

Putting Skeptics in Their Place

THE NATURE OF SKEPTICAL
ARGUMENTS AND THEIR ROLE IN
PHILOSOPHICAL INQUIRY

John Greco
Fordham University



PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK <http://www.cup.cam.ac.uk>
40 West 20th Street, New York, NY 10011-4211, USA <http://www.cup.org>
10 Stamford Road, Oakleigh, Melbourne 3166, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain

© John Greco 2000

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2000

Printed in the United States of America

Typeface Bembo 10.5/13 pt. *System* DeskTopPro_{UX} [BV]

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

Library of Congress Cataloging in Publication Data

Greco, John

Putting skeptics in their place : the nature of skeptical arguments and their role in
philosophical inquiry / John Greco

p. cm. – (Cambridge studies in philosophy)

Includes bibliographical references and index

ISBN 0-521-77263-X (hardcover)

1. Skepticism. I. Title. II. Series.

B837 .G74 2000

149'.73-dc21

99-042103

ISBN 0 521 77263 X hardback

Contents

<i>Preface</i>	<i>page</i> xiii
1 The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry	1
1. The Three Theses	2
2. Strategy and Outline of the Book	6
3. Three Criteria for an Adequate Theory of Knowledge	15
2 Skepticism about the World: Part One – Reconstructions	25
1. An Argument from Hume	25
2. An Argument from Descartes	34
3. Need the Skeptic Challenge All of Our Knowledge at Once?	39
4. Two More Ways Not to Understand Descartes' Argument	51
5. Discriminating Evidence	54
6. Supporting Evidence	58
3 Skepticism about the World: Part Two – Dismissive Responses	61
1. Charges of Self-Refutation	61
2. Pragmatic and Rhetorical Responses	64
3. More Dismissive Responses	69
4. Transcendental Arguments	71
4 Skepticism about the World: Part Three – Dualism, Realism, and Representationalism	77
<i>I. What Is Necessary for Skepticism about the World</i>	83
1. The Modern Ontology and Skepticism	83
2. Realism and Skepticism	88
3. Representationalism and Skepticism	90

II. <i>What Is Sufficient for Skepticism about the World</i>	94
1. The Big Mistake	95
2. The Modern Ontology Again	102
3. Representationalism Again	103
4. Conclusions	106
5 The Argument from an Infinite Regress of Reasons	108
1. The Regress Argument and Strong Particularism	108
2. Foundationalism	118
3. Contextualism	121
4. Coherentism	126
5. The Objection to Traditional Coherence Theories	128
6. The Objection to Nontraditional Coherence Theories	130
7. The Role of Sensory Appearances	134
8. Conclusions	135
6 Hume's Skepticism about Unobserved Matters of Fact	137
I. <i>Hume's Arguments and Dismissive Responses</i>	138
1. The Arguments in Sections IV and VII of the <i>Enquiry</i>	138
2. Objection (a): Hume Relies on an Inadequate Theory of Ideas	141
3. Objection (b): Hume Requires Absolute Certainty for Knowledge	145
II. <i>The Standard Objection against Hume</i>	146
1. The Objection	147
2. Deductive, Inductive, and Nonsupportive Inferences	149
3. An Objection Considered (against Our Interpretation of Hume)	155
4. The Lesson of Hume's Skeptical Argument	159
7 Agent Reliabilism	164
I. <i>Simple Reliabilism</i>	165
1. Simple Reliabilism: The Big Idea	165
2. Why Skeptical Arguments Go Wrong	167
3. From Processes to Virtues	174
II. <i>Agent Reliabilism and the Question of Subjective Justification</i>	180
1. Knowing That One Knows	181
2. Understanding That One Knows	184
3. Sosa's Perspectivism	187
4. From Perspectives to Dispositions	190
5. The Place of Epistemic Norms	192
6. The Place of Epistemic Responsibility	200
7. Conclusions	202

8	Agent Reliabilism and the Relevant Sense of “Relevant Possibility”	204
	1. Two Tasks for Working out This Approach	205
	2. The Relevant Sense of “Relevant Possibility”	207
	3. A Theory of Virtues and a Virtue Theory of Knowledge	211
	4. Agent Reliabilism and Relevant Possibilities	217
9	Moral and Religious Epistemology	220
	<i>I. Religious Epistemology</i>	221
	1. Plantinga’s Rejection of the Evidentialist Objection to Religious Belief	222
	2. Alston’s Defense of Religious Perception	226
	<i>II. Moral Epistemology</i>	231
	1. Three Kinds of Moral Epistemology	232
	2. A Theory of Empirical Perception	235
	3. A Theory of Moral Perception	241
	4. Some Traditional Objections to Moral Perception	244
	<i>III. General Conclusions</i>	248
	<i>Bibliography</i>	255
	<i>Index</i>	261

1

The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry

This book has three major theses: (1) that a number of historically prominent skeptical arguments make no obvious mistake and therefore cannot be easily dismissed; (2) that the analysis of skeptical arguments is philosophically useful and important and should therefore have a central place in the methodology of philosophy, particularly in the methodology of epistemology; and (3) that taking skeptical arguments seriously requires us to adopt an externalist, reliabilist epistemology. More specifically, it motivates a position that I call “agent reliabilism,” which is an externalist version of virtue epistemology.

