis X. To know X (or that X is true), one must rule out every possibility P that one knows to be incompatible with one's knowing X (that X is true).

¹ The term “justified” is not widely applied to beliefs in ordinary discourse. It may be a philosophers' term of art, meaning: “belief not defectively formed or sustained in a way that reflects poorly on the believer's mind.” Of evaluative terms ordinarily applied to beliefs, “reasonable” comes perhaps closest to this.

None of your three beliefs – about the husky, about Tom’s honesty, about the oasis – satisfies this requirement, which may now be offered as a way to explain why none amounts to knowledge.

We may now see Descartes’s skeptical reasoning as fitting the same pattern. Suppose yourself in a situation where you know that you face a fire only if you know it perceptually (Descartes’s situation). Take now any skeptical scenario – evil demon, brain in a vat, etc. – in which, through direct stimulation of one’s brain, or through direct control of one’s experience by the demon, one has fully detailed and rich experience as if one faced a fire, although this experience is quite independent of any fire. Such a skeptical scenario is a possibility that one knows to be incompatible with one’s knowing that one faces a fire; unable to discern perceptually the case where one really faces a fire from the skeptical scenario, therefore, one lacks perceptual knowledge of any such fire, and therefore lacks such knowledge altogether in one’s actual situation, where by hypothesis one knows only if one knows perceptually.

A version of the argument may be laid out as follows (supplemented):

- 1 All your sensory experience and information at t is compatible with your dreaming at t .
- 2 So you need some test “indicating” that you are not then dreaming, where you know that both (a) it is satisfied and (b) if it is satisfied then you are not dreaming.
- 3 But how could such a test ever be available to you, if it is a condition of your knowing (perceptually) *anything* beyond your experience that you know yourself *not* to be dreaming? If so, you could not know perceptually that you satisfy the test itself unless you (already?) on some occasion knew yourself not to be dreaming. And how could you know yourself to satisfy the test except, at least in part, *perceptually*? Besides, how could you know that if the test is satisfied then you are awake, except by relying on being awake (and not dreaming) either at that time or at some earlier times (when you gathered the data, presumably observational, enabling your conclusion that in general one satisfies the test only when awake)?

This nowhere assumes that in order to know that p one must be *absolutely* and justifiably *certain* that p . We are not setting the standards too high and then complaining that we can’t possibly measure up. Our argument is designed to show rather that, when compared with

the possibility that we are just dreaming, our thought that we really see is based on *no good reason whatever*. A sufficient reason must enable us to rule out the alternatives that clearly would preclude our knowing, as above. The dismissal of the skeptic as setting the bar too high is superficial because it overlooks this argument.

8.2 Epistemic Externalism and Internalism

A famous dialectic consists of three theses:

- D (a) Knowing that here is a hand requires knowing that one is not dreaming.
- (b) I know that here is a hand.
- (c) I do not know that I am not dreaming.

Moore and the skeptic both accept D(a), but then part ways, the skeptic rejecting D(b), while Moore rejects D(c) instead. Given Moore's internalist, givenist proclivities, it is not surprising that he accepts D(a). Contemporary externalists share no such proclivities, and many join in rejecting D(a) along with other traditional assumptions.² How can one be justified in one's belief about the hand, however, except by relying on an assumption that one is awake and not dreaming? And how could this assumption ever get justified? It is not something that could be known just a priori, nor is it directly introspectable. It pertains rather to a contingent causal relation between oneself and one's surroundings. What could be one's basis for thinking that this relation holds? If it is not known a priori and is not known through any kind of direct perception or introspection, then how is it known? More specifically, if one's epistemic justification is not of the armchair variety, nor of the introspective or perceptual variety, then how *is* one justified in believing oneself causally connected with one's surroundings in the ways required for epistemically justified beliefs, e.g. about how it is out there on the basis of how it seems from here? This prompts questions about the nature and status of such epistemic justification. What reason might be offered in support of an internalism such as Moore's? Here are two familiar supports.

² Let "dreaming" here stand for "philosophically dreaming" defined as follows: *S* philosophically dreams that *p* iff *S* experiences as if *p*, but *unveridically* so (and the same goes for all of *S*'s sensory experiences at the time).

Cartesian internalism of justification:

Justification requires only really proper thought on the part of the subject: if a believer has obtained and sustains his belief through wholly appropriate thought, then the believer is justified in so believing – where the appropriateness of the thought is a matter purely internal to the mind of the subject, and not dependent on the environment.³

Chisholmian internalism of justification:

The “concept of epistemic justification . . . is *internal* . . . in that one can find out directly, by *reflection*, what one is justified in believing at any time.”⁴

The appropriateness of one’s thought is “purely internal to one’s mind” only if what makes and *would* make (would make by necessity if ever it occurred) one’s thought appropriate – that *in virtue of which* it is and would necessarily be appropriate – involves only matters “internal to one’s mind.” And these are matters constituted by one’s mental properties, including relations intrinsic to the mind, as when one introspects a headache, and also “propositional attitudes.”

How are these forms of internalism related?

Thesis:

That Chisholmian internalism follows from Cartesian internalism, given only the following assumptions (the “accessibility of the internal,” AI; and the “accessibility of epistemic supervenience,” AS):⁵

- AI Matters purely internal to the mind, such as occurrent mental properties and propositional attitudes, are always open to discovery by reflection.
- AS If a belief attains, and would necessarily attain, a particular epistemic status in virtue of its non-epistemic properties, that it would do so is open to discovery by reflection.

3 Properties and propositions, and necessary relations binding them, are meant to count, not as part of the subject’s “environment,” but as matters themselves “internal.” Perhaps one should read “internal” here as “nonexternal.”

4 Roderick Chisholm, *Theory of Knowledge*, 3rd edn. (Englewood Cliffs, NJ: Prentice-Hall, 1989), p. 7.

5 The demonstration is left to the interested reader. By definition, “reflection” involves either (a) introspection, (b) rational intuition, (c) memory, (d) deduction, or (e) induction or ampliative reason which builds only on materials provided by (a)–(d).

Here again is how Chisholm himself sees the matter:⁶

The usual approach to the traditional questions of theory of knowledge is properly called “internal” or “internalistic.” The internalist assumes that merely by reflecting upon his own conscious state, he can formulate a set of epistemic principles that will enable him to find out, with respect to any possible belief he has, whether he is *justified* in having that belief. The epistemic principles that he formulates are principles that one may come upon and apply merely by sitting in one’s armchair, so to speak, and without calling for any outside assistance. In a word, one need consider only one’s own state of mind.

What might justify this assumption? Why assume with internalism that unaided armchair reflection could always reveal whether or not one is justified? An answer to that question is needed since, for one thing, neither assumption AI nor assumption AS seems just obvious on inspection. Take AI. Even if individual mental states are always transparent to reflection, which may be doubted, complex combinations of them might not be transparent. As for AS, who knows how complex the ways of supervenience might be, or how open to discovery by our limited minds?

8.3 Chisholmian Internalism

What does it take for a belief to be epistemically justified? What sort of status does a belief have in being epistemically justified? Deontology, (epistemic) justification, and internalism are said to be closely connected, and Chisholm’s internalism has been explicated by means of these connections.⁷ By “internalism” we shall mean mostly Chisholmian, epistemic internalism, i.e., the view that we have special access to the epistemic status of our beliefs (or at least to their status of being justified), special access by means of armchair reflection.

According to the deontological view of epistemic justification, such justification consists in an appropriate relation to one’s epistemic obligations or duties. Thus:

6 Chisholm, *Theory of Knowledge*, 3rd edn., p. 77.

7 Cf. Alvin Plantinga, *Warrant: The Current Debate* (Oxford: Oxford University Press, 1993), chs. 1–3. I discuss Plantinga’s critique of Chisholmian internalism in “Plantinga on epistemic internalism,” in J. Kvanvig (ed.), *Warrant in Contemporary Epistemology* (Lanham, Md.: Rowman and Littlefield, 1996).

DEJ One is epistemically justified (and not blameworthy) in \emptyset 'ing if and only if in \emptyset 'ing one does one's epistemic duty – or, at least, that is so for a large, important, and basic class of epistemic duties.

The plausibility of DEJ depends on how we interpret “epistemic justification.” On a strong deontological interpretation:

EJ S is epistemically justified at t in believing that p IFF S at t freely opts to believe that p , in the knowledge that he ought then to believe that p .

Even on this conception, however, where's the internalism? For example, how do we arrive at the result that when one knows what one ought to believe one knows it by reflection? (Only if one can know this by reflection, presumably, can one know by reflection that one *is* epistemically justified.)

One might appeal here to the intuition that the beliefs of the evil demon's victim would be no less justified than our ordinary beliefs, and on the same basis. But the victim must be justified in virtue of properties *ontologically* internal to him. And since for any actual thinker there is such a victimized double, therefore the justification of actual people must also derive from properties internal to them.

If the victim knows himself to be justified, he must know this by knowing of certain ontologically internal properties of his. But why must this knowledge be by *reflection*? Well, “reflection” usually means: “some combination of introspection and memory, along with intuition and inferential reason.” It is not unreasonable, then, that if the victim knows what makes him justified by knowing of certain ontologically internal properties of his, then he must know it through these reflective faculties. Even if the argument is not absolutely airtight, it does seem persuasive.

This line of reasoning from a deontological conception of epistemic justification to an epistemic internalism *assumes* ontological internalism. It assumes both that the demon's victim would be no less and no more justified than his flesh and blood counterpart, and that any justification enjoyed by the victim would derive at most from his internal properties. Only on such assumptions can we thus derive a doctrine of epistemic internalism: the doctrine that one could always know *by reflection* about one's own epistemic justification. On this line of reasoning, therefore, it would not be true that one could just start with a

deontological conception of epistemic justification, and derive *either* or *both* of epistemic internalism *and* ontological internalism from that. On the contrary, only by first assuming ontological internalism can one then derive epistemic internalism from a deontological conception of epistemic justification.⁸ What is more, we still face (1) how unobvious it is that one's mind should be *wholly* transparent to one's own reflection, and (2) how unobvious it is that the supervenience base of one's epistemic obligation must always be simple enough to be accessible through unaided reflection.

Moreover, the conception of epistemic justification as deontological has further limitations. Most epistemologists, of whatever persuasion, externalist or not, will surely grant that for a large and important class of cases, what has been claimed on behalf of the internalist is likely to be true. Take any case where one knows something or one is epistemically justified in believing something purely on the basis of reflection. Many epistemologists, externalist or not, will agree that in any such case one *can* know internalistically, by reflection, that one's belief has its positive epistemic status. Real disagreement concerns rather knowledge that is either perceptual or depends essentially on the perceptual. Of course, even here one must presumably be able to know or believe correctly, simply on the basis of reflection, that one's belief is justified, but now the argument will have to rely on an independently held *ontological internalism*; as was argued above.

The conception of epistemic justification as deontological is also limited in other ways. The problem is not so much that there isn't such a notion, nor is it that the scope of such epistemic justification is too narrow, given our limited control over what we believe. These are problems, I believe, but my main problem here is rather this: Does the defined notion of "epistemic justification" capture an epistemic status of primary interest to epistemology, *even to "internalists" generally*? Don't we fall short in our effort to throw light on internalist epistemic justification if we stop with the explicated notion of deontological justification? After all, there are other desirable statuses for our beliefs to attain. Thus it is good that one's beliefs not derive from a deliberate effort on one's part to believe what is false, from a sort of "epistemic masochism." And it is a good thing that one's beliefs not derive from

8 Actually, it is not ruled out that ontological internalism receive some independent support; the important point is that, so far as we have been able to determine, it seems needed as a prior premise for any discernible argument that moves from deontologism of justification to epistemic internalism.

uncaring negligence. These are moreover sorts of things that one could know about the derivation of one's beliefs on the basis of armchair reflection. So one could secure internalism, and define a sort of "epistemic justification" even without invoking full-scale epistemic deontology. But surely we would then fall short in clarifying matters epistemic. Avoiding epistemic masochism and uncaring negligence are only two of the things involved in epistemic excellence. And a similar question could now be pressed against the deontologist: In opting for a deontological conception of epistemic justification, are we capturing all relevant aspects of "internal" epistemic excellence that we wish to illuminate? Or are we falling short in ways analogous to the ways in which we fall short if we stop with mere masochism-avoidance or negligence-avoidance conceptions of justification?

If we focus on the concept of epistemic justification defined, I believe we can see that it will inevitably miss respects of epistemic excellence, of internal epistemic excellence.

Consider our correct beliefs about what duty requires in various circumstances, and our correct beliefs about the particulars of our own situation that determine the call of duty for that situation. These beliefs must also be assessable as epistemically justified or not. What is more, one could hardly attain deontological *epistemic* justification if the duty-specifying beliefs that govern one's epistemic choice are wholly *unjustified*. So, these beliefs must themselves be epistemically justified in some sense. But, on pain of vicious circularity, the sense involved cannot be that of deontological epistemic justification. It would seem viciously circular to define epistemic justification as abiding by rules that one is *epistemically justified* in accepting.

Secondly, what are we to say of beliefs that do *not* result from a knowledgeable choice on the part of the believer? Some beliefs one either could not have avoided, *or* could not have known to be wrong beliefs to hold in the circumstances. Either way one bears no responsibility for one's lack of pertinent control or knowledge or correct belief. This could perhaps happen even with wholly internal beliefs that are introspective or reflective and independent of perception. If we restrict ourselves to deontological justification, then all such beliefs turn out to be indistinguishable in respect of justification. None can be said to be justified.

That consequence is hard to accept. An unavoidable belief can amount to certainty of the highest grade: the belief that $1 + 1 = 2$, for example, or the belief that one has a headache when one has a migraine, beliefs one cannot help having – and the great class of such

obvious, non-optional beliefs. Compare the convictions of someone brainwashed, however, or the beliefs of a *naif* with a crude conception of what justification requires in a certain ambit, who acquires his beliefs in ways that he is convinced are methodologically sound, simply because he was raised in a culture where such ways are instilled, so that now, through no fault of his own, our *naif* does not properly know what to believe and what not to believe in that ambit.

Is there a sense of justification in which these two classes of belief are on a par? If so, that can hardly be a sense of main importance for epistemology. Clearly, we must go beyond our concept of deontological justification, if we wish to explicate the fact that some indubitable beliefs are epistemically justified while other beliefs are not, despite being equally indubitable at least to the brainwashed or naive. The claim is *not* that there is *no* sense in which both sets of beliefs are on a par in respect of internal justification. The claim is *rather* that there is likely some further important sense of internal justification that remains uncaptured by our notion of deontological justification. With this notion we fall short, in a way in which we would fall short in our effort to explicate internal justification if we were to stop with just the notion of belief non-negligently acquired and sustained. The broader notion of internal justification would need to take account both of the avoidance of negligence and of the securing of deontological justification. Not that these would have no relevance, each in its own way, to internal epistemic justification. But, as we have seen, the broader notion would need to go beyond these nonetheless. The contours of that broader notion have yet to be traced.⁹ If we agree that a belief can be irresistible without being epistemically justified, however, then we should be willing to hold open the possibility, worth exploring, that there is such a broader notion, even if its contours remain obscure.

Again, the objection here is *not* that there is anything intrinsically wrong with the definition of epistemic justification as deontological. Some such notion can surely be defined in some such way as those explored above. The doubt I have voiced is rather about the capacity of this notion to cast a powerful light on an interestingly *epistemic* status to which our beliefs can aspire. My point is that any such notion will fall short as we have seen the defined notions of epistemic “masochism” and epistemic negligence to fall short.

9 Some suggestions towards this project may be found in chs. 8 and 16 of my *Knowledge in Perspective* (Cambridge: Cambridge University Press, 1991).

8.4 Justification and the Internal

Compare yourself with a counterpart victim of the evil demon. Suppose the two of you indistinguishable in every current mental respect whatsoever: If you are having a current sensory experience, so is your counterpart; if you have a certain belief, so does your counterpart; if you would defend your belief by appeal to certain reasons, so would your counterpart; and vice versa. The two of you are thus point by point replicas in every current mental respect: not only in respect of mental episodes, but also in respect of deeply lodged dispositions to adduce reasons, etc. Must you then be equally epistemically justified, in some relevant sense, in each such belief that by hypothesis you share? If either of you is epistemically justified in believing that you face a hand, must the other be equally justified in so believing? What could a difference in justification derive from? Each of you would have the same fund of sensory experiences and background beliefs to draw upon, and each of you would appeal to the same components of such a cognitive structure if ever you were challenged to defend your belief. So how could there possibly be any difference in epistemic justification? All of this agrees with Chisholm's internalism, as we have seen. For Chisholm, justification is a matter of the rational-cum-experiential structure of one's mind. It is constituted by one's experience together with one's dispositions to respond in certain ways under sustained Socratic dialectic.

According to Chisholm, for every epistemically justified belief, there must be a rational structure whose presence in the subject's mind yields that justification. This rational structure may be brought to light through a process of Socratic dialectic, which will press for the reasons one might have for believing as one does, and for the reasons for those reasons, and for deeper reasons in turn, until one reaches a bedrock where one can only appropriately repeat oneself, saying "What justifies me in believing that p is just the fact that p ," or the like. In other words, subjects alike in their dispositions to respond to Socratic dialectic, to demands for a display of one's justification, will be alike in their attained justification (so long as their bedrocks are equally secure). But your demonic counterpart is imagined to be just like you precisely in all such underlying rational structure: you each would adduce the same reasons in response to challenges (and your experiential bedrocks are equally secure). Each of you will hence share the same epistemic justification for your corresponding beliefs. Call this "Chisholm-justification," a sort of present-moment justification.

Compare this: Mary and Jane both arrive at a conclusion *C*, Mary through a brilliant proof, Jane through a tissue of fallacies. Each has now forgotten much of her reasoning, however, and each takes herself to have established her conclusion validly. What is more, each of their performances is uncharacteristic, Jane being normally the better logician, Mary a normally competent but undistinguished thinker, as they both well know. The point is this: Jane would seem currently only better justified in taking herself to have proved *C*, as compared with Mary. As of the present moment, therefore, Jane might seem as well justified as is Mary in believing *C*. We know the respective aetiologies, however; what do we say? Would we not judge Jane's belief unjustified since based essentially on fallacies? If so, then a belief's aetiology can make a difference to its justification. Call this sort of justification PA-justification ("personal-aetiology justification").

Mary and Jane are equally Chisholm-justified, then, but only Mary is PA-justified. Chisholm-justification depends just on your present ability to adduce reasons, and on your present structure of experiences and beliefs. So it depends on present-time-slice, internal faculties such as introspection and current ability to adduce reasons. By contrast, PA-justification goes beyond that to encompass also the operation of past sensory experience and introspection, along with pertinent temporally extended reasoning and the operation of memory. Although we have extended our scope to encompass also one's mental past, however, there is a clear sense in which PA-justification is still internal to the subjectivity of the subject, past as well as present. What reason might there be for preferring either of these concepts of justification, the Chisholmian present-restricted one or the past-encompassing one? We leave this question open, as we press on to further extensions.

Suppose a teacher lapses into reasoning that $(x^n)^n = x^{n+n}$, and on that basis reports that $(2^2)^2 = 2^4$. You, a schoolchild, believe accordingly, just on the teacher's sayso. Are you justified in so believing? Suppose the conditions for acceptance of testimony to be optimal: the teacher is normally exceedingly reliable, and the circumstances are otherwise unremarkable; it's just a normal school morning in every other relevant respect. Are you then well justified in your belief? You *are* Chisholm-justified, and also PA-justified. But is there any sense in which you are not as well justified as you might be? If so, a belief's justification would then seem to depend on factors beyond its present-time-slice profile, and even beyond its personal aetiology, to encompass also the quality of the information derived from testimony. This sort of justification is "social-aetiology justification," since it depends not only on the quality

of the belief's present supporting rational structure, and not even just on that together with the quality of its personal aetiology, via such faculties as memory and reasoning, but also on the social aetiology of that belief, wherein testimony may feature crucially.

How far might justification depend on the external? Consider some examples. Suppose you face a friend named "Mary" believed by you to be so named; your counterpart will then face a counterpart friend and will believe her to be so named. But what if you remember her name because she is a lifelong friend, whereas your counterpart is oblivious to having been just created by a playful demon with just that set of beliefs and experiences (and, let's say, either there is no friend there at all, or if there is one it just derives from the demon's caprice). Are you equally well justified in your respective beliefs?

Second case: You remember having oatmeal for breakfast, because you did experience having it, and have retained that bit of information through your excellent memory. Your counterpart self-attributes having had oatmeal for breakfast, and may self-attribute remembering that he did so (as presumably do you), but his beliefs are radically wide of the mark, as are an army of affiliated beliefs, since your counterpart was created just a moment ago, complete with all of those beliefs and relevant current experiences. Are you two on a par in respect of epistemic justification?

Here is a third case. You believe *C* as a conclusion of long and complex deductive reasoning, but your counterpart believes it only owing to the demon's caprice, although you both would now report having deduced *C* through complex reasoning, and each of you could now produce on demand some limited fragment of such reasoning. A relevant difference between you, compatible with your remaining perfect mental twins as of the present moment, is this: you have deduced that conclusion through a flawless proof; your counterpart is far from having done so. Are you both equally well justified?

Return now to that innermost sanctum of rationalist internalism which includes rational intuition itself. Given the human mind's liability to fall into paradox, wherein reason itself apparently leads us astray, how are we to think of the justification provided by intuition? In the throes of paradox, reason delivers beliefs B_1, \dots, B_n , but also the belief, B_{n+1} , that the earlier n beliefs are logically incompatible. At least one of these beliefs must be false. Such paradoxes show, therefore, that rational intuition leads us astray in the sort of way in which perception sometimes leads us astray, through illusion of one or another sort.

Suppose B_i to be the falsehood among your $n + 1$ beliefs. Can it still be justified in spite of that? Well, it does *seem* about as well justified as any of the other, true, beliefs in the set constituting the paradox; assent to B_i is not some special mistake that some particular subject has fallen into, through inattention or intellectual negligence or special deficiency. It is a mistake to which the human mind itself is inherently subject. It is like a perceptual illusion such as the Müller-Lyer lines, or the “bent” stick in water, or the “oasis” in the distance. Subjects unaware of the special circumstances – inexperienced children, for example – would *not* seem “unjustified” in “trusting their senses” when thus misled. Collateral information normally prevents such mistakes, but may be missing through no fault or defect of one’s own, as with the inexperienced child. So in some sense one is then epistemically justified in assenting. And assenting to the deliverances of rational intuition when in the grip of paradox seems quite analogous. But there is an alternative reaction; one is also pulled in the opposite direction. (Paradoxically?)

Recall that fallacious reasoning does not plausibly justify, even if one may be justified in believing the reasoning to have been valid and sound. And faulty or absent memory yields a similar result: plausibly, you are justified in your belief about your breakfast, whereas your counterpart is not. That is so, both for memory and for deductive reasoning, in some natural sense, *even though in the Chisholmian present-time-slice sense, despite the faults in one’s reasoning or in one’s memory, one may still be fully justified*. Apply that now to our case of paradox-enmeshed intuition. Can’t we appropriately distinguish two corresponding senses? Of course, the distinction must now be made *within the present time slice*, since there is here no time-encompassing faculty at work, unlike our earlier cases of inferential reason or memory.

Here is a way to draw the distinction. Let B_i be the one false belief contained in the paradox-constituting set $\{B_1, \dots, B_n, B_{n+1}\}$. Let B_j be a true belief in that set. B_i and B_j are alike in being, let us say, *subjectively* justified, but only B_j might be *objectively* justified. What is this difference? B_i is a belief that gives every appearance of being justified. The human mind (and not just our subject’s mind) is drawn to affirm B_i just on the basis of understanding its content. And the human mind is, let us suppose, quite reliable in what it accepts on such a basis, which is itself believed justifiedly by our subject, or so we may suppose. In such circumstances, our subject is “subjectively” justified in assenting to B_i . He takes himself to be justified, or may be supposed

to do so, in a way that coheres with his intellectual self-conception. But there is still something importantly wrong with his assent to Bi. People are presumably reliable in what they accept just on the basis of understanding it, as we are thereby non-accidentally put in touch with what is true (either because such assent is truth-constituting, or because such assent in such circumstances is reflective of what is independently true). Nevertheless, in accepting Bi specifically, one *fails* to be thus non-accidentally in touch with the truth: Bi being false, one is thereby not at all in touch with the truth. More revealingly, through Bi one fails to be appropriately responsive to the truth: in assenting to its content one is not believing something in such a way that one would not go wrong. One would not go wrong either when one's believing as one does is, in the circumstances, truth-constituting, or when one's so believing is in the circumstances truth-reflecting, i.e., reflective of a fact of the matter, a fact independent of the current state of one's mind. In neither way is one's assent to the content of Bi correlated or in step with the truth, however, since Bi is not so much as true. Bi does *not* track the truth in the sense that one would hold that belief if, and only if, it were true.¹⁰ Something goes wrong, then, with one's assent to the content of Bi, but it is something in a natural sense *external* to one's mind, something pertaining to a fact that in some way goes beyond what is determined by the contents of one's mind, because it is independent of that, either in the way the roundness of the Earth is thus independent, or in the way community rules are independent of any one mind. Nevertheless, it is not something that need go beyond the present time, for it is at this present time that Bi is false and (hence) fails to track the truth. Although the rationally intuitive belief Bi is thus different from mistaken beliefs attributable to memory or inference in that temporal respect, the three sorts of belief are also importantly alike: In each case – faulty “memory,” faulty “inference,” and faulty “intuition” – a normally reliable faculty is ostensibly operative, but is not *really* operative.¹¹ Regarding memory and inference, there is a strong pull in opposite directions: towards the verdict that the subject is justified and also towards the verdict that the subject is not really well justified. This suggests an ambiguity: the subject is *subjectively*

10 This is Cartesian tracking, not the Nozickian tracking which requires not that one would believe that *p* only if it were so that *p*, but rather that if it were not so that *p* then one would not believe it.

11 Not that these faculties could not on occasion really be operative while still misleading us. I simply have in mind our specific examples (such as the oatmeal case) where by hypothesis memory *is* ostensibly operative but is *not* really operative.

justified, but *objectively* unjustified. The subject is perhaps subjectively justified in that he quite reasonably attributes the belief to the operation of a normally reliable faculty of his own. This applies in all three cases: that of memory, that of inference, and that of intuition. But the subject is objectively unjustified in that the relevant faculty is not really operative in the case at hand. It is not really true memory, or true deductive inference, or true intuition that is at work.

If that much is right, it prompts two questions. First, is there an important externalist element even in the faculties of reflection: namely, those of memory, and reason, deductive and intuitive (along, perhaps, with introspection)? Is there some element that is in some way epistemically “external,” one on whose brute presence we must rely, one not automatically open to our reflective view, even when we look while cool, collected, alert, etc.? Second, if the “external” is thus involved in such justification, why stop with these as faculties that provide justification? Why not include also environment-involving perception, and even neighbor-involving testimony? The importance of any intuitive internal/external divide is thus put in question. This opens the way for an account of knowledge and justification that dispenses with any such deep distinction, while showing how its intuitive basis may be accommodated nonetheless, on a derivative, more superficial level.

A Virtue Epistemology

9.1 Aptness and Adroitness

A virtue epistemology V would involve the following components:

- Va X is an intellectual virtue *only if* X would produce a high ratio of true beliefs
- Vb B is a justified belief *only if* B is a belief acquired through the exercise of one or more intellectual virtues

Here is a way to interpret Vb (where ‘w’ ranges over possible worlds):

- J-APT (For all w) (B is apt-justified in w *only if* B is acquired in w through the exercise of one or more intellectual virtues that are virtuous in w)

Proposals like J-APT run against the new evil demon problem,¹ as follows. The victim of Descartes’s evil demon seems *not* deprived of *ordinary* justification, since his beliefs still derive from sources that we recognize as justification-conferring: namely, sensory experience, memory, etc. The environment changes radically, but the victim retains and uses a repertoire of intellectual virtues. True, because the environment is so radically abnormal and wrong for his normal virtues, those virtues may *not* qualify as virtuous relative to *that* environment. But, despite J-APT, the victim’s beliefs are still justified, in some relevant sense. So J-APT is not a full and illuminating enough account of

¹ A problem for externalism due to Stewart Cohen and Keith Lehrer: “Justification, truth, and coherence,” *Synthese*, 55 (1983), pp. 191–207.

all that might be involved in a belief's being "justified." Here now is a further component:

J-ADROIT (For all w) (B is adroit-justified in w *only if* B is acquired in w through the exercise of one or more intellectual virtues that are virtuous in our *actual* world)²

This is not open to the objections lodged above against J-APT. In particular, the beliefs of the victim can still be adroit even if they are not apt.³

Let now V-APT and V-ADROIT be principles that combine Va with J-APT and J-ADROIT, respectively, amounting to the following:

V-APT (For all w) (B is apt-justified in w *only if* B derives in w from the exercise of one or more intellectual virtues that *in that world* w virtuously would produce a high ratio of true beliefs)

V-ADROIT (For all w) (B, in w , is adroit-justified *only if* B derives in w from the exercise of one or more intellectual virtues that *in our actual world* virtuously would produce a high ratio of true beliefs)

Against V-APT and V-ADROIT (henceforth "the V accounts") it has been objected that the folk are not really inclined to relativize in the ways that these V accounts require.⁴ How damaging is this objection?

2 On the whole it seems to me preferable to take "actual" here as an indexical; with J-ADROIT we hence relativize through the contextual nature of our affirmation, where the context is set as the world of the subject who affirms (through speech or thought).

3 My use of the words "apt" and "adroit" here is of course largely stipulative, especially with regard to the implied contrast between the two, with a belief's "aptness" deriving from the epistemic quality of its sources *in the world wherein it is yielded by those sources*, whereas a belief's "adroitness" derives rather from the quality of its sources in the world wherein that belief is assessed by the assessor. Alvin Goldman's only objection against J-ADROIT, found in section II of his paper "Epistemic folkways and scientific epistemology" (in *Liaisons* (Cambridge, Mass.: MIT Press, 1992), pp. 155–63) is that "there is no evidence of relativization on the part of the folk" nor is epistemic "quality relativized to each possible world or environment." It is presumably for this reason that he would reject both J-APT and J-ADROIT. This objection is examined closely in my "Goldman's reliabilism and virtue epistemology," *Philosophical Topics*, 29 (1 and 2), pp. 383–400.

4 For example by Goldman in "Epistemic folkways."

The V accounts flow from “conceptual analysis” of epistemic justification; they derive from a priori reflection. Such conceptual analysis comes in two varieties worth distinguishing. First there is *meaning analysis*, which leads to conclusions that no one could possibly reject without failing to understand one or another of the constitutive concepts, and hence to conclusions that no one could possibly reject *period*.⁵ Secondly, there is *substantive analysis*, which leads to conclusions that are a priori and necessary all right, yet difficult enough that a mistake would not necessarily evince failure to understand.

Presumably the V accounts are offered as results of substantive analysis. We need to remember this when considering the claim that the folk are not inclined to relativize virtues and vices. Compare: “He is tall,” as applied by a child; by a basketball pro. “It is daytime,” as said in New York; as said in Tokyo. Or take “It is snowing.” Where do we check for falling snow? Often enough we relativize unawares, through contextual features not present to the consciousness of the speaker(s). Perhaps that is the sort of relativization proposed by the V accounts. Perhaps the folk do relativize often enough, in the context-driven way of indexicals, as in our examples.

Thus the V accounts explicate epistemic justification, at least partially, combining for an account secure against the new evil demon problem. Although it is objected against the V accounts that “there is no evidence that the folk are inclined to relativize virtues and vices,” this is outweighed by the fact that the relativizing may be contextual and implicit.

9.2 Externalism and Justification

We are discussing externalism: some main objections to it and various ways to modify or supplement it so as to meet those objections. It seems to many that epistemic justification is an internal matter isolated from the luck of your contingent emplacement. This is redolent of the Kantian view that moral quality derives from how it is in the subject’s mind and not, for example, from contingent external features yielding causal consequences of his policies and decisions, thus determining how well these turn out in the actual world.

5 That no one could possibly reject at least by denying it *in propria persona*. One could of course reject *indirectly* propositions one does not understand, for example as in “the proposition stated on a certain page of a certain tract.”

There is in epistemology an analogue of such Kantianism: epistemic worth should derive from qualities of the subject's mind and not from contingent external features determining (1) the causal origins of the subject's beliefs and inference patterns, and (2) how reliably truth-conducive *such* causal determination of beliefs and inference patterns is generally in our actual world. Suppose epistemic justification does indeed depend only on such internal matters independent of the subject's contingent emplacement. How, more positively and fully, should we then conceive of such justification? Some have tried to understand it as a matter of being blameless or of avoiding any wrongful violation of epistemic norms required for right belief formation. But this way of understanding epistemic justification takes us only so far, and comes up short: if you are brainwashed or brought up epistemically in a superstitious community, you may bear no blame for beliefs and inference patterns that still fall short epistemically. You are not then "wrongfully violating" any norms, since how you proceed is then *not* "up to" you, deriving as it does only from aspects of your psychology before which you are then helpless. Internal epistemic justification hence seems *not* fully explicable through mere blamelessness. How then *is* it to be understood? Before taking up this question, we face an objection.

Objection:

Suppose we view warrant as, approximately, whatever must be added to true belief to yield knowledge, and we insist on the hardline reliabilist analysis of the concept as follows:

WR S 's belief that p is warranted in world w iff S 's belief that p is produced by a belief-forming process that is reliable in w .

WR faces the new evil demon problem. But why not take a "heroic" line in response? "Yes, the inhabitants of the demon world are *unwarranted* in most of their beliefs." If you think this response is outrageous, consider the following. Unless one is going to maintain that all our intuitions form a nice, neat, coherent whole, it seems pretty clear that any epistemological theory will have to revise at least some of them. The heroic response would count as an instance of such revision. The question then becomes: Is it a revision that demands too much? (Or, perhaps: Are the intuitions that might lead one to favor reliabilism obviously trumped by the intuition that the demon-world victims are warranted in their beliefs?) Some pre-theoretic epistemic intuitions are of course obdurate, or at least much harder to abandon: thus, perhaps,

the intuition that it is possible for one to know that one sees a (real, flesh-and-blood) hand. How obdurate is the intuition that the victims of a demon world are warranted in their beliefs? Well it is, after all, a pretty far out possible world. Is it clear that any intuition about how matters stand epistemically in such a world will be strong enough to make one want to abandon reliabilism and join the opposition?

A reliabilist who gives the heroic response might also add that, while the demon-world victims are unjustified, they certainly aren't blameworthy for believing as they do. Does the intuition that the demon-world victims are warranted in their beliefs remain strong if this deontological conception is distinguished firmly from the what-must-be-added-to-true-belief-to-get-knowledge conception?⁶

Reply:

That objection does have force if our evil demon world is one in which our present virtues *would* lead us astray sufficiently often. Take for example a world where people, and even rational beings generally, would normally go astray in taking their experience at face value. In this especially hostile demon world, your beliefs plausibly *lack* warrant even if formed by taking experience at face value, or in some other way that is reliable in our actual world. It is this notion of warrant that is presumably invoked when we countenance possible higher beings who gain knowledge by properly and reliably forming ("warranted") beliefs, despite the fact that their epistemic ways, while successful in their world, would be miserably inadequate in ours.⁷ Correlative to this is the intuition that the victim of the demon is *not* justified, which fits the "heroic" response. It is this intuition that is catered to by our V-APT account, according to which the beliefs of such a victim would *not* qualify as apt-justified. (So the concept of apt-justification is our correlate for the sort of "warrant" that figures in the "heroic" approach by a hardline reliabilist.)

A further issue is now pressed by internalist epistemologists. Is there not some notion of epistemic justification, perhaps some *other* notion, that is *also* intuitively gripping (however it may relate to knowledge, and however it may relate to the kind of warrant that is necessarily

6 David Matheson gave me this objection, without necessarily endorsing it.

7 It seems to me that the folk would naturally understand such speculation, and would not hesitate to attribute knowledge and proper belief formation to the superbeings. So I doubt that one could plausibly cleave to Goldman's hardline anthropocentric epistemology, according to which folksily we recognize only homebound justification that rigidly binds virtue to our lists of what works in the actual world.

contained in knowledge), some other notion according to which the victim of the demon *would* still be justified. This is the notion that the classical foundationalist explicates in terms of a foundation where one takes the given and a superstructure that one builds on that foundation through valid reasoning. Sometimes internalist epistemologists even claim that the traditional issues of epistemology concern only such rational justification, and that our Gettier-related focus on “knowledge” and reliability is an unfortunate diversion.⁸ Traditionally the project has been to answer the skeptic, and this has been largely a matter of showing how, despite what the skeptic has to say, we still remain “justified” in certain of our important beliefs about the world around us, our friends and their thoughts, etc. It is *this* kind of justification that, according to internalists, you might still have even if unlucky circumstances make your animal mechanisms unreliable.

Given this, the project for the externalist is to show how, within a broad externalism, within truth-connected epistemology, we can understand such a notion of justification. This is a project engaged by our V accounts: by the inclusion specifically of V-ADROIT, and not only of V-APT.

9.3 Contextual Justification

Consider the crucial context-dependent expression, “actual” contained in V-ADROIT. It is this context-dependence that may help us do justice to internalist intuitions. Compare how the following is more than just accidentally true:

X is here iff X is where I am.

Yet it would be false to say, about a hypothetical situation, that we would determine whether an object would then be *here* by determining whether it would in that situation be (*there*) co-located with me.

On this account we can say that, if one *were* a demon’s victim, one’s intellectual performance might still be adroit (adroit-justified). That is to say, in a world *w* in which one was a victim, one’s beliefs acquired through our normal perceptual faculties would come out of processes that in our actual world *are* reliable, and hence those beliefs of ours in that world *would* be adroit (adroit-justified).

⁸ Thus Laurence Bonjour in his book on the a priori, *In Defense of Pure Reason* (Cambridge: Cambridge University Press, 1998).

If we attribute such contextual character to our concept of justification, however, it may then seem coordinately wrong to wonder what it is to be *really* justified. At a minimum, it will bear scrutiny exactly what that might mean. What does it mean for a time to be *really now*, or for a place to be *really here*, or for a person to be *really me*. Here as elsewhere such expressions may lack any clear sense until a contrast is specified: “really” as opposed to what? Perhaps we should drop the “really” whenever possible. Even once we do so, nevertheless, a sense of unease still lingers over such questions as “What is it to *be here* (what constitutes the hereness of being here)? What is it to *be now* (what constitutes the nowness of being now)? What is it to *be me* (what constitutes the me-ness of being myself)?” And similar unease would then extend to “What is it to *be justified*?” This question may well be appropriate only given contextual or other cues to clarify what is really meant. If what we are after is clarification of the *concept* of epistemic justification, accordingly, then the question remains in order even if the concept turns out to be indexical or otherwise contextual. But it is more problematic how to respond to a question that wants scientific or metaphysical elaboration of the *property* of being here, or that of being now, or that of being me. And we may then similarly wonder about the question of what it is to be justified, understood as a request for clarification of the nature of a supposed property. This is by no means to say that there is no metaphysically interesting content involved in the concept(s) of epistemic justification (content that for the virtue epistemologist will of course involve intellectual virtue), any more than the concepts of here, now, and myself are devoid of their metaphysically interesting content, involving respectively space, time, and the self.

9.4 The Demands of Internalism

It may be replied that we have not really done justice to internalist intuitions concerning epistemic justification, if internalists are worried about subjective justification, or subjectively adequate belief. Here is some reasoning to that effect.

Objection:

Internalists require a kind of internal rationality as necessary for knowledge, even if more is required. In any case, even if it is not required for knowledge, such internal rationality is an important epistemic status that a belief might or might not enjoy. And the V accounts

do not capture it. Specifically, V-ADROIT does not capture it, and V-APT fails even more clearly. Put another way, there is an intuition that there is something right about the demon victim's beliefs *from her own point of view*, so to speak, something subjectively adequate about them. But the idea that her beliefs are from faculties that would be virtues in *our world* does not capture this.

To continue the same point, consider the later claim above: "*Justified* too may be subject to various different sorts of contextual determination. That is to say, the indices relative to which uses of 'justified' get semantic assignments may be determined in various ways." But it seems that no assignment captures the idea that S's beliefs are subjectively adequate or appropriate, or subjectively justified in that sense. If the assignment makes a justification claim be about reliability in S's world, then the claim will be false. If the assignment makes a justification claim be about reliability in our world, it won't concern subjective appropriateness. Why should the subjective appropriateness of S's belief be a matter of whether her faculties would be reliable in our world?

In sum, the upshot of the defense of V-ADROIT is that it is supposed to explain our intuitions that the demon victim has justified beliefs. But it does not do that if our intuitions are to the effect that the victim's beliefs are in some sense subjectively adequate or appropriate, and that this is an important kind of positive epistemic status. Internalists have the intuition that there is something good about S's beliefs relative even to her world; that they are subjectively justified in some sense, and it is not just that her beliefs are formed by faculties that would be virtuous in our world.

The frequent talk by internalists about responsibility and duty testifies to their focus on some such notion of subjective appropriateness, even if they are wrong to cash it out in terms of doing one's duty, doxastic voluntarism, etc.

Reply:

There are three epistemic statuses (at least) worth distinguishing here. One corresponds to V-APT, one to V-ADROIT, and one to Foley-rationality, i.e., to the status that a belief has when it would survive deep reflection by the subject in the light of her deepest epistemic standards (roughly).⁹ Of course, the subject in the evil demon scenario may

9 Cf. Richard Foley, *The Theory of Epistemic Rationality* (Cambridge, Mass.: Harvard University Press, 1987), and *Working Without a Net* (New York and Oxford: Oxford University Press, 1993).

well be Foley-rational, and that is an epistemically good way to be, so far forth, but someone deeply superstitious and not fully rational might equally qualify as Foley-rational. Nor will it do to appeal simply to being blameless, since someone deeply epistemically flawed, *internally* flawed, may still be blameless. So the question is: what further internal state could we possibly be appealing to in supposing the victim to be internally, rationally justified? We might say: well, in arriving at his belief and sustaining it, he *reasons* well, and he *well takes his experience into account*, and he *remembers* well. The problem is that these are all matters that derive from states of the subject and from the causal network of such states, or from subjunctive or other modal relations of such states to the subject matter concerned, which will include the layout of the external world.

In reasoning thus, I assume that the notions of *adroitly taking one's experience into account* (and even that of reasoning *inductively* well) cannot be entirely independent of the modal relations between the contents of input states and output states, as one takes account of one's experience and as one reasons from beliefs to further beliefs. Thus, *visual experience as if there is something white and round before one* is a reason for believing that there is such a thing there before one, but only because in the actual world such a visual experience is reliably related to there being such a thing there: that is to say, what is required is that in the actual world such a visual experience *would* in normal conditions reveal the presence of such a thing before the perceiver. Most naturally one would want the state of the perceiver describable as "visual experience as if there is something white and round before him" to be necessarily such that it *would* normally reveal the presence of such a thing before the perceiver. Otherwise that same state would not be properly describable as such an experience.

Alternatively one might suppose that there is a state with an intrinsic mental character whose intrinsic mental character makes it properly characterizable as a state of experiencing thus (as a state of experiencing a white, round item, or the like), and this with logical independence of any modal relation that such a state, with such an intrinsic character, may bear to the presence or absence of white and round items. The problem for this view will be to explain how such states could possibly give a reason to believe that there is something white and round before one. It may be held that *it just does*. And now one will face the following prospect: the need for a boundless set of principles each with fundamental status, connecting various intrinsically characterized mental states with paired external facts of specific

sorts. Even though there is no modal relation between a given intrinsic mental state and its paired external fact type, that state might nevertheless serve as a reason to believe in its paired fact. And this will be so even when, as in the demon world, that sort of intrinsic mental state may in fact – when combined with the modally stable presence of the controlling demon – be a modal *counterindication* of its paired sort of external fact. Each such rational relation uniting the supposed mental/external pairs would be postulated as holding primitively, despite there being no modal relation among the mates, and even when the mental state modally counterindicates the paired external fact (assuming the presence of the demon to be modally quite stable relative to that world).

If we are to avoid such an unpleasant prospect, it is hard to see what the alternative could be other than an appeal to mental/external pairings by modal connection in the actual world, whether these pairings are already determined even by the very content of the experience or not. If the pairing is already so determined, then of course taking experience at face value would be assured of reliability simply by the fact that it would involve believing *that such and such* based on experiencing *as if such and such*, where neither the experience nor the belief could have the content (such and such) if they did not bear an appropriately reliable causal connection to external situations where it was the case that such and such.¹⁰

9.5 Skepticism and Circularity

A further question may be pressed on behalf of the skeptic: What if the actual world is *itself* a demon world? Are we then justified? Consider this conditional:

- (D) If the actual world is a demon world, then our beliefs acquired through our normal perceptual processes are justified.

Is that a true conditional? Well, either it is an indicative, material conditional or it is a disguised subjunctive (or in any case modally strong) conditional. If it is an indicative, material conditional, then it is true

¹⁰ And something similar would seem in order with regard to inductive reasoning, though this would be an even more complex matter, with a vast literature.

vacuously, since its antecedent is false. And if it is rather a subjunctive conditional, then we are back to our earlier reasoning. Conditional D thus seems true regardless of interpretation. (Of course, *also* vacuously true is the opposing material conditional that shares D's antecedent but denies its consequent.)

It might be replied that this is question-begging against the skeptic, since it assumes that the actual world is *not* a demon world. But the most formidable skeptic does not want to take it as a premise that the actual world *is* a demon world. Rather, he takes it as a premise that the actual world *might* be a demon world. And his premise is *not* that *for all we know* the actual world is a demon world. No, this is rather his conclusion, or close to that. His premise is only that *metaphysically* (or at least logically) the actual world *might* be a demon world. And we can of course agree with *that* premise, while rejecting the skeptical conclusion nonetheless.

There is a less formidable skeptic who is answered more easily, if he needs to be answered at all. This is the skeptic who says that we might metaphysically be in a demon world, and that we cannot just assume or presuppose that we are not, since it is part of his skepticism not just to argue that we do not know what we take ourselves to know, but *also* actually to put in question all of our supposed knowledge of the contingent, external world around us. Mark well, this means not just to raise the question as to whether we are wrong in the likes of the following single hypothesis: *that most of our beliefs are true*. No, also put in question, along with that hypothesis, is each of our contingent beliefs, particular or general, about the external world. So it would not do for the dogmatist to just *take for granted* certain contingent conditionals. The skeptic is precisely *not* granting any of that. On the contrary, he is putting all such conditionals in question, along with every other contingent truth about the external world. So it would beg the question against him to take for granted an answer to that which is in question.

True enough, if we allow the skeptic to put all of that in question, in one fell swoop, then there is no defeating him without begging the question. Well, perhaps there is, actually, since we can press against the skeptic a charge of arbitrariness.¹¹ Why is he putting in question all

11 A famous Reidian charge, to which Descartes also seems quite alive (ironically, since Reid's charge is made as an objection to Descartes), as suggested by the true depth of the skepticism that he takes up, which is in fact Pyrrhonian and not just Cartesian.

propositions about the external world, and not those about the realm of the necessary or about his own states of mind? After all, it is not as though we are infallible about *everything* in these realms. It is not as though here we could never make a mistake. And even if we take ourselves to be infallible and incorrigible about certain of our beliefs at least in these realms, how would we know which of them are truly so protected against error? Would we not need to employ, in order to demonstrate our prowess, the very faculties that a skeptic does or might put in question? How then could we defeat such a skeptic without begging the question?

If we are at least right about these restricted areas in which we are assured of success if careful enough, it might be replied, we can then see how we enjoy, with regard to these restricted domains, something truly special. We are there exempt from dependence on the luck of our emplacement, on the luck of a clear channel, of a favorable medium. With regard to these special domains, there is no such dependence on channel or medium. So our control and our freedom from the favor of luck is at a maximum. No wonder the knowledge that we enjoy in such cases has been so exalted in our tradition, and has even been raised to the level of the only *true* knowledge. No wonder so many have been tempted to be “skeptical” about any other knowledge. All other knowledge would then reasonably fall below the clearly enough understood exalted level, since all other knowledge would be dependent on matters beyond our true control, and subject to the luck of channel or medium, or other external circumstances.