If these theses are correct, then many philosophers have misunderstood the nature of skeptical arguments and their role in philosophical inquiry. For example, many philosophers think that skepticism poses no philosophically interesting problem. According to this view, skeptical arguments rest on some obvious mistake, such as a quest for absolute certainty or a demand for immutable foundations, and can therefore easily be dismissed. Others think that skepticism rests on a substantive philosophical mistake, but that skeptical arguments teach no epistemological lessons. For example, many philosophers think that skepticism is rooted in a bad ontology. On this view skeptical arguments assume an ontological dualism between knowing mind and material object of knowledge and can therefore be rejected by rejecting the offending dualism. Others have thought that skepticism is rooted in representation-alism, and still others that it is rooted in realism. Finally, some philosophers have appreciated that skepticism is indeed an epistemological problem, but have tried to solve it by remaining within a traditional, internalist epistemology. Against all of these positions, I argue that the recent externalist revolution in epistemology is necessary for a quite

traditional reason: to adequately address a range of well-known skeptical arguments.¹

1. THE THREE THESES

My first thesis is that a number of historically prominent skeptical arguments make no obvious mistake. On the contrary, such arguments begin with assumptions about knowledge and evidence that seem eminently plausible, outside the context of philosophical inquiry. Often they are assumptions that we ourselves accept either explicitly or implicitly. But by reasoning that is seemingly cogent, such arguments “prove” a conclusion that is outrageously implausible, even incredible in the literal sense. Accordingly, skeptical arguments are powerful in the following sense: it is not at all easy to see where they go wrong, and rejecting them requires one to adopt substantive and controversial theses about the nature of knowledge and evidence. This is not to say that they are powerful in a psychological sense – that they have the power to persuade. In this respect skeptical arguments are like arguments for God’s existence: it is doubtful that any has ever produced a convert.

My second thesis is methodological and is closely tied to the first. Specifically, I argue that the analysis of skeptical arguments is philosophically useful and important. This is not because skepticism might be true and we need to assure ourselves that we know what we think we know. Neither is it because we need to persuade some other poor soul out of her skepticism. Rather, skeptical arguments are useful and important

1 Some philosophers do take skeptical arguments seriously, giving them pride of place in their own methodology. For recent discussions that endorse my first two theses, see Peter Klein, *Certainty: A Refutation of Scepticism* (Minneapolis: University of Minnesota Press, 1981); Barry Stroud, *The Significance of Philosophical Scepticism* (Oxford: Clarendon Press, 1984); Robert Fogelin, *Pyrrhonian Reflections on Knowledge and Justification* (Oxford: Oxford University Press, 1994); and Michael Williams, *Unnatural Doubts* (Princeton: Princeton University Press, 1996). None of these authors endorses my third thesis, however. Williams argues that skeptical arguments mistakenly assume a thesis he calls “epistemic realism,” while Klein argues for a defeasibility theory of knowledge. Stroud and Fogelin suggest that certain skeptical arguments cannot be answered in any satisfactory manner. In the context of the more usual attitudes regarding the nature and usefulness of skeptical arguments, however, my disagreements with these authors come fairly late in the day. Much of what I say in the book, in fact, overlaps with one or another of them. Finally, a good number of analytic epistemologists – reliabilists and externalists among them – engage in the analysis of skeptical arguments episodically as part of their methodology, and it is fairly clear that their own accounts of knowledge and evidence are motivated by this. In what follows I try to give this common practice an explicit articulation and systematic defense. I also recommend a more consistent application of it.

because they drive progress in philosophy. They do this by highlighting plausible but mistaken assumptions about knowledge and evidence, and by showing us that those assumptions have consequences that are unacceptable. As a result we are forced to develop substantive and controversial positions in their place. On this view skeptical arguments are important not because they might show that we do not have knowledge, but because they drive us to a better understanding of the knowledge we do have. There is another side to this coin: the price of ignoring skeptical arguments, or of rejecting them for the wrong reasons, is that we miss the lessons that the arguments can teach us.

Again, if these two theses are correct, then philosophers from a wide range of traditions have badly misconceived the nature and place of skeptical arguments. As we have already noted, many philosophers think that skeptical arguments make some obvious mistake and may therefore be easily dismissed. One example of this position is the view that skeptical arguments require some high level of certainty for knowledge – perhaps infallible premises, or incorrigible ones. Another is the view that skeptical arguments allow only deductive inferences. According to these diagnoses, to reject skepticism we need only reject its implausible standards for knowledge. Yet another widely held view is that skepticism is self-refuting, either because the skeptic makes incompatible claims (she knows that no one knows), or because skepticism cannot be lived and so the skeptic is caught in a “performative” contradiction.

If any one of these diagnoses were correct, then the analysis of skeptical arguments would not be philosophically useful or important. Such easy refutations would teach no lesson at all, or perhaps only one learned long ago – that rationalism is false. But against these views I argue that historically prominent skeptical arguments – for example, ones from Aristotle, Descartes, and Hume – make no obvious mistake and therefore cannot be easily dismissed. These skeptical arguments can be interpreted so that they involve no contradiction, performative or otherwise. Moreover, a close investigation of such arguments reveals that they run on assumptions much more dear to us than infallibilism or deductivism. Their lesson is not that some vestige of rationalism is false, but that something much more plausible will have to be given up.

There is another view of skepticism that is widely popular among philosophers, and that is incompatible with the one I am defending here. Philosophers from a surprising range of traditions claim that skeptical arguments presuppose a Cartesian ontology of internal subjects and external objects. Once the offending dualism is given up, these philoso-

phers claim, skeptical arguments cannot even get off the ground. A strong version of this diagnosis implicates not only skepticism but epistemology as well. The idea is that epistemology is essentially the activity of constructing solutions to skeptical problems. But if skeptical problems cannot arise in a post-modern world, then epistemology is robbed of its purpose, and therefore of its existence.

Other philosophers think that skepticism is grounded not in a bad ontology but in a bad philosophy of mind. These philosophers make representationalism the root of all skeptical arguments. Still other philosophers think that a bad theory of reference is the problem, and still others think that the skeptic makes some kind of linguistic mistake. Against all of these positions I want to argue that skeptical arguments run on mistaken assumptions about the nature of knowledge and evidence. A close analysis of skeptical arguments drives positive epistemology, not ontology, or philosophy of mind, or philosophy of language.