Even if we go along quite far with such reflections, however, should we not equally wonder about the luck of proper teaching, of proper upbringing, of proper constitution? And if these matters are to a substantial extent also beyond our ability to control, if how smart we are, how persistent, how patiently and deeply reflective, if these are all, to an impressive degree, matters of constitutional luck, as with raw IQ, then how well we do even in such matters will depend, for a clear example, on whether one is Ramanujan or a mental defective. And *that* would seem largely a matter of luck, something outside one’s control. So why exactly is the luck of our emplacement so detrimental to true knowledge, whereas the luck of our constitution is of no concern? What if it might more easily turn out that our constitution be altered in a knowledge-affecting way, than that our emplacement be altered with such a consequence? What if our relevant external channels and media are quite stable modally, whereas our relevant respects of constitution

are somewhat stable modally, but much less so? Are beliefs dependent on the stable channels still to be assessed as lesser knowledge by comparison with beliefs dependent on our less stable relevant constitution? Maybe so, but if so, we need a reason to believe it.¹²

9.6 An Objection to Contextual Justification

Finally, there is one further problem for our V accounts. Consider what thinkers in world *w* would be thinking were they to think that a belief *B* in that world was *justified*. They would have to be thinking that *B* derived from a process that is in fact virtuous enough in our actual world. But that seems absurd.

Compare, however, what thinkers elsewhere have to be thinking if they think that a certain table is *here*. What they have to think is that the table is where we are now. This suggests that epistemic terminology such as “epistemic justification” is not quite like *indexical* terminology such as “here.” But, even if this conclusion is hard to deny, our epistemic terminology may still be highly contextual. Compare “nearby.” It might be thought that

- (a) “*x* is nearby” is equivalent to “*x* is near here” (where for simplicity we interpret co-location as a limiting case of nearness).

But this is questionable if we compare the following two claims:

- (b) If I were far away from here, this table would still be here.
- (c) If I were far away from here, this table would still be nearby.

Claim (c) seems false, claim b true. But if we use the equivalence (a), then claim (c) is equivalent to (d):

- (d) If I were far away from here, this table would still be near here.

¹² Let it not be said, anyhow, that we are essentially constituted just as we are, so that possible variation in constitution is not to be countenanced. Surely we could have been quite a bit different both in circumstantial *and* in constitutional respects. We could have had somewhat less or somewhat more intelligence, patience, persistence, and so on, and yet have remained numerically the same, though of course qualitatively different.

But this follows from (b), and hence cannot be false if (b) is true. So the supposed equivalence (a) is false after all.¹³

“Justified” too may be subject to varieties of contextual determination. That is to say, the indices relative to which uses of “justified” get semantic assignments may be determined in various ways. Compare for another example, the term “tall.” When someone is said to be “tall” there is usually some contextually determined reference class. But when we attribute a belief that someone is tall, this very attribution may so change the context that a new reference class is selected. Thus an anthropologist may tell us that pygmies who knew they were tall would never call attention to their own height. Clearly the anthropologist would not be using our normal standards of tallness in making that attribution. Especially is it not the case that the beliefs of the pygmies are to be assessed for truth using us as the reference class. Something similar may happen as we move from us in the actual world, assessing our beliefs as justified or not, to the victim of the demon, who is supposed to make similar assessments. When the victim assesses a certain belief as justified, his world, not ours, is the relevant contextual index. Yet, despite this, we can, from our contexts, attribute such beliefs to pygmies and demon’s victims.

9.7 Conclusion

A fuller virtue epistemology would include not only the *V* accounts, *V_a*, and *V_b*, along with principles *V-APT* and *V-ADROIT*, but also the following on knowledge and its relations to belief, truth, justification, and faculties.

If a faculty operates to give one a belief, and thereby a piece of direct knowledge, one must have some awareness of one’s belief and its source, and of the virtue of that source both in general and in the specific instance. Hence it must be that in the circumstances one would (most likely) believe *P* only if *P* were the case; i.e., one’s belief must be *safe*; or, more strictly, one’s belief must be based on an indication, a safe deliverance of a virtuous source. And, finally, one must grasp that one’s belief non-accidentally reflects the truth of *P* through accepting an indication of *P*, thus manifesting a cognitive virtue. The fuller

13 Or so it seems at least on *one* way of interpreting “nearby.” We could perhaps read “nearby” so that (a) is true, and then (c) would seem true rather than false. To my ears that seems a possible reading.

account therefore would combine requirements of *safety*, of *reliable virtues*, and of *epistemic perspective*.¹⁴ Although this essay is not a complete defense of that fuller account, I hope, and expect, to have said enough to help provoke interesting debate.

14 More needs to be said, however, about the sort of doxastic ascent apparently required. It would be absurd to require at *every* level that one ascend to the next higher level in search of justification. Yet it seems no less absurd that a meta-level belief B' might help justify an object level belief B even though B' itself falls far short of justification. What the proper epistemic status and worth of a belief in fact require is rather that it be part of a body of beliefs with adequately comprehensive coherence. Such a comprehensively coherent body of beliefs would need to include meta-beliefs about object-level beliefs, about the faculties giving rise to them, and about the reliability of these faculties. Nevertheless, we would need also to allow that, at some level of ascent, it will suffice for the epistemic status and worth of a belief that it be non-accidentally true because of its virtuous source, and through its place in an interlocking, comprehensively coherent system of beliefs, *without* needing to be in turn the object of higher-yet beliefs directed upon it. That sketches my preferred alternative. Obviously, we have considered only some necessary conditions of knowledge or justification. As we saw earlier, it would be a mistake to suppose that such comprehensive coherence by itself would suffice to yield justification for its connected beliefs, since, for one thing, a requirement of mesh with experience must also be met. So it is clear that the sketch is both incomplete and incompletely defended. I plan to complete the job elsewhere.

III

Replies

Reply to Sosa

Laurence Bonjour

I will focus in this reply mainly on the last three of Sosa's four chapters, beginning with the last two (in sections 10.1 and 10.2), which contain Sosa's objections to internalism and his preferred semi-externalist alternative, and then returning (in section 10.3) to Sosa's critique of internalist foundationalism in his chapter 7. Despite my doubts about the concept of knowledge (see section 1.5 in my main essay), I have little quarrel with most of what Sosa says in his chapter 6, though a few points of disagreement will emerge in the discussion of these other issues. I will conclude my reply (in section 10.4) with some brief further reflections on the issue of skepticism as it relates to both of our positions.

10.1 Internalism and Externalism (yet again)

Sosa begins his consideration of internalism by focusing on two versions of the view, one that he attributes to Descartes and the other of which is due to Chisholm. The allegedly Cartesian version (I have some doubts about the correctness of this attribution) holds that the justification of a belief depends only on the "appropriateness" of the thought that generates and sustains it, where the considerations that determine such appropriateness pertain only to matters internal to the mind of the subject. The Chisholmian version, in contrast, holds that justification is internal in the sense that one can always find out merely by reflection whether a particular belief is justified. Sosa's suggestion is that the connection between these two views depends on two theses, neither of them in his view obviously correct: first, the thesis (his AI) that things that are internal to the mind are always "open to discovery by reflection"; and, second, the thesis (his AS) that the fact of

“supervenience” between such properties pertaining to the mind, or indeed any non-epistemic properties that justification might depend on, and justification itself is also thus discoverable by reflection. Given these two questionable theses, he suggests, the Chisholmian version of internalism follows from the (allegedly) Cartesian one.

But why, asks Sosa, should either of these theses be accepted? It is, I take it, in the search for an answer to this question that he moves to a consideration of deontological conceptions of justification, about which I will have more to say momentarily. But I first want to suggest that none of the versions of internalism mentioned so far, neither the Cartesian version nor the Chisholmian version nor the deontological conception of justification, really succeeds in getting at the heart of the internalist view or at its essential motivation. These emerge when we ask instead the questions raised at the beginning of my main essay: Do we have any good reasons for thinking that our beliefs about the world are true (or at least approximately true)? And if so, what specific form or forms do those reasons take? I suggest there that the issue posed by these questions is the most central one in epistemology. But while I still think that something like that is true, it will be enough for present purposes to insist that it is, when judged on either a substantive or a historical basis, at least one very central epistemological issue. In particular, I think that this is the main issue that animated Descartes and motivated his resulting internalist view. And I find it striking that Sosa has almost nothing very explicit to say about this issue in the present essay, that at most it barely emerges in a couple of brief passages (most clearly at the end of his section 8.1 and again in section 9.4).

The first thing to notice about this issue is that it is essentially a *first-person* issue, one that is in fact better captured in a first-person singular formulation than in the first-person plural formulation just employed. The basic question (which each person must in the end ask for himself or herself) is whether *I* have good reasons for thinking that *my* beliefs are true (and, if so, what form those reasons take). And the reason that this leads to an internalist view is that the reasons in question are supposed to be reasons that *I have*, not in the impossible sense of having them explicitly in mind at every moment, but in the sense of their being more or less immediately available or accessible.

It is common among internalists to use the word “justification” to refer to the possession of reasons of this sort for a belief (and to claim further that justification in this sense is one of the requirements for

knowledge); and it is this practice that leads to what is perhaps the most standard understanding of the internalism/externalism issue, according to which the basic question is whether it is internalism or externalism that gives a correct account of that species of justification that is relevant to epistemic issues (and is perhaps one of the ingredients of knowledge). But while I have adopted this way of putting things in my main essay, it is important to see that the fundamental importance of the issue of whether I have good reasons for my beliefs does not depend on its supposed relevance to an account of “epistemic justification” (nor, still less, on its supposed relevance to the concept of knowledge). For the denial that such reasons exist amounts already to a very severe and intuitively implausible version of skepticism, no matter what may be said further about the concepts of justification (and of knowledge). If I have no reasons for thinking that my beliefs are true, then I am “flying blind” in an obvious way, and the seriousness of this result is not diminished in any clear way by adding that in some other, perhaps externalist sense my beliefs are (or rather *may be* – see below) still justified.

This main internalist account of justification (for it will be convenient to continue to refer to it in this way) provides an obvious rationale for the Chisholmian version of internalism. If the question is whether I have good reasons for my beliefs, then the answer must appeal to reasons that I genuinely have and which are thereby available or accessible to my reflection. There is still an issue here about whether I might have full access to a reason but still somehow fail to recognize it as such or whether I might mistake something to which I have such access for a good reason when it is not (see further below), but a supposed form of justification that supervenes in the way that Sosa suggests on some complex set of facts in a way that I am unable to comprehend or that depends on complicated facts about my mental states that entirely elude my grasp will clearly not constitute a reason *that I have* for thinking that the belief in question is true. And, even more obviously, varieties of justification that depend even in part on entirely external facts will not yield such reasons. Thus the Chisholmian view plausibly reflects the particular sort of justification that the internalist is primarily concerned with.

On the other hand, the connection between the internalist understanding of justification and deontological notions like duty or responsibility is much more problematic. While I am, alas, one of those responsible for the idea that being epistemically responsible or satisfying one’s epistemic duties is tantamount to being justified in the

internalist sense,¹ it is in fact relatively easy to see that this is wrong, indeed that being epistemically responsible or satisfying one's epistemic duties (hereafter I will use only the latter of these formulations) is not even *sufficient* for internalist justification. As I have discussed more fully elsewhere,² one main reason for this is the possibility of situations of what I will refer to as epistemic poverty. Suppose that a group of people is in a situation where the kinds of evidence or the methods of inquiry available to them are so limited as to make it difficult or even impossible to come up with strong evidence or good epistemic reasons for answers to many important questions. In fact, it is clear that many human beings living in earlier eras have found themselves in such situations, and perhaps less clear than is often thought that we are not still in one to a considerable extent. It is quite implausible to insist that in a situation of this kind, epistemic duty still requires accepting only beliefs for which there are strong reasons, since this would mean that the people in question are required to have no beliefs at all about many important issues. To insist on such a view is in effect to give the avoidance of error an absolute and unwarranted priority over the discovery of truth.³ But while it is plausible to think that it is no violation of epistemic duty to accept beliefs on a weaker basis in such a situation, it remains the case that the people in question fail to possess good reasons for thinking that beliefs accepted in this way are true and hence that their beliefs are not justified to any substantial degree in the main internalist sense, even though the requirements of epistemic duty have been satisfied. (I think that virtually all internalists would accept this result and that their occasional suggestions to the contrary are a result of focusing on situations of at least approximate epistemic plenty.)

Cases of epistemic poverty are cases in which it seems possible to fulfill one's epistemic duty without being epistemically justified in the main internalist sense. Are there also cases of the opposite sort, cases in which a person has good, truth-conducive reasons for his or her

1 Mainly in *SEK*. The most explicit passage is on p. 8 (though the better formulation of internalism in terms of having good reasons to think that one's beliefs are true is also present there).

2 In my paper "The indispensability of internalism," *Philosophical Topics*, forthcoming.

3 I take the foregoing to capture one part of Alvin Plantinga's criticism of the deontological view of internalism in his *Warrant: The Current Debate* (New York: Oxford University Press, 1993), chs 1 and 2; he errs, however, in thinking that internalism in general is committed to such a view.

beliefs, but still fails to satisfy the requirements of epistemic duty, thereby showing that the satisfaction of duty is also not *necessary* for internalist justification? I think that there are indeed such cases, and indeed that Sosa comes pretty close to one of them in section 8.3, where he mentions the possibility of a person who is guilty of “uncaring negligence” in the way he or she arrives at a particular belief. It is plausible to suppose that such a person is violating an epistemic duty, but it also seems entirely possible that he or she might still happen to have good reasons for the beliefs that result and so be justified in the internalist sense. (This would mean that what I am calling the main internalist account of justification fails to capture all aspects of epistemic excellence, even of internal epistemic excellence. I am unable, however, to see why this would constitute an objection to it. On the contrary, it seems to me that nothing is likely to be gained by insisting that all aspects of epistemic excellence must be lumped together into one conception.)

For these reasons, it seems to me clear that epistemic justification cannot be simply identified with the fulfillment of epistemic duty as the deontological conception claims. What is true, I would suggest, is rather something substantially weaker: seeking good epistemic reasons and believing on the basis of them is, at least in situations of relative epistemic plenty, *one* important requirement of epistemic duty and arguably the most central of all. This makes it easy to understand how some internalists, myself included, were led to overstate the connection between the two concepts. But it remains the case that the idea of satisfying epistemic duty turns out to be quite distinct from the main internalist conception of epistemic justification. And much the same thing is true of the idea of “wholly appropriate thought,” which figures in the version of internalism that Sosa ascribes to Descartes: thinking in a fully appropriate way is neither clearly sufficient nor clearly necessary for having good reasons for the beliefs that result, so that this is also not a good way to formulate the central idea of internalism.

I turn to a consideration of some of the other difficulties and problematic examples that Sosa raises in relation to internalism. Given this altered understanding of the internalist view, many of them can be easily answered or set aside, while others will require more extended discussion.

Consider, first, the idea of *ontological internalism*: the thesis that justification depends only on properties that are *metaphysically* internal to the person in question (section 8.3). In fact, neither the specific version of internalism just suggested nor indeed, as far as I can see, any

reasonable version is committed to such a thesis. As Sosa himself half recognizes (footnote 3), reasons for beliefs will very often involve an appeal to a priori discernible inferential relations among propositions, and these are obviously not metaphysical features of the individual person in question. Thus ontological internalism would lead immediately to skepticism about any claim whose justification depends on any sort of reasoning.

Second, once the deontological view is set aside, the issue of doxastic voluntarism has no special relevance for internalism. One may have an excellent reason for an irresistible belief, and indeed the recognition of such a reason may well explain why the belief is irresistible.

Third, consider a person whose beliefs are arrived at in ordinary ways and the counterpart evil-demon victim whose mental life is entirely parallel (section 8.4). Are the corresponding beliefs of these two individuals equally justified? On what I am describing as the main internalist view, the answer is obvious: any reasons for thinking that such beliefs are true that are available to one of these individuals will be equally available to the other, and so they will be justified to the same degree. But is this the correct answer to the question? It is, if it is *their reasons* for their beliefs that we are concerned with, and this, I am insisting, is at least one extremely important epistemological issue. But obviously there are also important differences between the epistemic situations of these two individuals, differences that might possibly be captured by saying that their beliefs are not equally justified under some other conception of justification. I am not at all sure that this is the best way to capture the relevant differences. But the main point I want to make is that there is no very compelling reason to think that there is only one relevant kind of justification of which the internalist view and such possible alternatives must be viewed as giving competing accounts, so that only one (at most) of them could be correct. On the contrary, I suggest, the situation of the two individuals is similar in some respects and quite different in others, and either the samenesses or the differences can be formulated in terms of intelligible conceptions of "justification" (though it is doubtful that either really need to be so formulated).

Fourth, what about the case of the pupil whose normally reliable teacher gets confused on a particular day about the value of $(x^n)^n$, believing that the exponents should be added rather than multiplied (section 8.4)? The pupil comes to believe on this basis that $(2^2)^2 = 2^4$, thus arriving at a correct result on the basis of an incorrect general principle. Is this belief justified? If we assume (as we must to make the case

clear) that the pupil has no other knowledge or insight of any sort that is relevant, then it seems clear that he or she has a fairly strong reason for thinking that the belief is true: it is an immediate consequence of a more general belief derived from a source that is known by the pupil (I again assume) to be highly reliable in this area. This reason is just as strong as would be the comparable reason that a pupil (in the same situation in the indicated respects) would have had if the teacher had stated the correct general principle. Thus the pupil's belief is, from the internalist standpoint, fairly strongly justified. Of course, there are again also other things that can be said about it. Most obviously, its truth is a matter of luck or accident, so that it thus fails to satisfy one of the fairly standard anti-Gettier conditions of knowledge. Perhaps, though I am again not sure that this is the best way to put the point, there is also a further "social-aetiology" conception of justification under which the belief is not justified. And there are no doubt other possible conceptions of justification that might be applied as well (such as a simpler reliabilist conception, on which it is at least not obvious – depending on how the relevant belief-forming process is specified – that the belief is not justified). But none of these other conceptions seems to me to provide any real reason to question either the significance of the specific question that the internalist conception focuses on or the correctness of the answer that it gives to it.

Fifth, consider the case of Mary and Jane (section 8.4), who arrive at the same conclusion, "Mary through a brilliant proof, Jane through a tissue of fallacies," even though Jane is usually much better at the sort of reasoning in question. This seems to me to be the most difficult and interesting of the cases that Sosa describes, and I want to deal with it carefully and in detail.

Consider first the issue, which Sosa bypasses, of what to say about their two beliefs at the moment that they arrive at them, with the relevant reasoning (I will assume) still fresh in mind. Each of them presumably thinks that she has a good reason for the conclusion at which she has arrived, but are they correct about this? It is easier to answer this question with respect to Mary. Mary's argument is completely cogent and she, presumably, is capable of recognizing that this is so, with her lesser degree of competence meaning only that it was less likely that she would have hit upon such an argument.⁴ Thus Mary has

4 The case would be quite different if Mary's competence is so low that she cannot genuinely grasp the correctness of the argument, but only happens to go through a series of steps whose cogency could be appreciated by someone else.

a very strong reason for her belief and is, in the internalist sense, clearly justified to a high degree. Jane, on the other hand, has arrived at her conclusion via an entirely fallacious argument, where this is presumably something that she is capable of recognizing, even though she has for some reason failed to do so. Thus she fails to have a good reason for thinking her conclusion is true, despite thinking that she does, because her reason would not withstand careful scrutiny. (Perhaps having an argument that seems superficially to be cogent constitutes some reason for the conclusion, but not a very strong one.⁵)

What then about the later stage upon which Sosa focuses, the stage at which each of the two believes the conclusion in question, believes herself to have arrived at it via a valid proof, but has “forgotten much of her reasoning” – enough, I will assume, that neither the fallacious character of Jane’s reasoning nor the cogent character of Mary’s is any longer capable of being reflectively recognized by them. In this situation, I submit, neither of them has nearly as strong a reason for thinking that the conclusion is true as Mary originally had and thus neither of them is justified to that same very high degree from an internalistic standpoint. Each of them has *some* reason: they each remember having gone carefully through a seemingly cogent argument and each of them also presumably has good reasons for thinking (1) that she rarely if ever has “memory hallucinations” about having gone through such episodes, and (2) that she rarely is mistaken about the cogency of a carefully considered argument. How strong these reasons are and whether one of the resulting justifications is stronger than the other will depend on the details of (1) and (2) (how rarely?) and on the strength of the reasons supporting them.

But is there not still, as Sosa suggests, another obvious variety of justification (“personal aetiology justification”) in which Mary is justified and Jane is not because Mary’s belief originally resulted from cogent reasoning and Jane’s only from “a tissue of fallacies”? Perhaps so (though it is again worth noting that the most standard view of this general sort, namely reliabilism, does not yield in any clear way the result that Jane is not justified; this will again depend on how the belief-forming process is specified). But even if such an alternative view of justification can be specified (which I have no reason to doubt),

5 Suppose that the fallacy is so subtle that Jane is incapable of recognizing it, despite her very high degree of logical competence. Then, I would suggest, she does have a good reason and is (internalistically) justified, even though there is still something wrong with her epistemic situation, something that might again be captured by some other conception of justification.

this has no tendency at all, as far as I can see, to show either: (1) that the internalistic response to the internalistic question is wrong; or (2) that the internalistic question is philosophically unimportant or uninteresting; or, most importantly, (3) that there is any useful way to meld these different conceptions of justification and the results they yield into one neat package, or any real point in attempting to do so. One other point worth noting, whose significance I will return to later, is that neither Mary nor Jane will be able to tell on the basis of what they can remember whether or not their beliefs are justified in this personal aetiology sense. Nor indeed will anyone else who does not have reliable access to the relevant history.

Sixth, consider finally the somewhat similar case that Sosa discusses at the end of chapter 8: a case of rational paradox, in which rational intuition apparently yields both some set of beliefs B_1 through B_n , but also yields the further belief B_{n+1} to the effect that the other beliefs are logically incompatible, so that at least one of the results of rational intuition (including B_{n+1}) must be mistaken. It is reasonably clear what the main internalist conception of justification that I have been advocating would say about such a case. Assuming that there are no relevant differences among the rational intuitions in question (in terms of such things as degree of strength or clarity), there is for each of the relevant beliefs a substantial and equal reason to think that it is true, but one that is weakened by the further insight that at least one of them must be false.⁶ Clearly this is an intellectually unsatisfactory situation, and perhaps it should be added that the justifications in question are, as it were, under a special cloud because of the person's complete certainty that one of the beliefs in question is false – a situation that very rarely obtains (though weaker analogues of it are quite common).⁷

Sosa's suggestion about this case is that there is also an intuitive inclination to say that the false belief B_1 , *whichever* one it is, is *not* justified, and he further suggests a distinction between a "subjective" and an "objective" sense of justification: each of the relevant beliefs is subjectively justified, but only the true ones are (in the absence of any further problem) objectively justified. I have no objection to describing the internalist variety of justification as "subjective," in the sense that it pertains to the person's own subjectively accessible reasons, though

6 I assume that this further very simple insight, not the one that leads to B_{n+1} but rather the one that says that all of B_1 through B_{n+1} cannot be true, is not in question – which presumably means that it is much clearer or stronger than any of the others.

7 There is a discussion of cases of this sort in *IDPR*, section 5.4.

I would insist that the presence or absence of such reasons is just as much an objective fact as anything else, something that it is possible, at least in the short run, to be mistaken about.⁸ Nor do I have any objection to characterizing the obvious difference between that belief and the others in terms of a different, “objective” variety of justification, though I am again not sure that this is the most perspicuous way to put the point. In fact, Sosa’s further characterization in terms of the person’s being or not being “non-accidentally in touch with the truth” seems to me clearer and more helpful, and I would question whether there is any real point to insisting that this constitutes a further species of justification. Moreover, it is again worth noting that the person who is in such a situation will be unable to tell (while it persists) which of his beliefs, if any, are justified in this objective sense – and that it is also quite possible that no one else may be able to do, either at the time in question or perhaps ever.

The discussion so far may be summed up in the following three points: First, there is a perfectly intelligible internalist conception of justification, namely that of a person having good reasons to think that his or her belief is true. Nothing in Sosa’s discussion so far – but see further below – shows either that there is anything objectionable about this conception in itself or that the issue to which it speaks is not of great philosophical interest. Second, there are other sorts of differences pertaining to the epistemic situations in which people may find themselves that (1) are also of philosophical interest to varying degrees, (2) are not captured by this internalist conception of justification, and (3) may perhaps be captured by other notions or conceptions of justification. (Though there is no strong reason why all of the conceptions in question, including the internalist one, could not be formulated in an entirely adequate way without speaking of “justification” at all.) Third, there is no very compelling reason to think either (1) that these different notions of “justification” must somehow add up in any interesting way to one overall or resultant notion of “epistemic justification,” or (2) that one of these conceptions must be correct and the others mistaken as to the actual nature of “epistemic justification” itself. They are

8 In fact, Sosa seems to have something rather different in mind by “subjective justification,” since he takes it to depend in this case on the person’s having a justified belief that rational intuition is itself a reliable cognitive faculty. My own view is that a priori justification does not depend in general on such a higher order belief – and could not, on pain of skepticism. See *IDPR*, section 5.5.

simply different conceptions that speak to different issues and do not seem in any obvious way incompatible with each other. (As Sosa suggests at one point (chapter 8, n. 1), “justification” is more a technical term of philosophers than a term of ordinary language, so that it is not at all surprising that it can be used in different ways.⁹)

This is, of course, essentially the same reconciliatory conclusion that I reached at the end of chapter 2 of my main essay: There are many different and interesting epistemological projects and resulting conceptions of justification between which it is neither necessary nor desirable to choose in any general or absolute way. Thus I have no objection to partially or entirely externalist conceptions of justification (though I still doubt, for essentially the sorts of reasons indicated in that chapter, that “justification” is the best term to convey them). At the same time, however, I want to insist that the traditional internalist project that was the focus of my own main essay is at least equally legitimate, and that the tendency of many externalists to insist that their project is the only viable one should accordingly be resisted. Let us all practice epistemological tolerance!

But while I believe that all these epistemological projects are legitimate and that at least many of them are interesting, I want to close this rather long section by reiterating very briefly and in slightly different terms the point made at the end of chapter 2 of my main essay: The internalist conception of justification is indispensable and in a certain way central if any of the other conceptions are to be applied to actual cases in a more than hypothetical way. An externalist can say that *if* the specified external conditions are satisfied, then a particular belief is justified in his sense, but the question of whether either the believer or anyone else has any good reason to think that the antecedent of this conditional is in fact satisfied can in the end only be dealt with in an internalist way.

9 The most obvious reason for thinking that there *must* be one “right” conception of epistemic justification is of course the idea that there is one univocal conception of knowledge, which must then presumably invoke one of these conceptions in its warrant or justification condition. I have already expressed (in section 1.5 of my main essay) my doubts about the existence of such a univocal conception and, still more, about how much real significance it would have if it did exist. I will add, however, that even if there should after all turn out to be a univocal conception of knowledge that includes one specific conception of “epistemic justification,” this would not diminish to any serious degree whatever philosophical interest other conceptions of justification might seem on their own merits to possess.

10.2 Sosa's Virtue Epistemology (and yet more about Internalism and Externalism)

Sosa's own epistemological position is a rather complicated version of virtue epistemology, though one that is quite different from some other recent views to which that label has been applied in that it places little emphasis on what are most ordinarily described as intellectual virtues, i.e., such things as intellectual courage, open-mindedness, and the like.¹⁰ His description of the view (in his chapter 9) is also somewhat incomplete, with an important dimension of the position being indicated only at the very end, mainly in his final footnote. (I will have just a little to say about this dimension at the very end of this section.)

On Sosa's conception, an *intellectual virtue* is a source of beliefs, presumably a faculty or capacity, that produces or would produce a high ratio of true beliefs. (He mentions sensory experience and memory as possible examples.) His basic conception of justification is then that a belief is justified only if it is "acquired through the exercise of one or more intellectual virtues."¹¹ (This is obviously so far very similar to the familiar reliabilist conception, according to which a belief is justified if the cognitive process that it results from is reliable in the sense of producing a high ratio of true beliefs.)

It is obvious at once that the status of a particular faculty or capacity as a virtue and so the justificatory status of the beliefs that it produces according to this conception will depend on the surrounding context in which the faculty operates. Sense perception, for example, will presumably be a virtue in a world that has roughly the character that we commonsensically attribute to the actual world, but not in a world where sense experience is generated by, e.g., a Cartesian demon. This would seem to suggest that the perceptual beliefs of those who inhabit such a demon world are not justified, and according to Sosa there is a sense of justification, "apt-justification," in which this is so.

He acknowledges, however, the intuition that the demon victims' beliefs are in some sense justified, thus creating an apparent objection (the "new evil demon problem") to externalist views like reliabilism

10 For a view that stresses virtues of this sort, see Linda Zagzebski, *Virtues of the Mind* (Cambridge: Cambridge University Press, 1996).

11 There is a potential problem lurking here about what to say about faculties or capacities that yield beliefs only in combination with others (memory and reasoning are two obvious examples), making it difficult to see how to assess the truth-conduciveness of such a faculty in isolation. Sosa says nothing about this, and I will not pursue the issue further here.

and to a virtue epistemology that recognizes only apt-justification. Sosa's response to this potential objection is to distinguish a second sense of justification, "adroit-justification," according to which a belief is justified if the faculty or capacity from which it results is a virtue "in our *actual* world," even if not in the world of the believer, where "actual" is to be understood as referring indexically to the world of the person making the assessment. His suggestion is that the intuition that the demon victims are in some sense justified reflects the fact that they are (from our perspective) adroit-justified (or at least seem to be, since we could of course be wrong about the actual character of our own world).

I have several comments to make about this view. First, in accordance with the policy of epistemological tolerance endorsed above, I have no quarrel with Sosa's idea of an intellectual virtue as such or with the specific conceptions of justification that he bases on it. The concepts in question seem to be clearly defined, intellectually interesting, and quite possibly valuable for various purposes. I reject, however, for reasons already indicated, the idea that they capture the *only* correct conceptions of justification.

Second, however, along with Goldman, I have serious doubts whether the distinction between apt-justification and adroit-justification is to be found in common sense or ordinary thought. Goldman's specific objection is that ordinary people do not relativize justification in this way, and Sosa attempts to reply by arguing that relativization is often implicit and unrecognized. But while this may be true in the sorts of cases that Sosa cites as analogies, it seems to me quite implausible that it extends to anything as commonsensically outlandish as relativization to possible worlds.

Third, I am even more doubtful that the notion of adroit-justification adequately captures the intuition that the demon victims' beliefs are justified. Part of my reason for saying this derives from the previous point: The intuition in question does not seem to me to depend on even an implicit familiarity with the very sophisticated notion of adroit-justification. But I also think that there is a much more straightforward account to be given of the intuition in question, namely that the demon victims seem to us to be justified because they seem to possess *good reason* for their beliefs – indeed precisely the same reasons that we ourselves possess for our beliefs about our world. To be sure, the result of the demon's machinations is that their reasons lead them astray, but this has no tendency to show that those reasons are not good ones. Giving a philosophically adequate account of the

reasons in question and showing that they are indeed good ones turns out, notoriously, to be very difficult (and it is to this project, of course, that the latter stages of my main essay are addressed). But that these reasons do *seem* to be good seems to me beyond question at an intuitive level, and that the demon victims have the very same reasons is stipulated in the example, thus accounting much more straightforwardly for the intuition that their beliefs are justified.

The closest that Sosa ever comes to considering explicitly the internalist conception of having good reasons that I have been appealing to is in his section 9.4. There he adds to apt-justification and adroit-justification the idea of “Foley-rationality,”¹² according to which a belief is rational (and thereby in a sense justified) if “it would survive deep reflection by the subject in the light of her deepest epistemic standards (roughly).” In this last sense, as Sosa points out, the beliefs of a deeply superstitious person might well be justified, since those superstitions might be reflected in or partially constitutive of his or her “deepest epistemic standards.” Sosa seems to recognize, however, that none of these conceptions of justification is what the internalist is after: the first two because of their explicitly externalist character and the last because it is compatible with deep, albeit internally consistent and reflective irrationality. In my terms, the simple point is that a person might be justified in any of these senses while still not having any genuinely good reason for thinking that the beliefs in question are true.

But what then, Sosa asks, does the internalist have in mind: “what further internal state could we possibly be appealing to in supposing the [demon] victim to be internally, rationally justified?” His suggested answer is that this might mean that “he reasons well, and he well takes his experience into account, and he remembers well.”¹³ Though memory poses some special problems that I cannot go into here,¹⁴ this surely can be seen as pointing, albeit rather obliquely, to the idea that the person in question has *good reasons* for the beliefs that he arrives at. Sosa’s response, however, is that:

The problem is that these are all matters that derive from states of the subject and from the causal network of such states, or from subjunctive or other modal relations of such states to the subject matter concerned, which will include the layout of the external world.

12 Deriving, of course, from the work of Richard Foley, mainly in his books *The Theory of Epistemic Rationality* (Cambridge, Mass.: Harvard University Press, 1987) and *Working without a Net* (Oxford: Oxford University Press, 1993).

13 I have removed some italics here.

14 See my discussion in “The indispensability of internalism,” cited in n. 2.

What he is suggesting, I believe, is the thesis that the very notion of having a genuinely good reason for a belief that is internally recognizable is impossible or incoherent, in that the objective goodness of such a reason will always depend on external modal relations with the relevant subject matter, relations that cannot possibly be discerned or identified on a purely internal basis.

Here we have arrived at what seems to me to be the most fundamental issue between Sosa's view and my own. We may get a clearer grip on what is at stake by considering the more specific case that Sosa discusses, which also happens to be central to my main essay: the status of visual experience as providing good reasons for beliefs about the external world. Consider then, as an example, a "*visual experience as if there is something white and round before one.*" Such an experience, says Sosa:

is a reason for believing that there is such a thing before one, but only because in the actual world such a visual experience is reliably related to there being such a thing there . . . Most naturally one would want the state of the perceiver describable as "visual experience as if there is something white and round before him" to be necessarily such that it *would* normally reveal the presence of such a thing before the perceiver. Otherwise that same state would not be properly describable as such an experience.

I take this view to be essentially the same as Haack's view, discussed above in my main essay,¹⁵ according to which sensory experience should be described in terms of the external situations that are causally or lawfully (or modally) correlated with it. Sosa's suggestion is that only this sort of description captures the justificatory force of such experience in relation to claims about the external world. But the problem, as already noted in that earlier discussion, is that the applicability of such a description cannot be recognized from the inside in terms of the subjective character of the experience. Sosa, unlike Haack (who does not mean to adopt externalism), is quite content with this result, which as already indicated he takes to show (when suitably generalized) that there is no dialectical room for a genuinely internalist position – in my terms, that the very idea of a good reason for thinking that a belief is true that can be recognized from the inside is an impossible or incoherent one.

15 In section 5.2.

The only other option that Sosa can find to offer the internalist at this point is the idea that it is just a basic, not further explicable fact that sensory experience with a certain subjective character provides a good reason for a certain belief about the external world, in spite of the fact that there may be no general positive modal correlation between the occurrence of the experience and the relevant external fact; and indeed in a particular world (such as a demon world) may even be a negative modal correlation. I agree with him that this view¹⁶ is quite unacceptable, in that it would sever entirely the relation between justification (understood as the possession of a “reason” of this sort) and truth.

But what the internalist really wants, of course, is a view according to which the subjective character of sensory experience, as internally recognizable, provides a genuinely good reason, internally recognizable as such, for thinking that an appropriate modal correlation between such experiences and external situations of a particular sort genuinely obtains. It is this sort of view that was defended, albeit in only a provisional and schematic way, in the final chapter of my main essay, as it was also by Descartes (albeit via an intermediary appeal to God), Locke, and many others. It is admittedly (to say the least) not an easy view to defend successfully, which has led many over the years to opt instead for some version of idealism or more recently for externalism. For the moment, I will content myself with two remarks about the dialectical situation. First, I do not think that such an internalist view has been shown to be impossible, especially not in the general form that would extend beyond sensory experiences to other cognitive faculties, such as reason in particular. Sosa seems to me more to assume than to argue for such impossibility. Second, if such a view really is impossible in its general form, if we can never internally recognize good reasons for anything, then the skeptical implications are immediate and catastrophic – and in no way alleviated by the mere possibility, *which we will also have no accessible reason to believe to obtain*, that we still possess justification or knowledge in externalist senses. (I will consider this point a bit further in the discussion of skepticism below.)

Following up this last issue has taken us quite far from the discussion of Sosa’s virtue epistemology that was the official topic of this section. It does, however, put me in a position to make the central point

16 This is one interpretation of the view of Price and Chisholm, mentioned briefly as the second possibility in section 5.3 of my main essay.

that I want to make about his position. While it is quite unobjectionable in itself and quite possibly, like externalist views generally, interesting and valuable in relation to certain sorts of issues,¹⁷ it does not and cannot by its very nature speak to the issue that internalist views are primarily concerned with, which is, to repeat once more, whether we have good reasons to think that our beliefs, including any beliefs that we may have about our intellectual virtues, are true. And this would be true even if it were to turn out that a view that does speak to that issue cannot in the end be adequately stated and defended. It is in this way that a virtue epistemology of the sort that Sosa advocates is simply irrelevant to at least one central epistemological issue.

There is, however, as indicated earlier, one further dimension of Sosa's view to be considered, even though he says very little about it in the essay that I am responding to. In his earlier book *Knowledge in Perspective*,¹⁸ Sosa adds to the idea of "apt-justification" the further idea (mentioned also in his chapter 6) of a person's possessing a *coherent epistemic perspective*: very roughly, having a coherent body of beliefs, including both first-order beliefs and further meta-beliefs about his or her epistemic situation, such as beliefs about the reliability of his or her cognitive faculties. Indeed his view in that work seems to be that epistemic justification in the fullest sense requires both of these ingredients, that mere aptness is not by itself enough for justification.¹⁹ And here we have something that might seem to at least come close to satisfying internalist desires and intuitions. For the person whose beliefs add up to a coherent epistemic perspective and are also apt has both beliefs that are reliable, likely to be true, *and* a reflective awareness of that fact via meta-beliefs about the reliability of his or her cognitive faculties. What more could an internalist want?

But careful reflection will, I think, reveal that this picture is in fact much less satisfactory than it might at first seem. Such a person's beliefs add up to a coherent epistemic perspective, but that fact, as Sosa recognizes in chapter 6 (and elsewhere), is not in itself a good reason to think that they are true, basically for the same kinds of reasons that were considered in chapter 3 of my main essay. True, these beliefs are also, we are supposing, uniformly apt, but although the person by hypothesis *believes* this to be true, he or she also has no good *reason*

17 See the discussion in section 2.5 of my main essay.

18 *Knowledge in Perspective: Selected Essays in Epistemology* (Cambridge: Cambridge University Press, 1991).

19 See n. 14 in his chapter 9, and the associated text for a very brief indication of this view.

to think that any of these beliefs about aptness are themselves true. Thus it is unclear how the dimension of having a coherent epistemic perspective really adds anything to the variety of justification contributed by aptness alone, and we are still left with the conclusion that the person has no good reason to think that any of his or her beliefs are true. The basic problem is that aptness and coherence do not seem to combine in any clear way to yield any sort of justification beyond whatever they would each yield independently.

10.3 Sosa's Critique of the Appeal to the Given

In chapter 7, Sosa offers an extended critique of the internalist-foundationalist idea that foundational beliefs can be justified by appeal to the features of conscious experience, where I will understand this to mean that such experience can provide a good, internally accessible reason for thinking that such beliefs are true: essentially the same view as the one that I offered in chapter 4.

Since the dialectical path that we will have to traverse to come fully to grips with Sosa's criticisms is a bit complicated, it will help to have his conclusion clearly in front of us from the beginning. His thesis is that apart from cases where the content of the belief is too thin to be interesting, it is impossible for conscious experience to provide a complete, internally accessible reason for thinking that beliefs about it are true. Instead, the justification of any belief about experience with enough content to be interesting will always require an externalist appeal to reliability or to intellectual virtue. And if Sosa is right about this, then an internalist-foundationalist view of the sort offered in my main essay is after all doomed to failure.

The starting point of his discussion is a distinction between two sorts of awareness that one might have of a feature of one's experiences: (1) intellectual awareness or noticing ("n-awareness"), which involves believing or judging that feature to be present, and further requires that the belief in question be in some way justified or reasonable; and (2) experiential awareness ("e-awareness"), the sort of awareness that one has of the content of one's experience simply in virtue of having or undergoing it. The initial problem is then to understand in these terms how foundational justification is supposed to work. When the foundationalist appeals to a person's direct or immediate awareness of the content of his or her experience as a basis for foundational justification, the appeal cannot be to n-awareness, since this already presup-

poses beliefs (or judgments) that are justified. But to appeal to e-awareness does not work either, Sosa argues, since it is possible to be e-aware of an experience with a certain feature and still fail to be justified in believing that the experience has that feature. To show this, he appeals to Chisholm's famous case of the visual awareness of a speckled hen: assuming that the speckles are all clearly above some threshold of size, it is clear that one's visual image of the hen will have some definite number of speckles; but equally clear that if the number is fairly large, one can be e-aware of that image without thereby being justified in believing that the image has some specific number of speckles.²⁰ I think that this example does clearly establish the point, though, as we will see, nothing this elaborate is really required for this purpose.

I will first sketch a response to this general problem, in relation to which I can then proceed to consider Sosa's further objections. On my view, as set forth in my main essay, a foundational belief results when one directly sees or apprehends that one's experience satisfies the description of it offered by the content of the belief. Let me first set out a somewhat detailed example and then attempt to elaborate this view in a way that responds to Sosa's concerns.

Suppose then that I am standing directly in front of and looking at (with good lighting) a large abstract painting, containing a fairly large number of geometrical shapes of various sizes and colors, and bounded by a sharply delineated, rectangular black frame, where the resulting visual image is entirely contained within my visual field. The proposition is suggested, perhaps by someone else, that within the black rectangle in my visual field is a dark green, approximately equilateral, approximately triangular shape of sufficient size to be easily discernible. At first, given the size and complexity of the painting, I am unsure whether or not the proposition is true, but after some hunting around I spot such a shape and proceed to form the corresponding belief, a belief that, I am claiming, is thereby justified in the internalist sense I have been advocating throughout this reply: I have a good, indeed an excellent reason to think that the belief in question is true.

Let's analyze this example in Sosa's terms. Initially I have an experience, a visual image, that has the feature in question, which is to say that I am e-aware of that feature, along with the rest of the complicated image. But at the initial stage I am not n-aware of that feature: I have no belief that such a feature is present, and if I did, it would not be

20 See R. M. Chisholm, "The problem of the speckled hen," *Mind*, 51 (1942), pp. 368–73.

justified simply by virtue of my e-awareness. Instead I have first, as it might be put, to *attend* to various aspects of the experience, sorting through its details until I find the specific feature that I am looking for. This attending is not to be identified with becoming n-aware, but rather precedes and results in n-awarenesses of various sorts. Moreover, it is, as Sosa himself insists, quite possible to attend to a feature and still fail to form a belief, an n-awareness, that fully captures its character. The right thing to say, it seems to me, is that attending is not a distinct form of awareness from e-awareness, but is rather built upon it via what might be described as selective focus.

A bit further on in his discussion (section 7.3), Sosa considers such an appeal to attention. His response is that attention is insufficient to yield justification because it would be possible to (a) be e-aware of an experience with a certain feature, (b) believe that one's experience has that feature, and (c) attend to the relevant aspect of the experience, while still (d) not being justified in holding that belief.²¹ I agree that this is possible, though I am not convinced by the specific example that Sosa appeals to in trying to establish this point. He imagines a case in which a person who believes that his or her experience has a certain feature (containing a decagonal shape) still fails to possess a "phenomenal concept" of that feature (more on this idea below) and so is unable to recognize that it is present, despite attending to the relevant aspect of the experience. I doubt that such a case is possible because I think that it is essential to having a specific belief about the character of experience that one understand what an experience would have to be like in order to satisfy the description embodied in the belief (i.e., what it would be for the description to be true) and so have all of the relevant concepts. In my view, the reason that a case with features (a), (b), (c), and (d) is possible is rather that one can have an experience with a certain feature, understand completely what an experience has to be like to be correctly described as having that feature (and so possess all of the relevant concepts), attend to the relevant aspect or aspects of the experience, and still sometimes be unable to tell whether or not the description is actually correct, where at least the most obvious sort of case in which this happens is where the description is sufficiently complex as to make it hard to tell whether it applies even when the relevant feature is being attended to. Here the image of the speckled hen (does it contain exactly 51 speckles?) or Sosa's example of an 11-membered array of dots can serve as examples.

21 I have added the label (d) to Sosa's (a), (b), (c).

So what then needs to be added to conditions (a), (b), and (c) in order to yield internalist justification for the belief? The answer, I suggest, is quite simple: the person must apprehend or recognize the agreement or fit between the aspect of experience being attended to and the conceptual description given by the belief. In the earlier example, I have to apprehend or recognize the agreement between the aspect of the visual image upon which I have focused and the conceptual description of it as “a dark green, approximately equilateral, approximately triangular shape of sufficient size to be easily discernible.” In the case described, I am able to do this without much difficulty once I have succeeded in attending to the relevant bit of the image, but this might not be so if the conceptual description embodied in the belief were more complicated (where one important sort of case that I am including under complexity is a wholly or partially negative description: “a dark green, triangular shape that is below a yellow circle and not equidistant between two purple squares”). In cases where people fail to see or apprehend such agreement, the corresponding belief will not be justified, but this, I submit, has no bearing at all on the other cases where they are successful.²²

One possible objection to such a view, though not one that Sosa raises, would focus on this apprehending or recognizing and insist that it must be regarded as a further, independent cognitive state, just as much in need of justification as any other, and so of no help to the foundationalist. But any such view seems to me to be mistaken (as indeed would an analogous move in relation to attention). What is crucial here is the nature of e-awareness itself. When I am e-aware of an experiential feature, it is genuinely present in my consciousness. Because many other things are present there as well, I may have to selectively focus in order to pick out that feature and come to see that its character fits a particular description, but the awareness of it that allows me to do these things is just the original e-awareness itself. To insist that a distinct cognitive act, requiring an independent justification, was involved would be in effect to say that e-awareness in itself has no cognitive significance at all, that I can be genuinely conscious of the feature in question but be unable to do anything further with that consciousness. I can see no reason to accept such a view (even though I must admit that I was once tempted by it).

22 This seems to me essentially the same view that Fumerton has in mind in the account that Sosa discusses in his section 7.4, though I do not find the formulation in terms of “acquaintance” especially perspicuous (see my discussion in my section 4.4.)

Sosa's basic objection, as I understand it, is rather different. What he is demanding is in effect a general set of criteria, formulated in terms of the descriptive contents of the beliefs in question, that will distinguish the cases in which such an apprehension or recognition of agreement occurs, resulting in justification, from cases like the case of the speckled hen image or the 11-membered array of dots where it does not. I am inclined to agree that no such general account can be given, beyond the one that appeals simply to the fact of apprehension or recognition itself. I have described the main difference between the two sorts of cases in terms of "complexity," but there are many dimensions of complexity, and I see no reason to think that they can be sorted out into a neat general formula. What I fail to see is why such a formula is required or why any real objection follows from the inability to provide one.

One way to defend my view on this point is to appeal to an analogy with a more commonsensical sort of situation. If we set aside distinctively philosophical concerns about the problem of the external world and the like, it seems an obvious truth of common sense that people who are witnesses of physical events and situations are frequently able to formulate conceptual descriptions of them and to be justified in the resulting beliefs. But it also seems obvious that merely the fact that someone was a witness to an event or situation in plain view is not enough to guarantee or in general even make especially likely that the beliefs, even the true beliefs, that they may form about the observable aspects are thereby justified. (I am not interested here in public claims or public evaluation of claims, only in the beliefs themselves.) It all depends on whether or not the person is in fact able to apprehend or recognize clearly that the description embodied in a candidate belief is indeed realized by the event or situation, and here again complexity of various sorts provides an obvious explanation of why someone might not be able to arrive at such a recognition. Moreover, I doubt very much that it is possible to give any set of general criteria (beyond the appeal to apprehension or recognition itself) that will distinguish the cases in which beliefs are justified in this way from those in which they are not. But this surely has no tendency at all to show that beliefs are not frequently justified on just such a basis. And the same thing seems to me to be true for the foundationalist's beliefs about the content of experience. (Another feature that the two sorts of cases have in common is that the capacity for the relevant sort of apprehension or recognition will vary fairly widely from person to person, from one stage to another in the life of a given person, and no doubt on the basis

of other factors such as degree of effort or degree of fatigue as well, making the search for general criteria even more obviously hopeless.)