My third thesis is that taking skeptical arguments seriously pushes us in a particular direction in epistemology. Specifically, it pushes us toward externalism and reliabilism. Even more specifically, it pushes us toward agent reliabilism. The idea is this: reconstructed in their most plausible form, a number of skeptical arguments show, quite correctly, that there is no necessary relation between our beliefs and their evidential grounds. It is now a commonplace to recognize that there is no such deductive relation. The more interesting point, however, is that there is no necessary inductive relation either; it is not a necessary truth that the grounds for our beliefs make them even probable. For many philosophers this contention would be enough to entail skepticism. If there is no logical or quasi-logical relation between our evidence and our beliefs, as some would require for our cognition to be “within the logical space of reasons,” then a fundamental condition of knowledge goes unfulfilled.

There is in fact no such relation, however. This is one of the most important lessons that skeptical arguments teach us. A necessary condition for avoiding skepticism, therefore, is to rethink what it is for the grounds of our beliefs to be good evidence. Put another way, it is necessary to rethink what it is to be within the space of reasons. As it turns out, that space is neither logical nor quasi-logical. It is at most a contingent fact that the grounds for our beliefs are reliable indications of their truth, and any adequate epistemology must account for this.

The relevance of all of this to reliabilism is now easy to see. Taking skeptical arguments seriously provides a powerful motivation for reliabilism in epistemology, insofar as reliabilism can explain why evidence

need not be logical or quasi-logical. According to reliabilism, a belief has positive epistemic status (roughly) just in case it is in fact reliably formed. Put in terms of evidence, the grounds on which a person forms her belief amount to good evidence (again roughly) just in case those grounds are in fact a reliable indication that the belief is true. There is no requirement that the person know that her grounds are reliable, or even that she could know this on reflection. The latter requirements are plausibly fulfilled if we think that evidential relations are necessary. In that case, one might expect a kind of *a priori* insight into the fact that one's evidence indicates the truth of one's belief, either by entailing it or by necessarily making it probable. But there is no possibility of fulfilling such a requirement if the relation between our evidence and our beliefs is merely contingent. On the contrary, a requirement in that direction leads straight into skepticism. Reliabilism makes no such requirement and gives us a general approach to knowledge and evidence that explains why none is needed.

There are serious problems with generic reliabilism, however – serious enough to cause some philosophers to reject the position out of hand. One problem is that beliefs can be reliably formed by accident – for example, by arbitrarily adopting a method which, unknown to the believer, happens to be reliable. This would seem to violate a “no accident” condition on knowledge. A second problem is that beliefs can be reliably formed and yet subjectively inappropriate. This seemingly violates a “subjective justification” condition on knowledge, and reliabilism has been widely criticized on this point. Agent reliabilism addresses both problems by drawing on the resources of virtue theory. The main idea is to define knowledge in terms of virtuous cognitive character, and to define virtuous character in terms of proper motivation and reliable success. This takes care of the “no accident” condition on knowledge, in that true belief which is formed through an agent's reliable character is not an accident in any relevant sense. It takes care of the “subjective justification” condition as well, since there is proper motivation, and, as Aristotle would say, “the moving principle is within the agent.” Roughly, a belief is both subjectively and objectively justified, in the sense required for knowledge, when it is produced by a properly motivated, reliable cognitive character.

Agent reliabilism is therefore a kind of virtue epistemology, in the sense that it makes cognitive or intellectual virtue central in the analysis of important epistemic concepts. As such, it is an improvement over previous versions of reliabilism, including process reliabilism, method

reliabilism, and evidence reliabilism. All of these are subject to one or both of the two problems that I have mentioned, precisely because they fail to ground knowledge in the virtuous character of the knower. In defending agent reliabilism I do not pretend to offer a position that is either wholly original or fully worked out in its details. Rather, I am defending a general direction in epistemology. This direction has been taken by others, both historically and more recently. I am arguing that it is necessary in the context of well-known but underappreciated skeptical considerations.²

2. STRATEGY AND OUTLINE OF THE BOOK

So I have three major theses: one about the structure and content of skeptical arguments, one about their methodological role in philosophical inquiry, and one about where this methodology leads us. My strategy for establishing these is to engage in five tasks.

One thing I will have to do is to consider and reject dismissive responses to skepticism. I define a “dismissive response” as one that either (a) does not engage skeptical arguments at all or (b) engages them only superficially. Such responses are “dismissive” because they reject the skeptical conclusion without seriously considering the reasoning that leads to it. I include charges of self-refutation under “type-a” dismissive responses. Type-a responses do not consider skeptical arguments at all but rather react to the implausibility of the skeptical conclusion. But skepticism is not self-refuting in any philosophically interesting sense. To think that it is blinds one to the more subtle mistakes that skeptical arguments make, and that many non-skeptical philosophers make as well.

Under type-b dismissive responses I include the charges that skepticism assumes infallibilism or deductivism. We will see that these responses depend on uncharitable readings of the skeptical arguments and so are rightly classified as dismissive responses; if they engage the arguments at all, they do so only superficially.

My second task is to consider and reject non-epistemological diag-

2 Agent reliabilism has its historical roots in Aristotle, Aquinas, and Thomas Reid, among others. More recently, versions of the position have been defended by Ernest Sosa, *Knowledge in Perspective* (Cambridge: Cambridge University Press, 1991); Alvin Goldman, *Liaisons: Philosophy Meets the Cognitive and Social Sciences* (Cambridge, MA: MIT Press, 1992) and Alvin Plantinga, *Warrant and Proper Function* (Oxford: Oxford University Press, 1993).

noses of skepticism. These diagnoses often engage skeptical arguments quite seriously ultimately fail to understand them correctly. For example, skeptical arguments can be reconstructed so that they do not presuppose a modern ontology. Neither do such arguments depend on representationalism, or some version of a traditional theory of ideas. Neither do they depend on realism. Again, the lessons that skeptical arguments teach are about the nature of knowledge and evidence, rather than ontology or the philosophy of mind.