I have largely bypassed so far one major element of Sosa's discussion, namely his distinctions between different sorts of concepts that one might apply to experience: indexical concepts, "phenomenal concepts," and "simple geometric and arithmetical concepts" ("SGA concepts") (sections 7.2 and 7.3). I am doubtful, however, that this threefold distinction really sheds very much light on the issue we have been discussing.

Consider first the idea of "indexical concepts." While I have no doubt that we often pick out portions of our experience indexically, I doubt very much that we ever form fully indexical beliefs whose content could be captured in the form that Sosa suggests, beliefs that *this is thus*, with the predicate as well as the subject being indicated only in indexical terms, or that there would be any cognitive point at all to forming such a belief. Sosa says that the content of such a belief is "thin" or "minimal," but it seems to me to be essentially zero. It is of course possible that one might employ such a formulation while having some fuller idea of the predicate at least and probably of the subject as well in mind, but then the formulation would not really capture the content of the belief. Thus I propose to set aside the idea of indexical concepts that would figure as the predicates of such indexical beliefs. (Sosa seems to agree at least that they have no real epistemic importance.)

What then is the distinction between "phenomenal concepts" and "SGA concepts" supposed to amount to? According to Sosa, a phenomenal concept (Sosa's main example is the concept of a distinctive hourglass-shaped pattern of dots) is "defined in part by sensitivity to the relevant feature of which it is a concept" and "in part by the ability to tell when that feature is present and when [it is] absent in our experience"; whereas an SGA concept is not thus defined. Sosa's suggestion is that it is the fact that the concepts are defined in this way that accounts for the justification of beliefs involving phenomenal concepts, whereas no analogous justification can be given for beliefs involving SGA concepts. And without beliefs of the latter sort, the content of our justified beliefs about experience will be, he argues, too thin to allow the inference of much beyond them.

I doubt very much, however, that there is any genuine distinction of this sort to be drawn, because I do not think that there really are any "phenomenal concepts," as thus understood. My contrary suggestion, though there is no space for a full defense of it here, would be that no

descriptive concept of any sort is ever *defined*, even “in part,” in terms of an ability to apply it. Any such concept, no matter what it applies to, is defined *entirely* by its descriptive content, by what something of the relevant sort must be like for an application of the concept to be true or correct. If this descriptive content is sufficiently simple and straightforward and if a person has adequate access to the relevant object or situation, then it may be very easy to recognize that it does correctly apply, so much so indeed that failure to be able to recognize this may constitute excellent *evidence* that the content of the concept has not been adequately understood. But this, I would suggest, is never simply a matter of definition. And for *almost* any concept, it is possible to imagine circumstances in which its correct application would be difficult and hence uncertain. (Thus imagine trying to decide whether or not Sosa’s example of the concept of the hourglass-shaped pattern of dots is satisfied somewhere in a complicated visual array like the experience of the painting described earlier.) All this applies both to concepts that are arithmetical or geometrical in character and to those that are not. Indeed, it seems obvious that there are many clearly arithmetical and geometrical concepts that are as easily applied to experience as almost any others. What seems true is that concepts of this general sort make it especially easy, though not I think uniquely so, to arrive at concepts of sufficient complexity as to make it difficult to recognize that they apply. There is, however, as far as I can see, no sharp distinction here, but only a continuum in many dimensions.

My conclusion is that Sosa’s discussion in his chapter 7 yields no real objection to the sort of foundationalist views offered in my chapter 4. Foundational beliefs about the content of experience are justified when the person is able to recognize the agreement between the content of such a belief and the experiential content of which he or she is aware, and there is no need for any general account in terms of the content of the belief as to when this will be possible and when it will not. Nor is there any reason to think that beliefs involving arithmetical or geometrical concepts will not often be justified in this very same way. Thus no appeal to safety or virtuousness (see Sosa’s section 7.7) is needed to give an account of when such beliefs are justified. This is not to deny that the relevant beliefs will in general be both safe and virtuously formed, as Sosa explains these concepts. But their having these features will be a consequence of the way in which they are justified, not part of the main account of that justification. And, much more importantly, the possession of these two features by a belief is arguably recognizable by the believer “from the inside” on the basis of

the way the belief is justified, rather than being merely external properties to which he or she need have no access.

10.4 Some Further Reflections on Skepticism

Though the space available for this reply is nearly exhausted, I want to conclude by saying a bit more about the issue of skepticism, contrasting Sosa's view (to be found mainly at the end of his chapter 9, especially in section 9.5) with my own. Here the dialectical situation is very complicated and needs to be sorted out quite carefully.

Imagine that we are considering a belief that is included in one of the various categories of beliefs that common sense confidently regards as justified²³ and as constituting knowledge. What can Sosa's view say about the epistemic status of such a belief? First, it can say that it is *apt-justified*, i.e., that it results from one or more faculties or capacities that yield a high proportion of true beliefs and thus count as intellectual virtues. Or rather, for reasons that will emerge more fully here but have already been suggested above, it can say that such a belief *may* be apt-justified. Second, it can say that the person in question has a *cognitive perspective* in relation to such a belief, consisting of meta-beliefs about the ways that it was acquired and the virtuousness of those ways, higher-level meta-beliefs of the same sort about the first level of meta-beliefs, and so on, up to some highest level whose beliefs are not themselves the objects of further meta-beliefs of this sort. These various beliefs *may* again all be apt-justified. In addition, they *seem* to fit together with each other and with other first-level beliefs to make up a comprehensively *coherent* system and *may* actually have this status, though whether this is so or not depends on whether the various sorts of seemingly a priori insights that are ingredients in the concept of coherence are themselves apt-justified.

The existence of this apparently coherent epistemic perspective may even seem to give the person in question a *reason* to think that the original first-level belief is true. But if Sosa's account is the final story, this last seeming is definitely illusory: even genuine coherence would not by itself yield such a reason, still less mere apparent coherence; and the reason that the hierarchy of meta-beliefs might seem to yield is not genuine because there is no such reason for the beliefs at the highest

23 Assuming here that common sense can be ascribed clear views about justification (as opposed to rationality or the possession of reasons).

level, upon which all the rest depend. Moreover, this last result applies not only to the person who has the belief in question but also to any third-person assessment of his epistemic situation: the person making such an assessment also in the end has no genuine reason for thinking that it is correct, since his or her apparent reasons also depend in the end on a highest level of meta-beliefs for which there are no reasons and on coherence that may be only apparent. Thus the original belief *may* be apt-justified and *may* also be justified in the fuller sense that includes a genuinely coherent epistemic perspective all of whose elements are apt-justified, but neither the person who has the belief nor anyone else has *or even could possibly have* any genuinely cogent reason for thinking that either of these things are so.

Contrast this picture with the one offered by the pretty traditional version of foundationalism that was defended in outline in my main essay. On that view, there are at least two sorts of beliefs that are genuinely foundational in the sense that a person can have cogent reasons for thinking that they are true that do not depend on other beliefs in need of justification: (1) a priori beliefs about self-evident truths (including principles of inference of various sorts) and (2) beliefs about the conscious contents of one's own states of mind.²⁴ And the main foundationalist conviction, which seems to me to be entirely correct, is that only by beginning with beliefs that have this sort of foundational status is it possible to ever have good reasons to think that any other beliefs are true. If such a foundational status is ruled out, in the way that Sosa's view seems to do, then no amount of higher-level reflection on one's belief can ever yield genuine reasons.

But even if the foundation is as secure as I believe it to be, there is notoriously a wide gap that must be crossed in order to arrive at reasons for the various sorts of non-foundational beliefs that seem, commonsensically, to be justified and to constitute knowledge. In my main essay, I have offered a sketch of how this transition might go for what is perhaps the main category of non-foundational beliefs, namely beliefs about material objects and situations of the sort that seem to be justified on the basis of perceptual experience. Even if that sketch is approximately in the right direction, however, filling in the rest of the details is extremely difficult at best. And there are in addition other sorts of beliefs, beliefs about laws of nature, about unobservable scientific entities and processes, about the historical past, about the

24 Though I have offered an account of only the latter of these two sorts of foundational belief in my main essay here. For my account of the former one, see *IDPR*.

mental states of other people, etc., for which the story would have to be even more complicated and difficult.

Thus there are, I suggest, two major concessions that any foundationalist must make. First, while the exact degree of optimism or pessimism that is warranted could be debated endlessly, it is impossible at present to be rationally sure or even fairly confident that the foundationalist picture can be completed in a way that will yield genuinely cogent reasons for the various beliefs of whose positive epistemic status common sense is so confident. (I am quite aware that many will think that this understates enormously how unlikely such a result really is.) Second, even if such success were to actually be obtained, it is quite unlikely that the details or even the approximate outlines of the reasoning that would be involved could be said to be accessible to ordinary believers in more than the dimmest, most implicit way. (And here again many will think that this seriously understates the difficulty.)

Where then does this leave us? We have one account of the justification of our beliefs that may perfectly well be realized, but (1) that does not involve our having good reasons for thinking that those beliefs are true, and (2) that we could by its own lights never even in principle have a good reason to believe to be realized. And we have a second account that could in principle yield good reasons for our beliefs and even for thinking that we have such reasons, but (1) that may very well not ultimately be realizable, and (2) that would not in any case yield reasons that ordinary people could plausibly be said to have any very clear awareness of. And we have the commonsense convictions that our ordinary beliefs are justified and constitute knowledge. Which view of justification is more acceptable in light of this conviction? (Many will no doubt think at this point that there must be some third alternative that is less dire; I can say only that I can find no dialectical space for such an alternative that does not simply evade the issues.)

Though this is admittedly a difficult choice, I can think of two relatively clear reasons that seem to me to favor the foundationalist view at this point, and I will conclude this reply by giving a brief account of them. Both turn on the further issue of just what the commonsense conviction in question really amounts to or more specifically of what specific conception of justification common sense has "in mind" in holding that the beliefs in question are justified.

First, it seems to me quite doubtful that the commonsense conviction is that our beliefs are justified in the specific, quite complicated sense that Sosa has delineated, because it is doubtful that common sense really has any inkling of that specific conception of justification;

and only slightly less doubtful that the commonsense conviction involves a conception of justification that is sufficiently generic or unspecified to include Sosa's conception as one possible specific realization without picking it out specifically. On the contrary, the commonsense conviction seems to me to be simply that we have good reasons for thinking that our beliefs are true, in which case (1) Sosa's account does not genuinely accommodate that conviction and (2) only a foundationalist view could, for reasons already considered, genuinely do so.

Second, suppose that I am wrong about this first point and that the commonsense conviction really does involve a conception of justification that either agrees specifically with Sosa's or at least includes it as one possibility among others. It remains the case that if Sosa's brand of justification is the only sort that is genuinely available, then just as we can have no good reasons for other beliefs, so also we can have even in principle no good reasons for this commonsense conviction itself. Thus the sense in which Sosa's view can accommodate the commonsense conviction turns out to be extremely Pickwickian at best: if Sosa's account is correct, then while the commonsense conviction *may* be true, we can in principle have no reason to think that it is true – and so also in principle no reason to think that being able to accommodate this conviction counts in favor of an epistemological theory.

Can the foundationalist view do any better in this last regard? *If* (1) the details of such a view can indeed be worked out successfully, and (2) the resulting lines of reasoning can after all be said to be dimly or implicitly within the grasp of common sense, then foundationalism can both interpret the commonsense conviction in a way that allows us to have good reason to think that it is true and accommodate the conviction as thus understood. And this result, though obviously less than fully satisfactory from the standpoint of common sense, still seems clearly better than anything that Sosa's view can offer. I thus conclude that the commonsensical implausibility of skepticism turns out to offer no good reason at all for preferring Sosa's view of justification (or indeed any broadly externalist view) to traditional internalist foundationalism.

Reply to BonJour

Ernest Sosa

11.1 Introduction

According to Roderick Chisholm one engages in “epistemology” in order to assess and improve one’s body of beliefs, and one does so in the armchair. Sustained by faith that we can succeed, we seek principles for belief formation that will help us do so. Our data for selecting such principles derive, moreover, from intuitive convictions about what is or is not epistemically justified. Far from being idiosyncratic to Chisholm, that view of epistemology also fits the founder of modern epistemology impressively well. Descartes was indeed out to improve himself epistemically as he sat down to his meditations, and he wished to do so by reflection. And he did take himself to have reflective access to certain of his (subjective, mental) properties. And he did think that some of our beliefs may be seen intuitively and reflectively to be thus justified. Early in the Third Meditation, for example, he attributes to himself certainty about the *cogito*, and on that basis takes himself to establish a way in which a belief can gain such certainty, namely by according with what to the subject is then clear and distinct.

Both Chisholm and Descartes require special access to one’s own justification and its sources if that justification is to attach to one’s belief at all. In recent years this view of justification has been championed by Laurence BonJour. He, too, highlights some such requirement of special access. Like Descartes, moreover, BonJour is primarily interested, not so much in whether we normally enjoy epistemic justification for our ordinary beliefs, as in what, if anything, *would* give us better epistemic justification, or even the best epistemic justification. Whether our actual beliefs are already rationally justified is negligible next to the question whether we could ever hope to attain such justification, even at great intellectual expense.

The question posed by BonJour is whether we can have a good rationale (an accessible one) for our empirical beliefs, particularly for our beliefs about the external world. How surprising it would be, however, to find such curiosity detached from any interest in replacing ill-supported beliefs with some that are better supported, an interest that more explicitly motivates Chisholm and Descartes.

But why should we restrict sources of epistemic justification to those accessible in the armchair? According to Chisholm, we must look for criteria that “will tell us something about the conditions under which *S* is *justified* in believing that there is an *F*. The conditions in question will not themselves be normative facts: they will be non-normative facts (say, being appeared to in certain ways) which constitute sufficient conditions for the existence of certain normative facts.”¹ Later in that same book, he adds the following.

The usual approach to the traditional questions of theory of knowledge is properly called “internal” or “internalistic.” The internalist assumes that, merely by reflecting upon his own conscious state, he can formulate a set of epistemic principles that will enable him to find out, with respect to any possible belief he has, whether he is *justified* in having that belief. The epistemic principles that he formulates are principles that one may come upon and apply merely by sitting in one’s armchair, so to speak, and without calling for any outside assistance. In a word, one need consider only one’s own state of mind. But if we look at the matter, it seems clear that the approach to the questions of the traditional theory of knowledge can only be thus internalistic . . .²

Recent epistemology has witnessed a resurgence of the classical foundationalism that we are said to receive through Chisholm ultimately from Descartes. We find such foundationalism in main recent writings of Richard Feldman and Earl Conee, Richard Foley, and Richard Fumerton. BonJour, too, is part of this movement, with meaty papers in which he renounces his earlier coherentism and embraces classical, internalist foundationalism. While diverging from Chisholm in important respects – over what we can hope to believe justifiably about the world around us, for example, and over the conditions required for that

1 *Theory of Knowledge*, 3rd edn. (Englewood Cliffs, NJ: Prentice-Hall, 1989), pp. 42–3.

2 *Ibid.*, pp. 76–7.

happy outcome – nevertheless he forthrightly acknowledges his fundamental agreement with Chisholm.³

11.2 The Main Structure of BonJour's Monograph

The main structure is as follows.

(a) The spotlight is first trained on rational justification, with knowledge neglected as of subsidiary interest at best. The reasons given for this neglect fall short, however, since they seem initially to apply to rational justification just as much as to knowledge. According to the main reason offered, there is no way to specify a sharp threshold of justification requisite for knowledge.⁴ Unfortunately, it seems about equally difficult to specify any sharp threshold of justification requisite for justification. And if we *were* told (which we are not), along Chisholmian lines, that the sharply defined concept of justification is that of being more justified in believing than in withholding, we could plausibly adopt that sharp concept of justification as requisite *not only* for justification (beyond reasonable doubt) but *also* for knowledge.⁵

3 Among classical foundationalists, it is of course not only Chisholm and BonJour who disagree. No two of Descartes, Chisholm, BonJour, Conee and Feldman, Foley, and Fumerton agree on all important points. (In saying that, I do not distinguish between Conee and Feldman, who co-author the papers that put them in the circle, papers that defend the general outlook of internalist foundationalism, with an interesting twist in their recent rejection of access internalism in favor of mentalism, the view that the epistemically normative supervenes on the individually mental.) Only a family resemblance joins them together. And in important respects it is BonJour who most closely resembles Descartes and Chisholm, in an "inner-circle" of classical foundationalism. Of the others, Fumerton is closest, with Conee and Feldman, and Foley, somewhat more distantly related, though they all do definitely belong.

The monograph by BonJour in the present volume, to which this is a response, combines and develops a view presented in a series of papers over recent years. See L. BonJour, "Foundationalism and the external world," *Philosophical Perspectives*, 13: *Epistemology* (1999), p. 235. Also his "The dialectic of foundationalism and coherentism," in J. Greco and E. Sosa (eds.), *The Blackwell Guide to Epistemology* (Oxford: Blackwell, 1999), p. 134. And "Toward a defense of empirical Foundationalism," in M. DePaul (ed.), *Resurrecting Old-Fashioned Foundationalism* (Totowa, NJ: Rowman and Littlefield, 2001), along with the replies to John Pollock and Alvin Plantinga in that volume.

4 See BonJour's section 1.5.

5 Besides, if philosophy really required sharp concepts, that would condemn a lot more than just the theory of *knowledge*. In fact, it would be interesting to go through this very monograph of BonJour's with that red pen in hand. Take the concept of belief, for example. How sure must you be in order to count as a believer? Well, it may be said, let us give up on belief as well. Let's just stick to the mental act of judgment. But even judgment in the sense of occurrent, conscious assent is itself a state that admits of degrees.

Nevertheless, in these comments I will focus on justification, and will rarely refer to knowledge or *its* requirements.

(b) Three positions are specified in terms of the following two contrasts: that between internalism and externalism, and that between foundationalism and coherentism. These are said to be “dichotomies,” but when one looks at the definitions,⁶ it is clear that they are not, at least not *as defined*; they are at most contrasts, and I will in what follows regard them as such.

The three positions that serve to structure the monograph are: *externalism* (externalist foundationalism), *coherentism* (internalist coherentism), and *foundationalism* (internalist foundationalism). After an introductory chapter that lays out the problematic in terms of rational justification and puts aside knowledge, three chapters are devoted respectively to the three specified positions. One each attacks externalism and coherentism, and a third defends foundationalism. A concluding, fifth, chapter further defends foundationalism by attempting to show how it can escape external-world skepticism. That chapter aims to reach the conclusion that there is such a world, through valid reasoning in line with classical foundationalism.

11.3 The Attack on Externalism

Here first is a précis of the anti-externalism argument.⁷

No external condition, beyond our cognitive purview, could suffice for rational justification. The point is urged especially against the proposal of reliable belief formation as such a condition,⁸ but it could be made equally well against a Nozickian tracking requirement, or any

I can assent with greater or lesser assurance. And the degree of my assurance would seem internally related to whether my conscious state counts as one of assent. So how much assurance is required? Is there a sharp threshold?

Or take “reasons.” At some point we must go beyond reasons, whether pro or con, to the *balance* of reasons, and to whether this is weighty enough to “justify” belief, i.e., to make believing more reasonable than withholding. All right, but how weighty is that? Is there some dimension with a threshold below which your belief is not justified? It might be argued that the relevant threshold is adequately picked out by it being the point where believing is more reasonable than withholding. But then what is to stop the advocate of *knowledge* as a viable concept from adopting that same threshold, or from declaring that his relevant threshold is that of being sufficiently justified to know?

6 These are found in the first few pages of BonJour’s chapter 1.

7 As it appears in section 2.3.

8 As in Alvin Goldman’s “What is justified belief?” in G. Pappas (ed.), *Justification and Knowledge* (Dordrecht: D. Reidel, 1979).

other brutally external causal requirement. If, for example, one were suddenly blessed with clairvoyance, the beliefs deriving from its exercise would not become rationally justified. To the contrary, if all our evidence pointed against our having any such faculty, it would be irrational for us to accept the deliverances of clairvoyance at face value, no matter how reliable it might in fact turn out to be. Some externalists respond by granting that external reliability is not by itself sufficient, while adding a meta-requirement that the subject *not* have weighty evidence against their having any such faculty. But that strategy is effectively countered in this chapter through a case of someone suddenly blessed with clairvoyance in the *absence* of *sufficient* evidence one way or the other. This would be compatible with the person having some evidence in favor, counterbalanced by the evidence against, which would preclude their rationally opting either for the positive or for the negative on the matter at hand. So Bonjour is led to the requirement that one must have adequate positive reason in favor of a belief and not just lack weighty negative reason against it. This is what he supposes the externalist would have to accept, so that the concept of justification becomes bifurcated, with both externalist and internalist components. And this is said to be very odd, lacking any evident rationale, and crushed by the problem of how to make non-arbitrary trade-offs between the two constitutive sorts of factors, the external and the internal.

As a first reaction, I must record my doubts about that negative assessment of “bifurcated” evaluative concepts. If we reject such concepts, then more than bifurcated epistemic justification would have to go, a lot more. Indeed, not many ordinary evaluative concepts would seem unaffected by the fact that it is sometimes hard to make non-arbitrary trade-offs, that conflicts of values and hard choices so frequently trouble our lives and common endeavors. It is not easy to see why a “bifurcated” epistemic concept of justified belief, one admitting a variety of considerations in the epistemic assessment of belief, would be *particularly* objectionable through its tendency to require such choices.

There is also something rather odd about the line of reasoning used against externalism. Objections against the sufficiency of reliability for justification are first said to be more fundamental than objections to its necessity. The reason proffered is that objections to the necessity of reliability, such as the new evil demon objection, have no direct force against the externalist proposal that reliability is sufficient for justification. But this does not show the sufficiency issue to be *more*

fundamental since it is mirrored by the fact that objections to the sufficiency of reliability, such as the clairvoyance cases, have no direct force against the externalist proposal that reliability is *necessary* for justification. Thus even *if* we agree with Bonjour's conclusion that externalists will need some appeal to positive internal reasons, and cannot rest with mere reliable belief production, as with clairvoyance, *still* this does not rule out that reliability be necessary in addition to any such positive internal reason. This is indeed the position occupied by William Alston, who in papers including "An internalist externalism" has argued for precisely that sort of theory.⁹ According to Alston, clairvoyant beliefs that pop in out of the blue are unjustified since they are not properly based on any conscious state. In order for a belief to be justified, moreover, not only must it be based on such a ground; in addition, the ground needs to be reliably truth-conducive. But Alston insists that this fact can help justify one's belief without itself being grasped.

Clearly, Bonjour would reject this position of Alston's as unacceptably externalist. So the sort of externalism that he opposes is not restricted to claiming some external condition, such as reliable belief formation, to be *sufficient* for justification. Alston's externalism does *not* claim reliability, or any other external condition, to be on its own sufficient. And yet it is still opposed by Bonjour, surely. Here then is a better, more general, statement of the sort of externalism opposed:

An epistemology is *externalist* if and only if it entails that some factor can add essentially to the epistemic justification of a subject's belief even though it falls outside the reflective purview of that subject.

11.4 The Attack on Coherentism

We start, again, with a précis.

First, by appeal to various familiar lines of reasoning, it is shown just how implausible *pure* coherentism is, how implausible it is that epistemic, rational justification for beliefs about an external world should derive merely from internal relations among a subject's beliefs, with no requirement of any kind of input from, or sensory sensitivity

9 W. P. Alston, *Epistemic Justification: Essays in the Theory of Knowledge* (Ithaca, NY: Cornell University Press, 1989), esp. Part III: Internalism and Externalism.

to, the subject's surroundings.¹⁰ So it is allowed that probably no one has actually advocated a perfectly pure coherentism, and that a moderate coherentism is much more plausible. While giving some important role, maybe the most important, to coherence, the more moderate coherentist also countenances one or another further factor as determinative of justification, somehow in combination with coherence.¹¹

It is surprising, therefore, that BonJour allows no such additional factor to play any independent justifying role. Whenever a factor is entertained in his chapter as offering a way to develop a more moderate coherentism, it is argued that for a coherentist such a factor could have epistemic justifying effect only through coherence-inducing reasoning that sustained its reliability by means of internal reflection.¹² How can any such requirement be imposed at this stage, however, if it so obviously leads back to the position already rejected as intolerably implausible, the position that at bottom it is *only* internal coherence that does the epistemic work? The problem is that internal coherence will not do the job just by itself, whether directly or indirectly, not when the beliefs under justification are empirical beliefs about the independent external world around us. We need the help of some additional justifying factor with its own status, a status that it must not in turn owe entirely to internal coherence.

Here are two passages that impose only the weak coherentist requirement:

[Coherentism is] . . . a view according to which (1) there are no basic or foundational beliefs and (2) at least the primary basis for empirical justification is the fact that such beliefs fit together and support each other in a variety of complicated ways, thus forming a coherent system of beliefs . . .¹³

In fact, largely for the reasons just noted, there is probably no one who has ever seriously advocated a *pure* coherence theory of empirical justification, one in which the coherence of a set of beliefs is claimed to be by itself sufficient for justification. The historical coherentist project has rather been, in effect if not very explicitly, to supplement the appeal to coherence in a way that avoids or at least mitigates [certain objections to

10 See the three introductory paragraphs of chapter 3.

11 What is more, foundationalism itself must admittedly allow some room for coherence, at least as *one* important factor; if one wishes to go beyond the foundations to some superstructure of beliefs, one will need the help of coherence.

12 As this is presented mainly in chapter 3.

13 From the first paragraph of chapter 3.

radical coherentism] . . . while at the same time avoiding a relapse into foundationalism.¹⁴

In the introductory passages of chapter 1, foundational beliefs had been defined as “beliefs whose justification does not depend at all on that of other beliefs.” Accordingly, in order to avoid any relapse into foundationalism one need only reject the possibility of beliefs fully epistemically justified independently of support from other beliefs. Compatibly with this, however, one can still allow that a factor might help justify a belief by acting conjointly with the support of other beliefs. So in these passages it is allowed that there can be a source of epistemic justification that contributes *essentially* to the epistemic standing of our beliefs without this contribution depending on coherence with or support by some of the subject’s other beliefs. According with experience, for example, can be a factor essential for the justification of a belief, even though the support of other beliefs is *also* required for the justification of that particular belief. (Of course we would need to make more precise the way in which “according” with experience and with other beliefs can be respectively “essential.”)

Admittedly, in that chapter the core of internal coherentism is occasionally specified in a way that *seems* to allow for some such additional factor, one that will contribute on its own to the epistemic work done by sheer coherence. But the form of critique directed against coherentism in general belies that concession, and presupposes that any such factor must depend fundamentally on pure coherence. Indeed, if we insisted on the more moderate versions of coherentism, then BonJour’s own foundationalism itself resembles a variety of coherentism. After all, he had earlier granted that the foundationalist will need appeal to coherence in moving beyond his foundations. But in that case the coherentist could just return the favor by allowing the taking of the given as the further factor that he for his part needs in addition to coherence, thus bypassing the objections to the pure form of his doctrine.

I say that BonJour’s own foundationalism *resembles* a variety of coherentism. What still distinguishes the two is that such foundationalism holds foundational beliefs to be *exclusively* epistemically justified by a factor neither constituted nor essentially aided by the support of, or by coherence with, any other beliefs held by the subject. Even moderate coherentism will deny that beliefs are adequately justified

14 From the third paragraph of chapter 3.

simply by belief-independent factors. Given that beliefs would not so much as exist without an extensive supporting cast of related beliefs, there is an air of unreality about the foundationalist claim that beliefs might nevertheless be justified independently of other beliefs. It is hard to conceive of the hypothetical cases that one would naturally invoke in support of such a claim, for these would be cases where one held the target beliefs along with the supporting conscious states but without the supporting cast of other beliefs. But you could not possibly so much as host the target belief without a lot of the relevant supportive beliefs. Nor does it seem that you could enjoy justification for the target belief in the absence of justification for a good number of those supporting beliefs, absent which you could not hold the target belief at all. If one nevertheless insists that, despite this, in some sense the other beliefs do not help justify the target belief, one will surely be asked to explain this special relation of *justifying* that can fail to relate a belief *X* and a belief *Y* even when belief *Y* would not be justified (not as fully, or as well) in the absence of support by belief *X*.¹⁵

Whatever success might attend the effort to define such a relation, moreover, coherentism might still survive, if it is committed mainly to the following: that beliefs are not justified one at a time but in clusters, such that no member of the cluster could be justified in isolation, absent enough other members of the cluster. So for a given target belief in the cluster it would be true that its justification would depend on that of the others in the following straightforward sense: were enough of the others in the cluster to lack justification, the target belief could not possibly be justified either. So the coherent, and mutually dependent presence of the clustering beliefs in the mind of that believer is thus essential to the justification of members of the cluster. Supportive experience would not by itself adequately justify a target belief, but must join together with enough supportive beliefs in the cluster.

Accordingly, coherentism seems defensible from the objections so far considered. Of all the many objections pressed against coherentism, there is one, it is true, that is drawn from the core of BonJour's

15 There is also, I believe, a compelling line of reasoning that we enjoy a kind of reflective knowledge, and a kind of reflective justification, each on a higher plane than its animal counterpart, such that we cannot have such more exalted epistemic states absent *either* relevant experience *or* relevant supporting beliefs constitutive of an epistemic perspective. I cannot stop here to detail my reasoning, which the interested reader can find in two papers: "How to resolve the Pyrrhonian problematic: a lesson from Descartes," *Philosophical Studies*, 85 (1997), pp. 229–49, and "Reflective knowledge in the best circles," *Journal of Philosophy*, 94 (1997).

position and would apply not only to pure coherentism but also to more moderate versions that make room for some essential contributing factor other than coherence, such as experience. That further objection targets the requirement that a belief must fit coherently in a complex body of one's own beliefs, even when this is offered as only *one* factor that bears on whether that belief is justified. Against such a proposed factor internalists like BonJour will press this objection: that so complex a fact is unlikely to be accessible to one's armchair reflection. Presumably one would have to unearth the relevant beliefs in that complex body of beliefs, for one thing, so as to then carry out the assessment of coherence, which might require comparing potential rivals, which themselves would have to be thought of and held in mind. This all threatens to prove a task beyond the powers of ordinary believers.¹⁶

We shall return in due course to that objection, based on a requirement that justifying factors be internally accessible. But first let us consider a second crucial objection against coherentism, both pure and moderate, pressed by BonJour as follows:

Meta-beliefs specifying the contents of a person's various systems of belief are themselves obviously contingent and presumably empirical in character, and so must, according to a coherentist view, themselves be justified by appeal to coherence with some appropriate system of beliefs, presumably the system reflecting the person's putative introspective awareness of his own states of mind (for what other system of beliefs would be relevant?). Yet any such account of the justification of these meta-beliefs seems to be inevitably circular or question-begging. . . . In other words, if what is at issue is which beliefs I in fact have, then no appeal to coherence with any system of my beliefs can hope to deal with that issue in a non-question-begging way. Though most coherentists have (somewhat surprisingly) failed to notice this problem or at least to explicitly acknowledge it, it is still pretty obviously one that a viable coherentist position would somehow have to deal with.¹⁷

BonJour grants that once the radical holism of pure coherentism is abandoned, the formulation of this problem becomes more complicated, but not in any way that affects its ultimate seriousness, or so he believes; and he asks us to compare the discussion in the text with a passage in his book, *The Structure of Empirical Knowledge*, pp. 101–2, which culminates with the following:

16 See part iv of section 3.1, "The Main Ingredients of Coherentism."

17 Ibid.

The shift to holism is of no help here, since the very possibility of a nonexternalist holism depends on my having a cognitive grasp of my total system of beliefs and its coherence which is *prior* to the justification of the particular beliefs in the system.

Again in the new work the charge of vicious circularity is repeatedly brought against coherentism, as for instance in the following passage:

The upshot is that there is no non-circular way for a coherentist to appeal to sustained or long-run coherence, making it even more difficult – or, I think, impossible – to respond to the alternative coherent systems objection or to argue for the connection between coherence and truth.¹⁸

In these passages Bonjour regards the coherentist principle as one that is *used* by the thinker in order to gain justification. Briefly put, the principle, call it *C*, says that a belief is justified if coherent. Bonjour seems to suppose that principle *C* would help one attain justification for a given belief *B* by one's in effect *reasoning from it*. One would need to grasp the fact that belief *B* is (or would be) coherent, and must then accept belief *B* for that reason. Thus would the belief become justified, as the result of some such *application* of principle *C*. And to this extent Bonjour's thought is in line with Chisholm's and Descartes's. Chisholm indeed claims explicitly that his epistemic principles would need to be *applied* in the armchair. However, Bonjour now adds a further requirement, one that is, I believe, certainly *absent* from Descartes, and not clearly present in Chisholm: namely, that in such application one's justified belief in principle *C*, and in the truth of its instantiation for belief *B*, would have to enjoy *prior* justification relative to the justification of belief *B* itself. It is this requirement of priority that might make the circularity vicious. But there is no evident reason why we must impose it: the imposition seems arbitrary.

Even once we restrict our focus to *moderate* coherentism, a problem for coherentism pressed by Bonjour apparently stands, one that rests on an internalist requirement of special accessibility. The problem is that we seem to lack armchair access to the coherence of our own beliefs, and that, even if we were to enjoy such access through a period of deep and extensive reflection, the access would seem to be gained

18 Chapter 3, penultimate paragraph.

only via the coherence attained for the belief that one's set of beliefs was indeed sufficiently coherent. Moreover, *this* belief would need to join one's further belief that such coherence was sufficiently correlated with truth. But this would seem to enmesh one in a kind of circularity that BonJour repeatedly condemns as vicious. Accordingly, let us next turn to his own favored alternative: to his proposed account of the sort of epistemic justification that can attach to our ordinary beliefs, an account that is meant to improve on that of the coherentist through its ability to meet internal access requirements. We turn thus to his defense of a kind of *classical, internalist foundationalism*.

11.5 General Defense of Foundationalism

According to this, in order to gain justification for one's commonsense beliefs about the external world and about other minds, one must reason in highly sophisticated and complex ways *from* data about our given conscious experience *to* such commonsense beliefs.

Three main objections arise immediately. First, it seems implausible that ordinary mortals reason in such ways while working or playing or driving or shopping or dining or in any other ordinary situation. Second, there is no a priori way to reason validly from the given to the external. Accordingly, such foundationalism would force us into a radical skepticism that confines us to knowing our present conscious experience and little more. Third, there is the dreaded Sellarsian dilemma: either the foundational conscious states have propositional content, in which case they would seem to require justification in turn, and could not after all function as a foundation; or else they have no propositional content of their own, in which case it is hard to see how they could possibly provide epistemic justification for any belief founded upon them.

At a certain level of abstraction, BonJour's response to the first two problems is entirely Cartesian. Descartes, too, faced the objection that he was raising the bar unrealistically, absurdly requiring mathematicians to engage in rational theology before they could hope to know their mathematics. After all, only near the end of the *Meditations* does the thinker attain true knowledge, aided by the rational theology of the earlier meditations. Only by reasoning deductively from one's clear and distinct certainties can one attain true knowledge of that which is not already clear and distinct on its own. But even this does not suffice to explain Descartes's downgrade of the reliable unreflective thinker; after

all, a good mathematician may be supposed to reason precisely thus with no need of meditation about theology. What the mathematician still misses, absent rational theology, if we believe Descartes, is a proper perspective on his own cognitive doings, those that so securely deliver his theorems. So here is Descartes's response to the incredulity that meets his requirement of a theology-based perspective for knowledge of mathematics.

The fact that an atheist can be "clearly aware that the three angles of a triangle are equal to two right angles" is something I do not dispute. But I maintain that this awareness of his [*cognitionem*] is not true knowledge [*scientia*], since no act of awareness that can be rendered doubtful seems fit to be called knowledge [*scientia*]. Now since we are supposing that this individual is an atheist, he cannot be certain that he is not being deceived on matters which seem to him to be very evident (as I fully explained). And although this doubt may not occur to him, it can still crop up if someone else raises the point or if he looks into the matter himself. So he will never be free of this doubt until he acknowledges that God exists.¹⁹

At a certain level of abstraction, then, Descartes's response to the relevant incredulity is analogous to BonJour's response to a parallel incredulity. Both thinkers respond by granting that, even absent the reflection that they require, one can still gain some important epistemic status, be it *cognitio* or externally reliabilist justification. As a matter of fact, a good case can be made that *cognitio* just is, for Descartes, reliabilist justification, *infallibly* reliabilist justification. Such justification would seem to reside largely in cognitive faculties whose virtue derives essentially from their great reliability in getting us to the truth. And it is in at least one sense an externalist sort of justification, since it can exist even in the mathematician who, far from having a correct perspective on his relevant faculties and on how and why they are so reliable, has on the contrary a woefully mistaken (atheist) perspective.

There remains, however, a crucial difference between the two thinkers. Descartes presses his perspectival concerns all the way to his deepest, most favored faculties of the understanding, namely intuition, whereby we grasp what we then clearly and distinctly enough perceive, and deduction, which depends essentially on intuition. Even when he

19 This passage is from the Second Set of Replies as it appears in *The Philosophical Writings of Descartes*, ed. J. Cottingham, R. Stoothoff, and D. Murdoch, vol. II, p. 101. Where this translation says that an atheist can be "clearly aware," Descartes's Latin is *clare cognoscere*.

perceives something directly and with sufficient clarity and distinctness, Descartes believes that this will *not* suffice to give him real certainty unless he has, at that very moment, an epistemically adequate perspective on the operation and worth of such perception. BonJour, by contrast, takes a very different view of the epistemology of his foundations. About his foundational taking of the given, his beliefs directly descriptive of his current states of consciousness, these are emphatically exempt from the requirement that the subject know them to be reliably derived. BonJour does not suppose that the ostensible taking of the given is infallibly reliable. On the contrary, he allows that one can make mistakes in one's directly descriptive beliefs about one's own current states of consciousness. Nevertheless, any such belief automatically enjoys rational justification, which might of course be defeated by contrary factors yielding reasons to question that belief, reasons that a rational subject should heed.²⁰

We shall return to this disagreement between our two classical foundationalists, but first we turn to the other initial objection to the doctrine.

According to this second objection, there is no way to reason a priori from information about our own subjective states to the layout of any external world. This seems again to be a problem confronting both foundationalists, whose responses are again parallel, at least at a certain level of abstraction. Each rolls up his sleeves in search of the a priori reasoning that will refute the objector. But here too important differences arise, ones that again involve epistemic perspective. Descartes's effort takes place crucially at a remove, even if he does shuttle back and forth as he develops his rational theology and then puts it to epistemic use. He mainly tries to show how and why, if we are careful in our intellectual procedure and keep our assent within proper bounds, we can be sure of succeeding. His objective is thus to

20 Thus here is a passage from the concluding paragraph of section 4.3, "The Justification of Beliefs about Sensory Experience" (on the subject of apperceptive beliefs about the character of one's concurrent conscious states, which apperceptive beliefs constitute BonJour's epistemic foundations, wherein one takes the given):

Here again there is no reason to think that mistake is impossible and thus no reason to think that such an apperceptive belief is infallible or indubitable. But as long as there is no special reason for thinking that a mistake is likely to have occurred, the fact that such a belief seems via direct comparison to accurately characterize the conscious experience that it purports to describe apparently provides an entirely adequate basis for thinking that the description is correct and hence an adequate basis for justification.

establish the assured reliability of a certain sort of intellectual procedure. Note the perspectival character of the objective. We wish to establish the correctness of a certain optimistic perspective about our own cognitive doings. Only when we combine the positive results of this project with the claim that we do form our beliefs with due care and diligence, can we reason our way, through a rather unexpected circuit (as Hume was to quip), back to justified belief in our first-level world view. It remains to be seen what Descartes is really up to in such an unusual project, and how he could imagine that it might be carried out without vicious circularity.

BonJour's procedure is quite different. Avoiding any unexpected circuit, departing from beliefs about one's subjectivity, it moves straight to conclusions about the objective external world through first-level inductive or analogical reasoning. From information about one's own current states of consciousness, pure reason is said to yield contingent hypotheses about the nature and structure of the physical world around us.²¹

So that is how BonJour will try to rebut those two initial objections. Recall first the objection that it is unrealistic to suppose that ordinary people do all the reasoning required if they are to move from the foundations to justified beliefs about the external world. In response BonJour grants that indeed it is unrealistic to suppose that any such sophisticated reasoning must underlie our ordinary perceptual accomplishments. In his view, what people normally attain is perhaps some epistemically valuable state of reliably formed belief, one that ordinarily lacks the support of any such reasoning. But the internalist epistemologist will not settle for such belief. He aims rather for the traditional objective of belief that is not only reliably formed, but also reflectively, rationally justified, where *of course* the subject must have reflective awareness of the reasons that render his belief thus justified.

As for the objection that the requirement of such reflective, rational justification is unfulfillable for us humans, given the independence of objective reality from our subjective states of consciousness, that this obvious fact precludes any valid reasoning from such subjective states to the objective world around us: BonJour's response is to try to *produce* the a priori reasoning that will show not only how it *can* be done, but also that it *has* been done.

So much for the first two of the initial objections to classical foundationalism. Although he does take these seriously, in fact BonJour is

21 Chapter 5 sketches a way in which it might do so.

most concerned with the third objection, the Sellarsian dilemma, which may be put as follows. A proposed foundational state *F* either affirms some propositional content or it does not. If it does, it would seem to require justification of its own. If it does not, then it is hard to see how it could *provide* justification for anything.²²

Perhaps it is not immediately obvious why this should seem so problematic, so here is some support for the apparent dilemma. Recall first the dispute between epistemologically indirect versus direct realism. What justifies perceptual beliefs about externalia? Experiences can do so directly, says the direct realist, and not just via beliefs about them. Against this, the indirect realist (e.g., Moore and Bonjour) argue that only via our *beliefs* about them can our experiences provide justification for corresponding perceptual beliefs. Who is right?

In favor of direct realism it might be pointed out that experience can of course act subliminally to guide our conduct, as when we drive to our destination without paying much conscious attention. But this does not entail that it acts directly without mediation by beliefs. For beliefs are themselves often enough subliminally implicit or even unconscious. So why not say that in perception we do rely on beliefs about our relevant sensory experiences, but that such beliefs remain mostly subconscious?²³

Epistemological theorizing may bring to conscious reflection certain epistemic norms that we apply unreflectively and can apply reflectively. In either case we thus attain a kind of justification through the *application* of norms, even if it is not a conscious application. And this means that the antecedents of the norms must be *believed* to obtain.²⁴

22 See the second paragraph of chapter 4.

23 This is not Bonjour's own response. He prefers to grant that ordinary perceptual beliefs are *not* based on beliefs about one's own sense experience, and takes this to be further evidence that ordinarily we are *not* justified in our perceptual beliefs. In his view, only after we heed and adopt the philosophical reasoning of the classical foundationalist can we hope to attain such justification. (See his reply "Toward a defense of empirical foundationalism" to John Pollock, in M. DePaul (ed.), *Resurrecting Old-Fashioned Foundationalism* (Totowa NJ: Rowman and Littlefield, 2001).

24 Thus if we combine the first section of chapter 1 of Chisholm's *Theory of Knowledge*, 3rd edn, esp. page 1, with the second section of chapter 8, esp. page 76, we find the following view plausibly suggested. In order to be justified in a belief through reasoning one must so believe in the light of the balance of reasons that favors one's then so believing. One must believe in the light of and *for* those reasons, if one's belief is to be fully *rationally* justified. Accordingly, one must have some awareness (at least implicit) of those reasons, and of their justifying force, since for one's belief to be thus fully rational, one must believe guided by such awareness. If one happens to believe in a way that is in fact the best justified in one's circumstances, but one does so by luck,

What about the principles that govern the appropriate formation of beliefs about experiences *themselves*? Must these *also* operate *via* prior beliefs about those experiences? Hardly. How could they? These are *already* the beliefs about the experiences.

At this juncture some would say that it is the experience itself that causes and justifies the relevant belief about it. For Bonjour, however, it is not the experience itself but a special “awareness” of it. Who is right? Our next section considers Bonjour’s side on this.

11.6 Bonjour’s Foundationalism and its Rationale

Here is the key claim in Bonjour’s foundationalism:

[Where] . . . I have a conscious sensory experience, I am, as already argued, aware of its specific sensory content simply by virtue of being in that experiential state. And therefore if (i) an apperceptive belief that I entertain purports to describe or conceptually characterize that perceptual content, albeit no doubt incompletely, and (ii) I understand the descriptive content of that belief, i.e., understand what an experience would have to be like in order to satisfy that conceptual description, then I seem to be in a good, indeed an ideal, position to judge whether the conceptual description is accurate as far as it goes; and if it is, to be thereby justified in accepting the belief. Here again there is no reason to think that mistake is impossible and thus no reason to think that such an apperceptive belief is itself infallible. But as long as there is no special reason for suspecting that a mistake has occurred, the fact that such a belief seems to accurately characterize the conscious experience that it

while insufficiently sensitive to their presence and justifying force, then one’s belief falls short of proper *rational* justification. Internalist epistemology focuses on such accessible reasons that are or can be operative in making us rationally justified, and these must be reasons of which we can be relevantly aware so that they can guide our belief formation as required. Often enough we are thus guided, however, without full explicit awareness of the relevant reasons and of the principles that set out their rational force. Fortunately, philosophical reflection can help us bring such principles to consciousness, which can in turn help us to attain better awareness of their content and enhanced sensitivity to their force, which might well enable us to improve our epistemic situation by replacing ill justified beliefs by better justified epistemic attitudes of alternative belief or of withholding. (I will not pretend, however, that this line of reasoning is fully explicit in Chisholm’s approach, nor even that it is *obviously* implicit. I suggest only that it is a plausible view suggested by his words, and in any case one that helps support an interesting kind of internalism in epistemology.)

purports to describe provides, I suggest, an entirely adequate basis for thinking that the description is correct, and thus an adequate basis for justification.²⁵

So here is a way to put Bonjour's account of a sufficient condition for foundational justification for a belief about a concurrent sensory experience of one's own:

Bonjour's Foundationalist Principle (BFP):

One is foundationally justified at *t* in consciously believing ⟨*p*⟩ IF, at *t*, four conditions are satisfied:

- (a) One believes ⟨*p*⟩, which means of course that one understands the descriptive content of that belief;
- (b) One undergoes sensory experience *E* with perceptual content *C*;
- (c) In consciously believing ⟨*p*⟩ one is describing one's experience *E* as one with perceptual content *C*; and
- (d) There is no special reason for suspecting that a mistake has occurred.²⁶

Let us focus on condition (d). What exactly does it say? Presumably what it means can be put more fully like this:

- (d1) One has no special reason for suspecting that one has made a mistake in believing ⟨*p*⟩.

A reason that is available but which no one has or which someone else has, but one lacks, presumably would not be relevant. A relevant reason would be a reason that one has in one's own possession at the time in question. To *suspect*, moreover, seems a propositional attitude similar to believing but somehow weaker. It is perhaps to be *inclined to believe* or the like. And the kind of reason that figures in (d1) is presumably epistemic reason (not pragmatic or moral or any other practi-

25 L. Bonjour, "Foundationalism and the external world," *Philosophical Perspectives*, 13: *Epistemology* 1999, p. 235. Similar statements are found in his "The dialectic of foundationalism and coherentism," in Greco and Sosa, *Blackwell Guide to Epistemology*, p. 134, and in his contribution to our current debate, above.

26 See "Conscious Thought and Constitutive Awareness of Content," section 4.1, especially its concluding three paragraphs.

cal reason). So condition (d1) seems to amount to a kind of “negative coherence” requirement: that is, one cannot at the time in question have within one’s total mental state any elements that would strongly enough support the view that in believing ⟨p⟩ one is making a mistake.