These first two tasks constitute indirect defenses of my thesis about the nature of skeptical arguments – the thesis that such arguments run on plausible assumptions about the nature of knowledge and evidence. My third task is to make good on this claim by analyzing some historically prominent skeptical arguments, and by showing exactly what epistemological mistakes they do in fact make. Here I focus on four arguments in particular: the argument from an infinite regress of reasons, from Aristotle's *Posterior Analytics*; the argument for skepticism about the world, from Descartes' *Meditations on First Philosophy*; the argument for skepticism about the world, from Section XII of Hume's *Enquiry concerning Human Understanding*; and the argument for skepticism about unobserved matters of fact, from Section IV of the same work.

The task here is to look at these arguments as they actually appear in their texts, but also to reconstruct them so as to bring out their real force. The methodology I defend instructs us to put the arguments in their most powerful form rather than rely on historical particularities to score hollow victories. For example, Descartes begins his skeptical inquiry by doubting the validity of all of his knowledge at once. But is there anything essential in Descartes' reasoning that requires him to present it that way? I argue that there is not, and so rejecting it on that basis does not give us insight into the real force of Descartes' skeptical considerations. Likewise, Hume's skeptical arguments are couched in terms of his empiricist theory of ideas. But his reasoning does not essentially depend on that theory, and therefore Hume's arguments cannot be rejected on that basis.

This puts me in position to undertake a fourth task. Having identified several plausible but mistaken assumptions about the nature of knowledge and evidence, I offer a theory that shows why these are mistakes. The idea is that a theory of knowledge and evidence that explains why important skeptical arguments go wrong, and that therefore preserves our common sense intuitions about what we do and do not know, is

made plausible by virtue of having those features. This makes good on my second thesis: that the analysis of skeptical arguments is philosophically useful because it drives positive epistemology.

My fifth task is to argue that the methodology I defend can be extended to moral and religious epistemology. In this way I continue to defend the thesis that skeptical arguments repay analysis, now by driving us to deeper understanding of moral and religious knowledge. Here the methodology can be extended in two ways. First, it can be extended directly, by applying it exactly as we do in the investigation of empirical knowledge. When doing empirical epistemology, we assume that we have knowledge, and we use skeptical arguments to root out assumptions that entail that we do not. In this way we uncover mistaken assumptions about the nature of empirical knowledge and evidence. This same methodology can be applied in moral and religious epistemology as well. For example, we can start with the assumption that we do have moral knowledge. We can then use skeptical arguments that conclude otherwise as heuristic devices for rooting out mistaken assumptions about the epistemology of moral beliefs. The methodology is as legitimate here as it is in the empirical realm. It is implausible to claim that I do not know that here is a hand, and any argument that concludes that I do not is almost certainly mistaken somewhere. But it is equally implausible to claim that I do not know that killing innocent children is wrong. If a moral epistemology entails that I do not, then we have good reason to think that there is something wrong with that moral epistemology.

The methodology can also be extended indirectly. In this case we do not assume that we have moral or religious knowledge, or even that some of these beliefs are more reasonable than others. Rather, we examine arguments for moral and religious skepticism, and we look for assumptions that, if true, entail that there is no empirical knowledge. If an argument for skepticism in the moral or religious realm can be shown to involve such an assumption, then we are warranted in rejecting it for that reason. Some arguments against the rationality of religious belief have exactly this character. Such arguments are meant to show that rational religious belief is impossible. But if these same arguments entail that I am not rational in believing that here is a hand, then something in them is almost certainly wrong. Accordingly, they do not have force against the rationality of religious belief. Some arguments against moral perception have an analogous character: if they were correct, they would show that moral perception is impossible, but they would also show that

empirical perception is impossible. Since the latter claim is implausible, we have reason for rejecting the relevant objection to moral perception.

That is my five-point strategy for defending the three major theses of the book. The outline of the book is as follows: In Chapter 2 I look closely at two important arguments for skepticism about the world, one from Descartes and the other from Hume. The purpose here is to reconstruct the arguments in their most powerful forms. I conclude that the strongest version of Descartes' argument is an inarticulate version of Hume's. The main idea is that there is no good inference (deductive or inductive) from the way things appear to the way things are. Put another way, there is no good argument, not even an inductive one, from appearance to reality. Far from relying on implausibly high standards for knowledge, this argument offers powerful considerations for the conclusion that we lack even inductive knowledge of the world.

Another version of Descartes' argument, less powerful but still formidable, trades on the plausible assumption that knowledge must discriminate truth from alternative possibilities. For example, it is claimed that I cannot know that I am sitting by the fire if I cannot discriminate this state of affairs from the possibility that I am a disembodied spirit deceived by an evil demon. I suggest that a "relevant possibilities" approach is promising here. The central idea is that, intuitively, knowledge requires only that we can discriminate among some possibilities, while others can be ignored as irrelevant. But such a response requires development. We would like a principled account of what makes some possibilities relevant and others irrelevant. Such development is undertaken in Chapter 8.

In Chapter 3, I consider and reject several dismissive responses to the reconstructed arguments from Chapter 2. These include charges of self-refutation, as well as several other responses based on pragmatic and rhetorical considerations. I claim that all of these responses miss the mark, since the skeptical arguments retain their force even if there are no skeptics to put the arguments forward. What gives skeptical arguments their force is not that some *other* person is willing to defend them. Rather, it is that they begin from premises that *we* are inclined to accept, and that seemingly entail conclusions that we do not accept. If this is right, then skeptical arguments are a problem for us, and whether or not there are any "real" skeptics to defend them. Pragmatic and rhetorical considerations are therefore irrelevant for adequately answering skeptical arguments. They would be relevant if the problem of skepticism involved some *skeptic*, at whom such considerations could be directed. But

if the problem of skepticism is the problem of analyzing skeptical *arguments*, such considerations are simply irrelevant.