However that condition is to be understood more fully, two sorts of problems arise. In the first place, the sufficient condition specified in BFP is not of the sort that will help explain the supervenience of epistemic normativity in the way one would hope a foundationalist condition would do, since it will itself contain a claim of epistemic normativity, albeit a negative claim. If the only kind of foundational source of normative epistemic status we are able to specify itself contains an essential invocation of epistemic normativity (positive or negative), then we would seem *not* to be on the way to understanding just how epistemic normativity enters our cognitive doings, our beliefs, and other cognitive propositional attitudes, how it enters the sphere of our thought out of *non*-normative sources. This would be at most a sin of omission, however, and not of commission.

In any case, there is, secondly, the problem of the speckled hen.

Let us consider the visual sense-datum which is yielded by a single glance at a speckled hen. The datum may be said to “comprise” many speckles . . . [How] many speckles does the datum comprise? . . . If we judge that there are forty-eight, it would seem, at first consideration at least, that we might very well be mistaken . . . [Our] difficulty is not that there must be characteristics of the many-speckled datum which pass unnoticed; it is, more seriously, the fact that we are unable to make a reliable judgment about what we *do* notice.²⁷

Suppose one did make the judgment that the visual sense-datum comprised 48 speckles. Might one not satisfy the conditions specified by BonJour’s BFP, above, for foundationally justified belief? Consider the four clauses of BFP as applied to the present case:

One is foundationally justified at *t* in consciously believing ⟨This visual sense-datum comprises 48 speckles⟩ IF, at *t*, four conditions are satisfied:

- (a) One believes ⟨This visual sense-datum comprises 48 speckles⟩, which means of course that one understands the descriptive content of that belief;

27 R. M. Chisholm, “The problem of the speckled hen,” *Mind*, 51 (1942), pp. 368–73.

- (b) One undergoes a visual sensory experience *E* with the perceptual content *C* of a sense-datum with 48 speckles (since one is staring at the side of that hen in mid-day sunlight at arm's length and that is how many quite separate and distinct black speckles it has);
- (c) In consciously believing ⟨This visual sense-datum comprises 48 speckles⟩ one is describing one's experience *E* as one with the perceptual content *C* of a sense-datum with 48 speckles; and
- (d) There is no special reason for suspecting that a mistake has occurred.

Only through the potential escape hatch of condition (d) can BonJour's account BFP avoid granting foundational justification to one's belief about the number of speckles. But why is it so bad to grant foundational justification in this case? Answer: what makes it so bad is that we are after an account of *epistemic* justification, and not just an account of deontological justification. We are not just after a kind of justification that might exempt one from blame, that would perhaps make one excusable and not subject to any moral or quasi-moral criticism or assignment of discredit or blame. What we are after is rather the kind of positive epistemic status that makes one an epistemically reasonable believer, in a way that makes one's belief epistemically creditable to oneself as a cognitive agent. This is presumably a distinctively epistemic sort of accomplishment with its own sources, one distinct from any excusability or blamelessness that might derive from one's being brainwashed, or influenced by subliminal suggestion, or through a bad intellectual upbringing, or through innocent enough inattention, etc. Whether one does or does not consider the sort of epistemic justification at issue as of a sort relevant to propositional knowledge, clearly the kind of excusability-based justification that might derive from brainwashing is not of traditional epistemological concern, nor can it be the sort of epistemic rational state that we seek through inquiry into the rational status of our beliefs about the external world. None of the listed excusing factors would seem to be positively relevant to the epistemic worth of one's belief, even if they are highly relevant to one's excusability for believing as one does.

If that is right, then the escape hatch of (d) seems crucial to the defensibility of BFP. Keeping (d) has the sort of problem indicated earlier, of being at odds with the sort of explanation of the superven-

nience of epistemic normativity that in the end we would like to have. Even putting that problem aside, we face the following further problem: namely, that BFP, even with (d), seems satisfiable in a case where one clearly does *not* attain epistemic justification for one's belief about the number of speckles. After all, one might have no reason to suspect that one is an abysmal judge of the number of speckles. One might never have checked one's reliability on that matter or on any other matter believed to have any bearing on how reliable one is about this many speckles. If, nevertheless, through whatever mechanism of wishful thinking or whatnot, one finds oneself believing that the number of speckles is 48, and that is in fact the right number of speckles, this belief obviously gains no status as *epistemically* justified. And the fact that it is a true belief bears negligibly or not at all. Even if one describes the visual image as comprising that many speckles, and the visual image does comprise that many speckles, and one *possesses* no reason to suspect that one is making a mistake, this does not suffice to give one (foundational) epistemic justification for one's belief (about the number of speckles).

It might be objected that simply the fact that one is so highly unreliable about that many speckles already constitutes a reason for suspecting that a mistake has occurred, even if one does not take it in. So there *is* after all a reason for one to suspect that a mistake has occurred, and this blocks one from qualifying as foundationally justified, according to the conditions of BFP. But now the opposition to externalism is hard to understand. It is now hard to see just how the resulting position is so importantly different from externalist reliabilism. After all, it will now be the mere external fact that one is unreliable in one's belief formation about the number of speckles that makes the difference, irrespective of whether, as epistemic agent, one takes it into account. So what makes one justified in judging that it is three speckles would then seem to be at least in important part the fact that one avoids in such belief-formation the kind of unreliability that dooms one's belief about the 48 speckles. But this now is hard to distinguish from some sort of externalist reliabilism.

If that is a real problem, is there a way for Bonjour to overcome it? For example, is there some simple modification to his BFP that will work? Perhaps one could restrict the sort of content allowed, so as to preclude the speckled hen and other such cases. What sort of restriction might one try? One thing seems clear. If one is to avoid externalist reliabilism, the further condition cannot be of the sort: *formed*

reliably whether within or outside the purview of the epistemic agent. For this will be a capitulation to externalist reliabilism.²⁸

11.7 A Comparison: Externalism, Coherentism, Foundationalism

It will be useful, in these concluding pages, to compare Bonjour's critique of his rejected options, externalism and coherentism, with his defense of his favored option, foundationalism.

His main objection against externalism, recall, apart from its supposed intrinsic implausibility, goes back to his influential early paper that already lays out the objection from "clairvoyance out of the blue." Claire is suddenly blessed with clairvoyance, a reliable mode of belief formation. Is she thereby blessed with a source of justified beliefs, as, apparently, some externalists would have us believe? Surely not, not if she lacks an adequate perspective on the operation of her wonderful faculty. Specifically, it is argued, Claire needs not only to be free of a contrary perspective that impugns the reliability of her clairvoyance. She needs, in addition, an underwriting perspective that positively favors the reliability of her newly acquired faculty. Absent such underwriting, she will lack adequately rational epistemic justification for accepting the deliverances of her faculty, which will present themselves as beliefs out of the blue, arbitrary and groundless.

There is quite a contrast, however, between this assessment of Claire's clairvoyance, and Bonjour's attitude towards *the taking of the given*. Here again is a mode of belief formation, whose operation can be more or less reliable depending on the circumstances: on the kind of subject matter, presumably, and on the subject's attentiveness, sobri-

28 And there is in addition the following problem. What if one forms a *false* conscious belief about a present conscious state of one's own? In this case one presumably still has a foundationally justified belief, so long as one is not privy to the factor that spoils one's belief, or to the way in which it does its damage, or so the theory would seem to entail, since the erring subject may have no reason in his possession to suspect that his introspection is then going astray. But is one's conscious state then being *described*? Can there be *false* Bonjourian "description"? Is this allowed as the relevant relation of description postulated to account for foundational justification? Or must a description be true in order to count as a relevant description at all? Moreover, is the relevant sort of description one that derives simply from its truth, or must there be some sort of causal or tracking relation between the description and the described in order for a proper relation of description to relate the two? These are all issues that it would be interesting to explore in a fuller development of the theory.

ety, and fineness of discrimination about that subject matter, and so on. Reliability seems from the context and overall framework to be required, at least presumptive reliability, as is also mildly suggested by the comment that absent special reasons why one's faculty might lead us into *error* in the particular circumstances of its operation, it may be accepted as delivering beliefs that are rationally justified.

We may of course have much better reason to think that we are good introspectors (or "good takers of the given") than to think that we are good clairvoyants, even when our faculties happen to be equally highly reliable in fact. That seems beyond dispute. What remains still unclear is the rationale for the *epistemic* discrimination between clairvoyance and introspection. Why not distinguish between them just by saying that we do enjoy an endorsing epistemic perspective for the one, while lacking any for the other? Why distinguish between them rather, as BonJour does, by saying that although both are fallible, the deliverances of one come with built-in *prima facie* (though defeasible) justification, while the deliverances of the other do not enjoy such a status of privilege by default? What justifies such discrimination, what defends it from a charge of arbitrary dogmatism?

BonJour's response here is of the familiar, Kantian, "this or nothing" variety. We had better accept that introspection does deliver the epistemic goods unaided by any underwriting perspective, or we shall be plunged into the deepest skepticism, where nothing at all is justified, let alone our rich set of beliefs about a world around us.²⁹ But this is not convincing. For example, so long as we are going to allow the operation of faculties that are not underwritten by the subject's perspective, why not allow both introspection *and* perception? That too would avoid the dark pit of skepticism, and it would save us the search, perhaps quixotic, for valid reasoning from the subjective to the objective. And if we allow to memory, and to temporally extended valid reasoning, also their own measure of epistemic efficacy, how do we deny it to the likes of clairvoyance? Would the beliefs of the clairvoyant now be granted their own automatic *prima facie* justification, along with the justification granted the deliverances of memory and inference? And if not, then how do we discriminate non-arbitrarily among such faculties, all of which may be supposed reliable for the sake of argument?

29 See the concluding sentence of "Conscious Thought and Constitutive Awareness of Content," section 4.1, where we are told that to deny that BonJour's taking of the given "is ever possible is to guarantee vicious regresses in all directions, rendering the operation of the intellect inherently futile."

11.8 A Cartesian Alternative

Moreover, why should we believe that unless we allow our deepest faculties to yield justification unaided by any underwriting perspective, we shall be forced into the deepest skepticism? On this matter Descartes himself took a crucially different view. It was precisely his deepest faculties that he had firmly in focus when he saw the need for an underwriting perspective. And this perspective is precisely what he attempted to provide, in his *Meditations*. So why would Bonjour think that this is hopeless? He does not say in full, but we can piece together a reasonable hypothesis from the ways in which he objects to various epistemological positions. Evidently he believes that there is no way for such a Cartesian project to avoid vicious circularity or regress.³⁰ The Pyrrhonian problematic admits only three options: circle, regress, and foundations. Since the circle and the regress are inevitably vicious,³¹ we must make the best of the foundations option, and that is what he, Bonjour, is trying to do. This means that we cannot require an underwriting perspective for the justification of all our first-level beliefs. Once we get to a foundation, *by definition* we must not still require some further beliefs, those constitutive of the underwriting perspective, as further sources of justification. If we did so, then our so-called foundation would not be foundational after all, but would depend on these further beliefs for its justification.

Such reasoning is based on important confusions. One of them goes back to the supposition that foundationalism and coherentism as originally defined are dichotomous. This is just a mistake.³² *Moderate* coherentism after all requires only that there be a non-belief factor, e.g. experience, that provides positive epistemic status on its own, without

30 Thus in the seventh paragraph of section 1.4, “The Case against Traditional Foundationalism,” we find this: “A basic belief cannot be literally *self*-justifying unless the foundationalist accepts circular reasoning as a source of justification, a view that seems obviously wrong (and that would also undercut one of the main objections to coherentism).” And this is only one of several passages with anticircularity animus.

31 For interesting dissent, see Peter Klein’s “Human knowledge and the infinite regress of reasons,” in James Tomberlin (ed.), *Philosophical Perspectives*, 13: *Epistemology* (1999), pp. 297–327.

32 Here’s how the early paragraphs of chapter 1 introduce “the dichotomy between *foundationalist* and *coherentist* accounts of epistemic justification: does such justification derive ultimately from ‘foundational’ beliefs whose justification somehow does not depend at all on that of other beliefs, or does it derive instead from relations of coherence or agreement or mutual support among beliefs, with no appeal to anything outside the system of beliefs?” Failing to exhaust logical space, this is no dichotomy.

having to operate via the filter of belief. Such moderate coherentism thus allows the possibility of a further source of justification that satisfies both of the following conditions: first, it is belief-involving; and, second, it might join together with the non-belief foundational source in raising the epistemic status of one's most basic beliefs.³³

But how could one possibly acquire justification for the beliefs in the underwriting perspective, especially when the faculties at work are those at the deepest level? What is the point of acquiring a perspective that underwrites one's basic epistemic faculties (intuition and deduction, say, on a Cartesian perspective) if one *must* acquire that perspective by the use of those very faculties? Is not such bootstrapping bound to be viciously circular?³⁴ It must be because he takes the answer here to be yes that BonJour is so willing to exempt his taking of the given

33 This is indeed the sort of view that I have myself long advocated. Consider for example my paper, "How do you know?" (*American Philosophical Quarterly*, 11 (1974); included in my *Knowledge in Perspective: Selected Essays in Epistemology* (Cambridge: Cambridge University Press, 1991)). Its concluding pages contain a passage where it is emphasized that the theory defended in that paper "will not commit one to a picture of knowledge according to which there is a bedrock of self-evident propositions. It is perfectly consistent with the present theory that part of what makes *any* proposition evident be its coherence with a network of mutually supporting propositions." In "Nature unmirrored, epistemology naturalized," *Synthese*, 55 (1983) (also in *Knowledge in Perspective* Cambridge: Cambridge University Press, 1991), this is developed further as follows, for example in the third paragraph of section G: "[There is] . . . an immense variety of animal knowledge, instinctive or learned, which facilitates survival and flourishing in an astonishingly rich diversity of modes and environments. Human knowledge is on a higher plane of sophistication, however, precisely because of its enhanced coherence and comprehensiveness and its capacity to satisfy self-reflective curiosity. Pure reliabilism is questionable as an adequate epistemology for such knowledge. The challenge of doxastic ascent does seem pertinent at least here, and it signals the promise of comprehensive coherence as a source of epistemic authority." This sort of view, labeled "foundherentism," has also been defended by Susan Haack, in her *Evidence and Inquiry* (Oxford: Blackwell, 1993).

34 Much recent work in epistemology has focused on this issue and related issues of "begging the question," partly in connection with the issue of content externalism and self-knowledge, but also in its own right as a general issue in epistemology. See, for just a few examples: Jonathan Vogel, "Reliabilism leveled," *Journal of Philosophy*, 97 (2000), pp. 602–23; Stewart Cohen, "Basic knowledge and the problem of easy knowledge," forthcoming in *Philosophy and Phenomenological Research*; Brian McLaughlin, "Self-knowledge, externalism, and skepticism," *Aristotelian Society Supplementary Volume* 74 (2000), pp. 93–117; Martin Davies, "Externalism and armchair knowledge," in Paul Boghossian and Christopher Peacocke (eds.), *New Essays on the A Priori* (Oxford: Oxford University Press, 2000), pp. 384–414; and Crispin Wright, "(Anti-)sceptics simple and subtle: Moore and McDowell," forthcoming in *Philosophy and Phenomenological Research*, and "Wittgensteinian certainties," forthcoming in Denis McManus (ed.), *Wittgenstein and Scepticism*, (London: Routledge).

from the requirement of underwriting perspective. The reason why in his view we must here make a stand is that if we do not, then we will be forced into reasoning that must be circular (or regressive), viciously so.

For my part, I do not share that dark view of circular reasoning as thus inevitably vicious. In particular, the critics of Descartes on this score seem to me imperceptive, and Descartes perfectly correct to reason as he does. Here I do not mean to endorse the substance of his view, including the detailed rational theology. What I mean to endorse is the structure of his epistemology, with its recognition of the importance of reflective endorsement, and of the fact that this is not to be attained without eventual circularity, and the acceptance of the value of the outcome despite that fact, and of how this redounds to the epistemic benefit of the retail beliefs that compose the thinker's view of the world within and around him.³⁵

11.9 Further Problems for Internalist Foundationalism

Here, finally, we put aside, for the sake of argument, the difficulties just canvassed concerning the proposed foundational taking of the given, and the special status it is given as source of justification in no need of reflective endorsement or perspectival underwriting. Let us grant that taking the given can, at least in certain circumstances, function thus foundationally, as a source of rational justification. Is this enough to yield a viable classical foundationalism, one based on the requirements of reflective access and no-circularity? Well, not quite, since, recall, we still need to reason our way from the foundations, in doing which we shall need subtle and complex reasoning extensively aided by memory. Recall indeed the earlier concession that foundationalism must invoke considerations of coherence in moving beyond the foundations to the external world. And consider the intricacy of the sketched "proof of an external world" in chapter 5.

Let us take stock. Suppose the foundational given is accepted as an independent source of rational justification. And suppose further that there is indeed a valid argument leading from the foundational data to

35 This is an outlook already defended in my papers that are cited in n. 33, and developed, by comparison also with Descartes's epistemology, in a series of more recent papers.

a commonsense picture of the world around us, with its colored and shaped objects arrayed in space and time, and so on. Does that not give us all we need for a viable foundationalism?

It does not. Reflective access, it will be recalled, is a requirement not only for foundational justification, but for justification of whatever sort. Consider then the thinker emerging from a perfected chapter 5 of BonJour's monograph. He has now gone through the reasoning that supposedly gives him the desired justification for his beliefs about an external world. But if that reasoning is to have that effect then it must meet the requirement of reflective access. However, evidently it does *not* meet that requirement in the way in which the foundational takings of the given do so. For the undergoing of such reasoning is a process that may take considerable time, and that in any case is not itself a current state of the subject's consciousness. So how *does* one gain reflective access to the justifying power of such a process? Here one is of course tempted to try "bootstrapping," but that option has been firmly foreclosed by BonJour. In any case, even if one did after all allow it: (1) that would now open the possibility that other external modes of belief formation might also be similarly underwritten, and (2) it would also require re-examining the discriminatory treatment of foundational taking of the given as a source of rational justification with a special status; one would now wonder why such a foundational faculty could not itself gain its status as reflectively justified in the same sort of way as reasoning, perception, memory, etc., namely through appropriate bootstrapping.

In fact, consider even just a step of immediate inference at an instant *t*. Here the subject adduces certain reasons for assenting to a proposition *p*, the new step in the reasoning, which one reaches in that immediate inference. Consider now the reasons thus adduced, the premises and lemmas from which the new conclusion (intermediate or final) is drawn. Of these the subject must be at least implicitly aware, and perhaps explicitly aware. It depends on the level of explicitness of the bit of reasoning involved. However, is it also required that the subject be aware *that* he is drawing that conclusion from those premises or lemmas? Must he at least be aware *of* his drawing of the conclusion, of his basing that new belief on those adduced reasons? Well, what exactly is the question? When it is wondered whether he *must* be aware (whether *that* or *of*), what is the implicit purpose in play? For *what purpose* is the awareness supposed to be requisite? Is it for a proper bit of connected reasoning to take place? Surely not. For the reasoning to have epistemic efficacy? Well, what *sort* of efficacy? Is it not good

that a mathematician reason with validity in arriving at conclusions, so as to ensure that the conclusions are indeed theorems? Surely it is. But must one as mathematician be aware *that* one is so reasoning, or must one at least be aware *of* one's reasoning, in order for it to attain that sort of efficacy? That is far from obvious. Indeed it would seem more plausible that the awareness is required rather for the attainment of *reflective* justification, and, if all goes well, reflective knowledge. It is this that does require perspective and meta-awareness. This seems to me right, and importantly right. Correlatively, developing an internalist epistemology through an account of reflective knowledge seems a project worth pursuing, and one I intend to pursue.

Bibliography

- Alston, W. P., *Epistemic Justification: Essays in the Theory of Knowledge* (Ithaca, NY: Cornell University Press, 1989), esp. Part III: Internalism and Externalism.
- Axtell, Guy (ed.), *Knowledge, Belief, and Character* (Totowa, NJ: Rowman and Littlefield, 2000).
- Ayer, A. J., *Foundations of Empirical Knowledge* (London: Macmillan, 1940).
- Bennett, Jonathan, "Substance, reality, and primary qualities," *American Philosophical Quarterly*, 2 (1965).
- Blanshard, Brand, *The Nature of Thought* (London: Allen and Unwin, 1939).
- BonJour, Laurence, *The Structure of Empirical Knowledge* (Cambridge, Mass.: Harvard University Press, 1985).
- , "Is thought a symbolic process?," *Synthèse*, 89 (1991)
- , "Against naturalized epistemology," *Midwest Studies in Philosophy*, vol. 19 (1995).
- , *In Defense of Pure Reason* (Cambridge: Cambridge University Press, 1998).
- , "Foundationalism and the external world," in James Tomberlin (ed.), *Philosophical Perspectives*, 13: *Epistemology* (1999).
- , "The dialectic of foundationalism and coherentism," in J. Greco and E. Sosa (eds.), *The Blackwell Guide to Epistemology* (Oxford: Blackwell, 1999).
- , "Toward a defense of empirical foundationalism," in M. DePaul (ed.), *Resurrecting Old-Fashioned Foundationalism* (Lanham, Md.: Rowman and Littlefield, 2001).
- , "The indispensability of internalism," *Philosophical Topics*, forthcoming.
- Bosanquet, Bernard, *Implication and Linear Inference* (London: Macmillan, 1920).
- Bradley, F. H., *Essays on Truth and Reality* (Oxford: Oxford University Press, 1914).
- Broad, C. D., *Scientific Thought* (London: Routledge and Kegan Paul, 1923), Part II.

- Butchvarov, Panayot, *The Concept of Knowledge* (Evanston, Ill.: Northwestern University Press, 1970).
- Chisholm, Roderick M., "The problem of the speckled hen," *Mind*, 51 (1942).
 —, *Theory of Knowledge* (Englewood Cliffs, NJ: Prentice-Hall, 1966; 2nd edn. 1977; 3rd edn. 1989).
- Cohen, Stewart, "Justification and truth," *Philosophical Studies*, 46 (1984).
 —, "Basic knowledge and the problem of easy knowledge," *Philosophy and Phenomenological Research*, forthcoming.
- Cohen, Stewart and Lehrer, Keith, "Justification, truth, and coherence," *Synthese*, 55 (1983).
- Davidson, Donald, "A coherence theory of truth and knowledge," in Dieter Henrich (ed.), *Kant oder Hegel* (Stuttgart: Klett-Cotta, 1983).
- Davies, Martin, "Externalism and armchair knowledge," in Paul Boghossian and Christopher Peacocke (eds.), *New Essays on the A Priori* (Oxford: Oxford University Press, 2000).
- DePaul, Michael (ed.), *Resurrecting Old-Fashioned Foundationalism* (Totowa, NJ: Rowman and Littlefield, 2001).
- Evans, Gareth, *The Varieties of Reference* (Oxford: Oxford University Press, 1982).
- Foley, Richard, *The Theory of Epistemic Rationality* (Cambridge, Mass.: Harvard University Press, 1987).
 —, *Working Without a Net* (New York and Oxford: Oxford University Press, 1993).
- Fumerton, Richard, "The internalism/externalism controversy," in James Tomberlin (ed.), *Philosophical Perspectives*, 2: *Epistemology*, 1990.
 —, *Metaepistemology and Skepticism* (Lanham, Md.: Rowman and Littlefield, 1995).
- Goldman, Alvin, "What is justified belief?," in G. Pappas (ed.), *Justification and Knowledge* (Dordrecht: D. Reidel, 1979).
 —, *Epistemology and Cognition* (Cambridge, Mass.: Harvard University Press, 1986).
 —, *Liaisons: Philosophy Meets the Cognitive and Social Sciences* (Cambridge, Mass.: MIT Press, 1992).
 —, "Epistemic folkways and scientific epistemology," in *Liaisons* (Cambridge, Mass.: MIT Press, 1992); reprinted in Hilary Kornblith (ed.), *Naturalized Epistemology*, 2nd edn. (Cambridge, Mass.: MIT Press, 1994).
 —, "Internalism exposed," *Journal of Philosophy*, 96 (1999).
- Greco, John, *Putting Sceptics in their Place* (Cambridge: Cambridge University Press, 2000).
- Haack, Susan, *Evidence and Inquiry: Towards Reconstruction in Epistemology* (Oxford: Blackwell, 1993).
- Harman, Gilbert, *Thought* (Princeton, NJ: Princeton University Press, 1973).
- Kitcher, Philip, "The naturalists return," *Philosophical Review*, 101 (1992).
- Klein, Peter, "Human knowledge and the infinite regress of reasons," in James Tomberlin (ed.), *Philosophical Perspectives*, 13: *Epistemology* (1999).