Another kind of dismissive response considered in Chapter 3 charges that the standards for knowledge assumed by skeptical arguments are too high, requiring, for example, deduction from evidence or absolute certainty for premises. Here I argue that the arguments reconstructed from Hume and Descartes require no such thing, and that a close analysis of the arguments reveals this. Finally, several versions of transcendental arguments are considered as dismissive responses to skepticism, and all are rejected as inadequate.

In Chapter 4, I consider some non-epistemological responses to “no good inference arguments” for skepticism about the world. These include the diagnosis that skepticism requires a dualism between knowing mind and material object of knowledge. Alternatively, some philosophers see representationalism as the driving force behind skeptical arguments, while others claim that realism is the problem. I reject all of these diagnoses, arguing that “no good inference” arguments can be reconstructed without dualism, representationalism, or realism. In fact, even Berkeley’s radical idealism is consistent with a charitable reconstruction of the skeptical argument.

These various non-epistemological diagnoses of skepticism are rejected in favor of an epistemological one. Specifically, I contend that “no good inference arguments” misunderstand the way that sensory appearances act as evidence for beliefs about the world. What these arguments assume, and what many non-skeptical philosophers assume with them, is that all evidential relations are inferential relations. In other words, they assume that sensory appearances can act as evidence for beliefs about the world only if the latter are inferred from the former. But since no such inference is forthcoming, the arguments conclude that appearances cannot give rise to knowledge of the world. My position is that skeptical arguments are correct in claiming that there is no good inference from appearance to reality and are therefore wrong in claiming that beliefs about the world must be inferred from sensory appearances. The latter is the plausible but ultimately disastrous assumption that many skeptical arguments and many non-skeptics mistakenly share.

In Chapter 5, I consider the ancient skeptical argument from an infinite regress of reasons. This argument contends that all knowledge requires justification by adequate evidence, and that all such justification involves inference from good reasons. But since all good reasons require

further reasons for their evidence, knowledge requires an infinite regress of justifying reasons, and therefore no knowledge is possible. I begin by revisiting some dismissive responses to skepticism and showing that these miss the mark against the regress argument as badly as they do against Cartesian and Humean arguments. For example, the argument does not trade on a requirement for infallible reasons, or irrevocable ones. Second, I look at the three most popular non-dismissive responses to the regress argument: foundationalism, coherentism, and contextualism. Here I argue that coherentism is psychologically implausible, and that plausible versions of contextualism reduce to foundationalism. Accordingly, I defend a contextualist version of foundationalism – one that is not open to the usual objections to the foundationalist position.

“Contextualist foundationalism” might sound like an oxymoron, but only because insufficient attention has been given to what foundationalism requires. Here again a close analysis of skeptical arguments proves instructive, in this case showing what foundationalism does and does not require to stop the infinite regress of justifying reasons. For example, it does *not* require that foundational beliefs be infallible, incorrigible, devoid of contextual or social features, or even irrevocable. What it does require is, once again, that not all evidential relations are inferential. Perhaps all knowledge must be grounded in evidence, but some knowledge is not inferred from other beliefs. This is the defining characteristic of foundationalism and is what is needed to answer the age-old infinite regress argument.

In Chapter 6 I consider Hume’s skeptical arguments regarding unobserved matters of fact. Here I distinguish two arguments – one from Section IV of the *Enquiry* and one from Section VII. I argue that the former is the most powerful and is immune to dismissive responses that are, at best, effective only against the latter. For example, the argument is in no respect dependent on Hume’s empiricist theory of ideas. I also argue against the most popular interpretation of Hume, namely, that he is a deductivist regarding knowledge of the unobserved. On my interpretation Hume’s claim is that our beliefs about unobserved matters of fact are not even inductively supported by past observations. On this view Hume does not require that our evidence be deductive, but that there be some necessary inductive relationship between our evidence and our conclusions. This makes Hume’s argument far more powerful than is usually supposed. No one thinks nowadays that all evidence must be deductive, and if Hume’s argument ran on that assumption, then it could be easily refuted. But it is seemingly obvious that evidence must

be at least inductively relevant to give rise to knowledge. This is in fact false and is the real lesson of Hume's skepticism about unobserved matters of fact. It is also a good illustration of the claim that skeptical arguments drive positive epistemology. If we did not see its skeptical consequences, we would hardly be inclined to reject such a commonplace, seemingly innocent assumption. But once focused on the real force of Hume's argument, we see that the assumption must go in favor of some better understanding of the nature of evidence.

Chapters 4 through 6 constitute a kind of negative epistemology: rather than saying what knowledge is, negative epistemology largely restricts itself to saying what knowledge is not. This is a worthwhile activity, in that it disabuses us of plausible but mistaken assumptions about the nature of knowledge and evidence. But the ultimate goal of the methodology I defend is to construct a positive epistemology, or a positive account of what knowledge is and how evidence works. This happens when we construct a theory of knowledge that *explains* the largely negative conclusions drawn from our analyses of skeptical arguments. This project is undertaken in Chapter 7.

Notice that the progress from negative to positive epistemology corresponds to three degrees of success we might have in refuting a skeptical argument. The first and least satisfying degree of success is to find an assumption in the argument that we need not accept. If we see that an assumption leads to unacceptable consequences, and if there is no overwhelming reason to accept the assumption, then we are warranted in giving it up as mistaken. For example, suppose we identify as disastrous the assumption that knowledge of objects must be inferred, deductively or inductively, from knowledge of how objects appear. This assumption is something we need not accept, and so we should give it up once we see where it leads.