- Kornblith, Hilary (ed.), *Epistemology: Internalism and Externalism* (Oxford: Blackwell, 2001).
- Kvanvig, Jonathan L. (ed.), *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge* (Lanham, Md.: Rowman and Littlefield, 1996).
- Lehrer, Keith, *Knowledge* (Oxford: Oxford University Press, 1975).
- , *The Theory of Knowledge* (Boulder, Colo.: Westview Press, 1990; 2nd edn., 2000).
- Lewis, C. I., *An Analysis of Knowledge and Valuation* (La Salle, Ill.: Open Court, 1946).
- Mackie, J. L., *Problems from Locke* (Oxford: Oxford University Press, 1976).
- McLaughlin, Brian, "Self-knowledge, externalism, and skepticism," *Aristotelian Society Supplementary Volume* 74 (2000).
- Martin, Michael, "An eye directed outward," in C. Wright, B. C. Smith, and C. MacDonald (eds.), *Knowing Our Own Minds* (Oxford: Oxford University Press, 1998).
- Moser, Paul, *Knowledge and Evidence* (London: Cambridge University Press, 1989).
- Peacocke, Christopher, "Conscious attitudes, attention, and self-knowledge," in C. Wright, B. C. Smith, and C. MacDonald (eds.), *Knowing Our Own Minds* (Oxford: Oxford University Press, 1998).
- Plantinga, Alvin, *Warrant: The Current Debate* (Oxford: Oxford University Press, 1993).
- , *Warrant and Proper Function* (Oxford: Oxford University Press, 1993).
- Popper, Karl, *The Logic of Scientific Discovery* (New York: Harper, 1959).
- Price, H. H., *Perception*, 2nd edn. (London: Methuen, 1950).
- Rescher, Nicholas, *The Coherence Theory of Truth* (Oxford: Oxford University Press, 1973).
- , *Methodological Pragmatism* (New York: New York University Press, 1977).
- Rorty, Richard, *Philosophy and the Mirror of Nature* (Princeton, NJ: Princeton University Press, 1979).
- Rosenthal, David, "Two concepts of consciousness," *Philosophical Studies*, 94 (1986); reprinted in David Rosenthal (ed.), *The Nature of Mind* (New York: Oxford University Press, 1991).
- Russell, Bertrand, *The Problems of Philosophy* (Oxford University Press, 1997); first published 1912.
- Ryle, Gilbert, *The Concept of Mind* (London and New York: Hutchinson, 1949).
- Sartwell, Crispin, "Knowledge is merely true belief," *American Philosophical Quarterly*, 28 (1991).
- , "Why knowledge is merely true belief," *Journal of Philosophy*, 89 (1992).
- Schmitt, Frederick, *Knowledge and Belief* (London: Routledge, 1992).
- Sellars, Wilfrid, *Science, Perception and Reality* (London: Routledge and Kegan Paul, 1963).

- , “Givenness and explanatory coherence,” *Journal of Philosophy*, 70 (1973).
- , “The structure of knowledge,” in Hector-Neri Castañeda (ed.), *Action, Knowledge, and Reality: Critical Studies in Honor of Wilfrid Sellars* (Indianapolis: Bobbs-Merrill, 1975).
- , “The lever of Archimedes,” *The Monist*, 64 (1981).
- Shope, Robert, *The Analysis of Knowing* (Princeton, NJ: Princeton University Press, 1983).
- Sosa, Ernest, “How do you know?,” *American Philosophical Quarterly*, 11 (1974).
- , “Nature unmirrored, epistemology naturalized,” *Synthèse*, 55 (1983).
- , *Knowledge in Perspective* (Cambridge: Cambridge University Press, 1991).
- , “How to resolve the Pyrrhonian problematic: a lesson from Descartes,” *Philosophical Studies*, 85 (1997).
- , “Reflective knowledge in the best circles,” *Journal of Philosophy*, 94 (1997).
- , “Goldman’s reliabilism and virtue epistemology,” *Philosophical Topics*, 29 (1 and 2) (2002).
- Strawson, P. F., “Perception and its objects,” in G. F. Macdonald (ed.), *Perception and Identity* (Ithaca, NY: Cornell University Press, 1979).
- Vogel, Jonathan, “Reliabilism leveled,” *Journal of Philosophy*, 97 (2000).
- Wright, Crispin, “(Anti-)sceptics simple and subtle: Moore and McDowell,” *Philosophy and Phenomenological Research*, forthcoming.
- , “Wittgensteinian certainties,” in Denis McManus (ed.), *Wittgenstein and Scepticism* (London: Routledge, forthcoming).
- Zagzebski, Linda, *Virtues of the Mind* (Cambridge: Cambridge University Press, 1996).

Index

- ability, and knowledge, 101–2
acquaintance, 119–22
 Fumerton’s theory, 129–30,
 193 n.22
 see also direct apprehension
adroitness
 use of term, 157 n.3
 and virtue epistemology, 156–8,
 185–6
adverbial theories, 78–9 n.3, 88 n.18
Alston, William P. (1921–), 206
alternative coherent systems objection,
 43, 53–4
animal knowledge
 cognitive virtues approach, 2
 reliability, 2
 vs. reflective knowledge, 2
aptness
 use of term, 157 n.3
 and virtue epistemology, 156–8,
 184–5, 189, 197
atheism, 213
attention, 126–9
 mechanisms, 123 n.6
awareness, 119–22
 apperceptive *vs.* non-apperceptive,
 62–4
 and consciousness, 119–21
 and experience, 119–20
 foundationalist perspectives, 120–1
 higher-order, 62, 63, 65–8
 non-cognitive, 63
 pre-predicative, 20
 see also constitutive awareness;
 experiential awareness;
 intellectual awareness
Ayer, A. J. (1910–89), 8, 102, 121 n.4
basic beliefs, 12, 13, 14
 definition, 208
 epistemic justification, 15–19, 24–6,
 61–5, 69–74, 190–7
 self-evident, 16–17
 self-justified, 16
belief formation, 26, 107
 modes, 28, 222–3, 227
 principles, 201
 reliable, 206
beliefs
 adroit-justified, 156–8, 161, 185–6
 apt-justified, 156–8, 161, 184–5, 189,
 197
 Cartesian, 141
 cognitively spontaneous, 49–51, 54,
 57
 coherent system of, 42–3, 45–7,
 51–2, 57–8, 111–12, 207
 concept of, 135, 203 n.5
 conceptual, 123
 foundationally justified, 128, 196,
 198, 199–200
 introspective, 126
 and knowledge, 6, 30, 102–4, 113–14
 non-occurrent, 69
 object-level, 170 n.14
 occurrent, 61, 62–3, 69–70

- beliefs (*cont'd*)
 perceptual, 73 n.10, 216
 safe, 138–9
 and sensory experience, 69–74, 107,
 217–18
 true, 102, 104, 109–10, 159, 166,
 182
 types of, 113
 use of term, 9 n.6
 well-justified, 141
see also basic beliefs; meta-beliefs
- Bennett, Jonathan, 94 n.23
- Berkeley, George (1685–1753), 85–6,
 93, 94
- Berlioz, Louis Hector (1803–69), 100
- BFP (BonJour's Foundationalist
 Principle), 218–22
- Blanshard, Brand (1892–1987), 42, 49
- BonJour, Laurence, 161 n.38
 epistemology theory, 1–2
 foundationalism, 203 n.3
- BonJour's Foundationalist Principle
 (BFP), 218–22
- bootstrapping, 227
- Bosanquet, Bernard (1848–1923), 42,
 44
- Bradley, Francis Herbert (1846–1924),
 42
- Broad, Charlie Dunbar (1887–1971),
 86
- Butler, Joseph (1692–1752), 41
- Carnap, Rudolf (1891–1970), 117
- Cartesian tracking, 154 n.10
- Chisholm, Roderick (1916–99), 8, 48,
 75, 82, 84, 102
 on degree of justification required for
 knowledge, 22 n.18
 on epistemic justification, 202
 epistemology, 201, 211
 foundationalism, 202–3
 internalism, 145–9, 150, 173–5,
 188 n.16, 216–17 n.24
 speckled hen problem, 121, 191,
 219–21
- Chisholm-justification, 150–1
- circularity
 vicious, 44–5, 58–9, 211, 212,
 224–6
 virtue epistemology, 165–8
- clairvoyance, 28–33, 205, 206, 222, 223
- clarity, reliability, 2
- cogito*, 2
- cognitio*, 213
- cognitive virtues, 169
 and animal knowledge, 2
 issues, 2
- coherence
 comparative assessments, 46
 concept of, 46–8
 epistemic, 113–14, 115, 189–90
 explanatory relations, 47, 111
 internal, 207
 truth conductiveness, 54–5
- coherentism, 13–14, 42–59
 circularity, 54–5, 58–9, 211, 212
 criticisms, 53–9, 60, 206–12, 222–3
 doxastic presumption, 51–3
 empirical, 13
 and empirical justification, 107–8,
 207–8
 and epistemic justification, 109,
 189–90
 internal, 208
 isolation problem, 43, 53
 moderate, 211, 224–6
 motivations, 43
 nonlinear justification, 44–6
 and sensory observation, 48–51
 vs. foundationalism, 7–9, 109–12,
 204, 207 n.11, 224–6
 weak, 207–8
- concepts, 130–5
 indexical, 124–5, 126–9, 195
 nondemonstrative, 123 n.6
 perceptual, 125
 phenomenal, 126–9, 130–1, 195–6
 thin vs. thick, 124–6, 195
see also simple geometric and
 arithmetical (SGA) concepts
- conceptual analysis, types of, 158
- Conee, Earl, 202, 203 n.3

- consciousness, 137, 215
 and awareness, 119–21
 higher-order thought theory of, 65–8,
 72
- constitutive awareness
 of content, 61–5, 71, 72, 73
 vs. high-order thoughts, 65–8
- content
 constitutive awareness, 61–5, 71, 72,
 73
 of experience, 69–74, 77–83, 86
 non-conceptual, 71–4, 78–9, 81,
 84–5
 propositional, 137
 use of term, 70 n.6
- contextual justification, 161–2
 criticisms, 168–9
- contextualism, 7 n.4
- Davidson, Donald (1917–), 70, 71, 72,
 78
- demons, 161
 Cartesian, 27, 94, 105 n.3, 116, 156,
 184
 controlling, 165
see also new evil demon problem
- Descartes, René (1596–1650), 7, 67, 94,
 116, 156, 188
 criticisms, 166 n.11
 epistemology, 201, 202, 211, 224–6
 foundationalism, 203 n.3, 212–15
 internalism, 144, 173–4, 177
Meditations, 2, 117, 201, 212, 224
- Descartes's paradox, 141–3
- direct apprehension, 17–19, 74–6
- doxastic presumption, 51–3, 56, 57, 69
- doxastic voluntarism, 178
- e-awareness, *see* experiential
 awareness
- empirical justification
 and coherentism, 107–8, 207–8
 of knowledge, 107–9
- empiricism, and foundationalism,
 116–18
- epistemic efficacy, 227–8
- epistemic regress problem, 9–12, 60,
 224
 solutions, 26, 44
- epistemology
 anthropocentric, 160 n.7
 Cartesian, 201, 202, 211, 224–6
 Chisholm's, 201, 211
 definition, 1
 issues, 1–2, 7, 145
 naturalized, 7 n.4
see also virtue epistemology
- Evans, Gareth, 135
- Evans/Peacocke proposal, 135–7
- evil demons, *see* demons
- exclusion, principle of, 141–2
- experience, 130–5
 apperceptive, 62
 and awareness, 119–20
 and basic beliefs, 17
 conscious, 62–3, 69–70, 73–4, 190,
 214 n.20, 217–18
 content, 78–9 n.3, 86
 direct acquaintance, 17–18
 introspective, 13
 non-conceptual, 71–4, 78–9, 81,
 84–5, 87
 visual, 164, 220
see also sensory experience
- experiential awareness, 121, 190–2
 characteristics, 193
 concept of, 120
- external virtues, vs. internalist
 foundationalism, 1–2
- external world, 77–96, 226
- externalism, 13–14, 24–6
 content, 225 n.34
 criticisms, 26–30, 40, 60, 114–15,
 204–6, 222–3
 epistemic, 143–5
 and epistemic justification, 7–9,
 24–41, 109, 158–61
 qualified, 30–3
 and skepticism, 39–40
 vs. internalism, 7–9, 13, 34, 35–41,
 173–90, 204
see also reliabilism

- externalist foundationalism, 7, 108–9, 204
- Feldman, Richard, 128 n.8, 202, 203 n.3
- first-person perspectives, and
internalism, 37–8, 39, 174
- Foley, Richard, 163–4, 186, 202, 203 n.3
- foundational beliefs, *see* basic beliefs
- foundationalism, 12, 13–14, 60–76, 108–9, 116, 126, 127–8, 135, 198
BonJour's, 208, 217–22
Cartesian, 203 n.3, 212–15
criticisms, 55–6, 60–1, 70–2, 119–40, 212–17
defense of, 61–4, 66–7, 69–70, 190–7, 212–17, 222–3
empiricist perspectives, 116–18
and epistemic justification, 7, 109, 140
issues, 137, 203 n.3
modified, 129
rationalist perspectives, 116–17
redux, 116–18
sources, 202–3
vs. coherentism, 7–9, 109–12, 204, 207 n.11, 224–6
see also externalist foundationalism;
internalist foundationalism
- foundherentism, use of term, 225 n.33
- Fumerton, Richard, 17, 56 n.24, 129–30, 193 n.22, 202, 203 n.3
- Gettier, Edmund (1927–), 103–4, 161
- Gettier problem, 23 n.19, 109–10
and justification, 106–7
- goals
incoherence of, 105
of inquiry, 104–5
- God hypothesis, 85–6, 93, 94, 95
- Goldman, Alvin (1938–), 30 n.6, 157 n.3, 160 n.7, 185, 204 n.8
- Greco, John, 55 n.21, 57 n.25
- Haack, Susan, 81–2, 187
- Harman, Gilbert (1938–), 47 n.9
- higher-order thoughts, 62
theory of consciousness, 65–8, 72
- holism, 45–6, 51 n.16, 52 n.17, 210–11
- Hume, David (1711–76), 7, 117, 215
- Husserl, Edmund Gustav Albrecht (1859–1938), 20
- illusions, perceptual, 153
- immediate acquaintance, *see* direct apprehension
- incoherence, concept of, 105 n.3
- indexical affirmations, 122–3
- inference, 154–5
- inquiry, goals, 104–5
- intellectual awareness, 121, 190–2
concept of, 120
- intentionality, 130–5
- internal/external divide, and
skepticism, 141–55
- internalism
Cartesian, 144, 173–4
Chisholmian, 144, 145–9, 173–5
criticisms, 33–5
epistemic, 143–5
and epistemic justification, 7–9, 25, 30–1, 144, 150–5, 175, 179–80, 182–3
first-person perspectives, 37–8, 39, 174
indispensability, 39–41
ontological, 146–7, 177–8
rationalist, 152
and virtue epistemology, 162–5
vs. externalism, 7–9, 13, 34, 35–41, 173–90, 204
- internalist foundationalism, 7, 13, 60, 116, 204
Cartesian, 8
criticisms, 2, 8, 119–40, 173, 190
defense of, 8–9, 190–7, 198–200, 212–17
issues, 203 n.3, 226–8
- introspection, 118, 223

- introspective judgments
 justified, 138
 unjustified, 138
 intuition, 154–5
 isolation problem, 53
- J-ADROIT, 157
 J-APT, 156–7
 judgments, 203–4 n.5
 perceptual, 127 n.7
 safe, 138
 virtuous, 138
 see also introspective judgments
- justification, epistemic, 2, 5–6, 9–13,
 28–30, 34, 103, 107–8, 109–18,
 138
 a priori, 6 n.2, 198
 of basic beliefs, 15–20, 24–6, 61–5,
 69–74, 190–7
 of beliefs, 5–6, 9–12, 34, 107–8,
 109–18
 Chisholmian internalism, 145–9
 coherentist perspectives, 208–9
 and Bonjour's Foundationalist
 Principle, 220–1
 coherence theory of, 13, 42–60, 109,
 189–90
 concept of, 205
 concepts of, 36, 144–9, 162, 182–3,
 199–200
 conceptual analysis, 158
 deontological perspectives, 145–9,
 174, 175–6, 177, 178
 and externalism, 7–9, 24–41, 109,
 158–61
 externalist perspectives, 24–41
 foundational, 127–8, 226–7
 and foundationalism, 7, 109, 140
 foundationalist vs. coherentist, 7–9
 and Gettier problem, 106–7
 hybrid views, 7 n.4
 internalist perspectives, 25, 144,
 150–5, 175, 179–80, 182–3
 internalist vs. externalist, 7–9, 24–5,
 30–9
 and knowledge, 22–3, 139, 203–4
 meta-justification, 54–5
 nonlinear, 44–6
 objective vs. subjective, 153–4,
 181–2
 personal-aetiology, 151
 and reasoning, 216–17 n.24
 reflective, 209 n.15, 228
 reliabilist, 213
 reliability, 205–6
 social-aetiology, 151–2, 179
 sources, 201–2
 sources of, 107–9
 subjective, 153–4, 162–3, 181–2
 use of term, 5, 168, 174–5, 183
 see also contextual justification;
 empirical justification
- justified, use of term, 141 n.1, 169
- Kant, Immanuel (1724–1804), 7,
 223
- Kantianism, and moral quality,
 158–9
- Kitcher, Philip (1947–), 36
- knowledge
 and ability, 101–2
 anti-Gettier conditions, 179
 and belief, 6, 30, 102–4, 113–14
 concept of, 21–3, 35–6, 104–5,
 183 n.9
 conditions, 30 n.5, 34, 35
 empirical justification, 107–9
 foundational, 137–40
 foundations, 119–40
 Gettier-related perspectives, 161
 of the given, 122–6
 and justification, 22–3, 139, 203–4
 lucky, 114–16
 non-propositional, 99–102
 of objects, 102
 objectual, 100
 procedural, 100–1
 and skills, 101
 strong conception, 22
 superknowledge, 106
 and symbols, 100
 theory of, *see* epistemology

- knowledge (*cont'd*)
 and truth, 106
 weak conception, 22–3
 see also animal knowledge;
 propositional knowledge;
 reflective knowledge
- Lehrer, Keith (1936–), 42
- Leibniz, Gottfried Wilhelm
 (1646–1716), 119–20
- Lewis, Clarence Irving (1883–1964), 8,
 17, 48, 83 n.9, 102
- Locke, John (1632–1704), 7, 85–6, 96,
 188
- “looks” statements, 79 n.5
- Martin, Michael, 135–6
- masochism, epistemic, 149
- mathematics, and rational theology,
 212–13
- Matheson, David, 160 n.6
- meaning analysis, 158
- meliorative epistemological project,
 use of term, 36–7
- memory, 58–9, 154–5, 184 n.11
 hallucinations, 180
- mental states
 beliefs, 136, 198–9, 210
 conscious, 61–5, 68
 consciousness, 73, 75, 222 n.28
 intrinsic, 164–5
- mentalism, 203 n.3
- meta-beliefs, 51–2, 69, 115, 170 n.14,
 197
 justification, 61–2, 64–5, 210
- Molyneux’s problem, 90 n.20
- Moore, George Edward (1873–1958),
 143
- moral quality, Kantian, 158–9
- Moser, Paul (1957–), 78 n.1
- Müller-Lyer lines, 153
- musical notation, 100
- naturalism, 33 n.8
- n-awareness, *see* intellectual awareness
- negligence, 149, 177
- new evil demon problem, 156, 159–60,
 184–6
- Nozickian tracking, 154 n.10, 204
- observation, 57, 118
 coherentist conception of, 48–51
 observation requirement, 51, 58
- paranoia, 108
- Peacocke, Christopher Arthur Bruce
 (1950–), 135, 136
- personal-aitiology (PA) justification,
 151
- perspectives, cognitive, 197
- phenomenalism, 117
- physical world
 inference to from sensory
 experience, 83–96
 phenomenalist view of, 83
- physical-object appearances, 79, 80,
 81–3, 84–7
- Plantinga, Alvin (1932–), 30 n.5,
 145 n.7, 176 n.3
- Plato (c.428–c.348 BC), 1, 102
- Popper, Sir Karl Raimund (1902–94),
 70
- pragmatism, 55 n.23
- prediction, 117
- Price, Henry Habberley (1899–1985),
 84, 86, 87–91, 188 n.16
- privileged access, 121–2, 137
- propositional knowledge, 99–100, 101,
 102–7, 225 n.33
 early studies, 102–3
- propositions, 99–100
- Pyrrhonism, 166 n.11, 224
- Quine, Willard Van Orman
 (1908–2000), 42 n.1
- rationalism, and foundationalism,
 116–17
- rationality
 epistemic, 29–30, 32
 Foley’s, 163–4, 186
- realism, direct vs. indirect, 216

- reason, 155, 204 n.5, 206
- reasonable, use of term, 141 n.1
- reasoning, 184 n.11, 211, 215
and justification, 216–17 n.24
a priori, 1, 214–15
- reflection, 155, 173, 201
definitions, 144 n.5, 146
internal, 207
- reflective knowledge, 209 n.15
vs. animal knowledge, 2
- regress problem, *see* epistemic regress problem
- Reid, Thomas (1710–96), 166 n.11
- reliabilism, 25–6, 37, 159–60, 184–5,
221–2, 225 n.33
- reliability, 161, 205–6, 223
concept of, 26 n.2
- Rescher, Nicholas (1928–), 42,
55 n.23
- Rorty, Richard (1931–), 6 n.3, 70
- Rosenthal, David, 66–8
- Russell, Bertrand (1872–1970), 119–20
- Ryle, Gilbert (1900–76), 121
- Sacks, Oliver Wolf (1933–), 132
- self-evidence, 16
- self-knowledge, 225 n.34
- Sellars, Wilfrid (1912–89), 42, 47 n.9,
49, 60, 70, 79 n.5
Sellarsian dilemma, 19 n.16, 63 n.2,
70, 212, 216
- sense-data, 219–20
families of, 90–1
nuclear, 89, 90
and sensory experience, 78–9 n.3,
80, 86, 88–92
- sensory experience, 13, 60, 127 n.7,
188
and beliefs, 217–18
conceptualization of, 77–83
inference from, to physical world,
83–96
justification of beliefs concerning,
69–74, 107, 217–18
phenomenological characterization
of, 78–9, 80, 86, 87–92
physical-object characterizations of,
79–80, 81–3
sense-datum view, 78–9 n.3, 80, 86,
88–92
see also experience
- simple geometric and arithmetical
(SGA) concepts, 126–7, 129, 195
- introspective, 128
- skepticism, 9–10, 34–5, 56, 197–200,
223
and externalism, 39–40
and internal/external divide, 141–55
- skills, and knowledge, 101
- social-aetiology justification, 151–2,
179
- Socrates (469–399 BC), 150
- Sosa, Ernest, 10 n.8, 35 n.10
epistemology theory, 1, 2
- speckled hen problem, 121, 127,
131–3, 137–40, 191–2, 194
and BonJour's Foundationalist
Principle, 219–22
- Steup, Matthias, 63 n.2
- substantive analysis, 158
- superknowledge, concept of, 106
- supervenience, 174
- symbols, and knowledge, 100
- telos, *see* goals
- the given, Sosa's critique, 190–7
- theology, rational, 212–13, 226
- thoughts
conditional, 133–4
conscious, 61–5
indexical, 128
phenomenal, 128
wholly appropriate, 177
see also higher-order thoughts
- truth
and knowledge, 106
realist concept of, 5
- V accounts, 161, 162–3, 168, 169
determinants, 159
use of term, 157
- V-ADROIT, 157, 161, 163, 169

V-APT, 157, 160, 161, 163, 169
virtue epistemology, 156–70
 and adroitness, 156–8, 185–6
 and aptness, 156–8, 184–5, 189,
 197
 circularity, 165–8
 criticisms, 184–90
 and internalism, 162–5
 and skepticism, 165–8

virtues
 external, 1–2
 intellectual, 156, 184–5
 reliable, 170
 see also cognitive virtues
visual experience, 164, 187, 220
visual images, 121, 124, 129–30,
 131–4, 137, 219–21
 awareness of, 191–3