In the next degree of success we arrive at some reason why the skeptical assumption is mistaken. For example, we conclude that not all evidential relations are inferential, and that this explains why sensory appearances can be evidence for a belief even if the latter is not inferred from the former. This seems to be a plausible position, and we might look for further confirmations of it in cases unrelated to the skeptical argument at hand. But even at this second degree we do not have an explanation of a particular sort. What we really want is to have a theory of knowledge that explains why the skeptical assumption is false. In other words, we want an account that tells us what knowledge is and

how evidence works, and which thereby provides a theoretical explanation of why the assumption in question is mistaken.

In Chapter 7, I defend a virtue theory of knowledge that does just this. Again, by a “virtue theory” I mean one that makes the cognitive faculties and habits of persons central in the analysis of important epistemic concepts. As we saw, the theory is a version of reliabilism, in that a stable disposition of a person counts as a virtue only if it is reliably successful in achieving its end. In the case of the cognitive virtues, this means that the faculty or habit makes the person reliable in forming true beliefs of the kind relevant to the virtue in question. Agent reliabilism explains *why* the skeptical assumptions rejected in earlier chapters are false. Namely, it explains (a) why not all evidence is inferential, (b) how sensory appearances can function as evidence without functioning as premises in an inference, (c) how some knowledge can be foundational, and (d) how propositional evidence that is neither logical nor quasi-logical can give rise to knowledge.

We saw earlier that one of Descartes’ skeptical arguments trades on our inability to discriminate among various alternative possibilities. For example, if one’s evidence does not discriminate between being in front of the fire and being a disembodied spirit deceived by an evil demon, then one cannot know that one is sitting in front of the fire. A promising strategy in response to this kind of skeptical reasoning is to distinguish between relevant and irrelevant alternative possibilities, and to claim that knowledge requires only that we discriminate among the relevant ones. The problem, then, is to give a theoretical account of what makes an alternative possibility relevant or irrelevant. In Chapter 8, I argue that agent reliabilism can do just this.

The main idea is that virtues in general are abilities to achieve some result, and abilities in general are functions of success in relevantly close possible worlds. In other words, to say that someone has an ability to achieve X is to say that she would be successful in achieving X in a range of situations relevantly similar to those in which she typically finds herself. But then possibilities that do not occur in typical situations are irrelevant for determining whether a person has an ability in question. For example, it does not count against Babe Ruth’s ability to hit baseballs that he cannot hit them in the dark. Likewise, it does not count against our perceptual abilities that we cannot discriminate real tables and fires from demon-induced hallucinations. But then our inability to rule out hypothetical demon scenarios is irrelevant to whether we have

knowledge, and the skeptical scenario is not a relevant possibility in that sense. This account of “relevant possibility” confirms the plausibility of agent reliabilism and deepens our understanding of what we must mean by a cognitive virtue or ability.

Finally, in Chapter 9, I argue that the methodology I am defending can be extended to religious and moral epistemology. Many arguments against the rationality of religious belief trade on assumptions about knowledge and evidence that, if true, would count against rational belief and knowledge in the empirical realm. By exposing such assumptions and rejecting objections to religious belief on that basis, the epistemology of religious belief is advanced. In the first part of Chapter 9 I endorse recent work by Alvin Plantinga and William Alston as instances of exactly this methodology, and I offer some suggestions about how their views might be further defended along this line.

In the second part of the chapter I consider the possibility of moral perception. Here I argue that recent work in the epistemology and psychology of empirical perception opens up possibilities for moral perception. Specifically, to avoid skepticism about the natural world we must understand empirical perception as a non-inferential cognitive faculty, but one that is nevertheless influenced by background beliefs, special training, and the like. Second, we must have an account of how complex, dispositional properties can be objects of empirical perception. Accounts of how these features are possible for empirical perception suggest promising extensions to moral perception as well. For example, recent theory concerning the roles of personae and scripts in empirical perception suggests interesting applications to moral perception.

Chapters 4 through 8 suggest a moderate foundationalism. The main idea is that knowledge and justified belief arise from the cognitive abilities of reliable believers, and that some of those abilities must be characterized as non-inferential. The combined results of Chapter 9 suggest that the foundations of knowledge are broad as well. In other words, they suggest that there exists a wide variety of non-inferential sources of evidence, including evidence for moral and religious beliefs. What emerges is a broad and moderate foundationalism in which much evidence is non-inferential, and where even inferential evidence is seldom deductive or inductive in a sense that is commonly supposed. What matters for knowledge and justified belief is not the infallibility or incorrigibility of our premises, nor even the logical or quasi-logical relations among our premises and conclusions. Rather, knowledge and justified

belief arise from virtuous belief formation, where the notion of “virtue” must be understood in terms of the contingent causal and motivational features of our cognition, rather than the necessary or intrinsic features of propositions, evidential relations, or the like.

3. THREE CRITERIA FOR AN ADEQUATE THEORY OF KNOWLEDGE

Before closing this first chapter I want to talk about one more methodological issue. Specifically, I want to suggest three criteria for an adequate theory of knowledge, and to talk about how they are related to the methodology for epistemology that I have been proposing.

First, an adequate theory of knowledge should do a good job of organizing our pre-theoretical intuitions about what cases count as knowledge. In other words, the theory should count as knowledge those cases that intuitively seem to be knowledge, and it should count as not knowledge those cases that intuitively seem not to be. “But whose intuitions count as the right ones?”, our suspicious friends will ask. The answer is the intuitions of us all in our non-philosophical lives. An adequate theory of knowledge should explain why normal people, not people caught in the grip of a philosophical theory, count particular cases as knowledge and other cases as not. “Normal people” includes most non-philosophers, and most philosophers when they are not philosophizing. Among such people there is in fact very wide agreement about which cases do and do not count as knowledge.

Moreover, universal agreement is not required for the methodology being proposed. What is necessary is that there is a wide range of cases that most people would find intuitively obvious. If a theory of knowledge does a good job with these, then that is a strong consideration in its favor. However, vague and contested cases are important as well; a good theory of knowledge should explain *why* certain cases are vague and why certain ones are contested. Consider vagueness first. There will be cases where we are not sure what to say – where we have no strong intuition about whether the case is one of knowledge or not. If our theory can identify some aspect of the case that is vague as described and can tell us that just this aspect is important for knowledge, then that will count in favor of the theory. For example, we saw that agent reliabilism requires that knowers be reliable in forming true beliefs and avoiding error in a relevant domain of inquiry. But how reliable must one be to

have knowledge? This might be left vague by the theory, or might be left vague in the description of a particular case. Either way, this could account for the vagueness of an intuition.

Finally, a good theory of knowledge should account for *disagreement* among intuitions as well. For example, if knowledge requires reliability, and if we disagree over whether someone is reliable in a particular case, then we might disagree over whether the case is one of knowledge. To say that a theory of knowledge should account for our intuitions, then, does not require that our intuitions lack vagueness or enjoy unanimity. A good theory should account for agreement where we agree, disagreement where we disagree, and should explain why vagueness arises when it does.

The second criterion for an adequate theory of knowledge is that it be immune to skeptical arguments. This criterion is related to the first, because our pre-theoretical intuitions are overwhelmingly non-skeptical. Any theory that entails that there is no knowledge of objects in the world, or no knowledge of unobserved facts, or no moral knowledge does a horrible job with our pre-theoretical intuitions about what cases count as knowledge. This is so for *all* of us, since none of us are skeptics in our everyday lives.

For this reason, a theory that has radically skeptical consequences does not capture the concept of knowledge that is actually in use – or at least there is a very strong presumption that it does not. It is, we may suppose, possible in principle that our ordinary concept of knowledge has widely unrecognized skeptical consequences. But if there are competitor accounts that do not have such consequences, then that is an almost insurmountable advantage of those accounts. Again, this follows from the first criterion of an adequate epistemology: that it capture our pre-theoretical intuitions about which cases count as knowledge.

We must take the qualification of the previous paragraph seriously, however, in light of a sophisticated version of skepticism recently suggested by Robert Fogelin.³ Sophisticated skepticism claims that it *can* explain our non-skeptical intuitions, by virtue of certain aspects of our linguistic practices. The idea is that practical purposes make it appropriate to assert knowledge-claims in a wide range of cases, and that this linguistic fact is behind our pre-reflective intuitions that such claims are true. However, at least many of our knowledge-claims are literally false. Moreover, in contexts where practical considerations are put aside and

3 See Fogelin, *Pyrrhonian Reflections on Knowledge and Justification*.

our level of scrutiny regarding knowledge-claims is raised, our intuitions become skeptical, as when skeptical arguments are made or skeptical considerations pushed. In light of this Fogelin-type sophisticated skepticism, we cannot claim that skeptical arguments, by their very nature, cannot account for our pre-theoretical intuitions about which cases count as knowledge. Rather, we have to put emphasis on the claim that a non-skeptical theory of knowledge is strongly preferable, other things being roughly equal. That is, if a non-skeptical theory of knowledge is available that explains why the majority of our intuitions are true, this will be preferable to a skeptical theory, and even if that theory comes with an explanation as to why what seems obviously true is in fact false.⁴

One reason that the non-skeptical theory will be preferable, other things being equal, is that it does not need the extra explanation. Another reason is that our non-skeptical intuitions are both strong and persistent. Fogelin is correct that our intuitions are to some extent unstable – that under the pressure of skeptical arguments it can seem to us that we do not know such things as that here is a hand or this is a pencil. But these fleeting moments of doubt should not be over-emphasized. On the other side are the persistent and overwhelming intuitions of common sense, even among those philosophers who are convinced by skeptical arguments in the study.

For example, consider G. E. Moore's famous statement that he knows that here is a hand, made when his hand was held up in clear view. Any theory that entails that I do not know such a thing, even with an explanation of why it seems obvious that I do know, has its work cut out for it. In this respect, any skeptical theory will be in the position Russell's was in when Moore wrote the following:

What I want, however, finally to emphasize is this: Russell's view that I do not know for certain that this is a pencil or that you are conscious rests, if I am right, on no less than four distinct assumptions. . . . And what I can't help asking myself is this: Is it, in fact, as certain that all these four assumptions are true, as that I *do* know that this is a pencil and that you are conscious?⁵

- 4 I say "Fogelin-type" skepticism because Fogelin's Pyrrhonian principles do not allow him to actually endorse any philosophical theory, including skepticism with respect to ordinary knowledge claims. Therefore, the skeptical position I have just described is not literally Fogelin's, although his discussion strongly suggests it as a possible alternative to non-skeptical theories.
- 5 G. E. Moore, "Four Forms of Scepticism," in *Philosophical Papers* (London: Allen and Unwin, 1959), p. 222.

The more radical the skeptical consequences of a position are, the more strongly this point will hold.⁶

The third criterion for an adequate epistemology is that it be psychologically plausible. What I mean by this is that an adequate theory of knowledge ought to be consistent with our common sense judgments about our own cognitive abilities, and with our best cognitive science as well. This too is related to the first two criteria, because an account that is not psychologically plausible will generate skeptical arguments. We can see this if we look at the most basic structure that any skeptical argument must have.

Although skeptical arguments come in many shapes and sizes, all of them can be boiled down to two essential premises: one stating that knowledge requires that some condition or set of conditions be fulfilled, and one stating that these conditions are in fact not fulfilled. More formally, we have the following skeptical argument structure:

- (SAS) 1. $K \Rightarrow C$.
 2. Not- C .
 3. Therefore, not- K .

Any theory of knowledge that is psychologically implausible will generate an argument with this structure. First, any theory of knowledge whatsoever will entail premises corresponding to premise (1) of (SAS), since any such theory posits conditions that must be fulfilled in order to have knowledge. But the fact that a theory is psychologically implausible guarantees that it will generate a premise corresponding to (2): that is, a premise stating that conditions laid down by the theory in question are not satisfied by beings with our psychology. Accordingly, we will have a skeptical argument amounting to a *reductio ad absurdum* of the theory in question.

This points to an elaboration of the methodology I am defending. Specifically, we are not restricted to the use of historically prominent skeptical arguments and their reconstructions. We can make up new arguments to demonstrate the mistaken assumptions of alternative accounts of knowledge, or even alternative solutions to skeptical problems. This actually happens in contemporary epistemology – it has happened,

6 For an extended argument that our intuitions are not better explained by Fogelin-type warranted assertability maneuvers, see Keith DeRose, “Contextualism: An Explanation and Defense,” in John Greco and Ernest Sosa, eds., *The Blackwell Guide to Epistemology* (Oxford: Blackwell, 1999).

for example, in recent arguments against coherentism.⁷ Coherentism was originally proposed as a non-skeptical response to the argument from an infinite regress of reasons. The idea is that reasons can be mutually supportive, and so no infinite regress of reasons is necessary to ground knowledge in adequate evidence. But a major problem with coherentism is that it fails to give a psychologically plausible account of the role of experience in forming perceptual beliefs. Take, for example, my current perceptual belief that it is raining. The coherentist contends that all beliefs are supported by other beliefs that serve as its evidence. But against this position, it is not psychologically plausible that I infer my belief that it is raining from other beliefs that I have and that act as its evidence. Surely I know that it is raining because I can *see* that it is.

The coherentist will have a story about perceptual knowledge, but it will not be a psychologically plausible one. For example, the coherentist might say that we unconsciously infer beliefs about objects from beliefs about sensory appearances. But it is psychologically implausible that we typically *have* beliefs about sensory appearances, much less infer beliefs about objects from them. And even if we did typically have such beliefs, what evidence do we have for the supposed unconscious inferences? I certainly do not *seem* to make the relevant inferences in perception. On the contrary, the empirical evidence suggests that such inferences are a philosophical invention.

Where does the empirical evidence regarding our cognitive capacities come from, and who gets to decide what that evidence makes psychologically plausible? On the view defended here, there are two principal sources of empirical evidence: our own reflection and empirical psychology.

In many instances a philosophical position can be recognized as psychologically implausible as soon as the question is raised. For example, it seems obvious upon reflection that we do not typically form beliefs about sensory appearances; in the typical case, we form beliefs about tables and trees, not about how tables and trees appear to us. Sometimes a little experimentation can confirm what seems obvious upon reflection. As Thomas Reid points out, objects not in focus present a double image. This is confirmed by placing your finger in front of your face and then focusing on an object in the distance. Attention to appearances will reveal that your finger presents a double image. Alternatively, if you

7 For several examples of this kind of critique of coherentism, see John Bender, ed., *The Current State of the Coherence Theory* (Dordrecht: Kluwer, 1989).

focus on your finger, the object in the distance will present a double image. Reid argues that we must be presented with double images almost all of the time, since almost always some objects are out of focus for us. And yet we do not notice this, confirming the point that we do not typically form beliefs about sensory appearances.⁸ This fact constitutes a devastating objection to coherentism, or to any other epistemology on which empirical knowledge requires beliefs about appearances. For since we do not typically have such beliefs, any such theory has the consequence that we typically lack empirical knowledge.

Finally, we can learn about our cognitive abilities from more rigorous empirical research. For example, some coherence theories have a “total evidence” requirement for knowledge and justified belief, laying down a requirement that rational belief acquisition must be sensitive to the total evidence that the person has at the time. But empirical research shows that people are sensitive only to a small number of their total beliefs at any one time. Here again we see a devastating empirical objection to an epistemological theory; if a theory requires sensitivity to all of the beliefs we have, and if our cognition is not capable of that kind of sensitivity, then the theory has unacceptable skeptical results. In this case the theory results in total skepticism, since the psychologically implausible requirement is a completely general one.⁹

The methodology that I am defending here is an extension of what Roderick Chisholm calls “particularism.”¹⁰ Chisholm argues that we should follow philosophers like Reid and Moore in testing philosophical

- 8 “Thus you may find a man that can say, with a good conscience, that he never saw things double all his life; yet this very man, put in the situation above mentioned, with his finger between him and the candle, and desired to attend to the appearance of the object which he does not look at, will, upon the first trial, see the candle double, when he looks at his finger; and his finger double, when he looks at the candle. Does he now see otherwise than he saw before? No, surely; but he now attends to what he never attended to before. The same double appearance of an object hath been a thousand times presented to his eye before now, but he did not attend to it; and so it is as little an object of his reflection and memory, as if it had never happened.” Thomas Reid, *Philosophical Works*, ed. H. M. Bracken, 2 vols. (Hildesheim: Georg Olms, 1983), vol. 1, p. 164b.
- 9 The point about total evidence requirements is made by Alvin Goldman in *Epistemology and Cognition* (Cambridge, MA: Harvard University Press, 1986), pp. 204–207. There Goldman cites John Anderson, *The Architecture of Cognition* (Cambridge, MA: Harvard University Press, 1983), and Christopher Cherniak, “Rationality and the Structure of Human Memory,” *Synthese* 57 (1985): 163–186. Goldman makes the point specifically against coherentism in his essay “BonJour’s *The Structure of Empirical Knowledge*,” his contribution to Bender, *The Current State of the Coherence Theory*, p. 112.
- 10 Roderick Chisholm, “The Problem of the Criterion,” in *The Foundations of Knowing* (Minneapolis: University of Minnesota Press, 1982).