

K **NOWLEDGE**

P **UZZLES**

*An Introduction
to Epistemology*

STEPHEN CADE HETHERINGTON

KNOWLEDGE PUZZLES

An Introduction to Epistemology

Stephen Cade Hetherington

The University of New South Wales



WestviewPress

A Division of HarperCollinsPublishers

All rights reserved. Printed in the United States of America. No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopy, recording, or any information storage and retrieval system, without permission in writing from the publisher.

Copyright © 1996 by Westview Press, Inc., A Division of HarperCollins Publishers, Inc.

Published in 1996 in the United States of America by Westview Press, Inc., 5500 Central Avenue, Boulder, Colorado, 80301-2877, and in the United Kingdom by Westview Press, 12 Hid's Copse Road, Cumnor Hill, Oxford OX2 9JJ

Library of Congress Cataloging-in-Publication Data

Hetherington, Stephen Cade.

Knowledge puzzles : an introduction to epistemology / Stephen Cade Hetherington.

p. cm.

Includes bibliographical references and index.

ISBN 0-8133-2486-6 (hardcover).—ISBN 0-8133-2487-4 (paperback)

1. Knowledge, Theory of. I. Title.

BD161. H45 1996

121—dc20

95-40996
CIP

The paper used in this publication meets the requirements of the American National Standard for Permanence of Paper for Printed Library Materials Z39.48-1984.

10 9 8 7 6 5 4 3 2 1

Knowledge Puzzles

*For Parveen,
who has taught me so much
worth knowing*

Preface

0.1 The Book's Aim

Epistemology is an intriguing area of an intriguing subject—philosophy. But epistemology is also one of the more difficult areas of one of the more difficult subjects—philosophy. Students should therefore welcome whatever assistance they can get in understanding the concepts and methods of inquiry important to the discipline. My motivation for writing this book is to provide that assistance.

The book is not intended to replace or to compete with the central primary writings in epistemology. Rather, it presents many of epistemology's main ideas in a way that will help you understand those primary writings. In each chapter I introduce one main theme by way of a few puzzles (examples, questions, and issues). I have aimed to stimulate and suggest, not to exhaust (either readers or a topic). Capturing the spirit—the flavor—of a given idea is the point of each chapter.

I have tried not to favor any one epistemological theory, but instead to find questions and puzzles about each. My goal is to present as many views as I can—certainly most of the major ones—as fairly as possible, regardless of whether I agree with them. By discussing each theory or idea nondogmatically and questioningly, I seek not only to convey epistemology's questioning nature but also to remind you time after time of that nature—and of your chance to take advantage of it. Approach each idea or theory with the aim of deciding what *you* think about it.

0.2 The Book's Structure

Those who read straight through the book should find its organization clear as they proceed, one chapter leading thematically into the next. But that is not the only way to use the book, as I shall explain. Before that, though, it might be helpful if I say something about which chapters “belong” together.

Chapter 1 introduces the basic epistemological project—the attempt to understand the nature of knowledge. Chapters 2 through 4 assemble three concepts (truth, belief, and justification) fundamental to that project. So begins our attempt to understand knowledge; should we end there, too? By understanding

these concepts, do we fully understand what knowledge is? Chapter 5 doubts that we do. It asks whether there is more to knowledge than is revealed in Chapters 2 through 4.

Well, is there? Chapters 6 through 15 grapple with that question: What is knowledge? These chapters contain a series of attempts to answer that question—and a collection of puzzles about those attempts. Does this plethora of possibilities suggest that it is not so easy to understand knowledge? Can we ever fully understand it? That is what Chapter 16 asks. Maybe we will never know all that could be known. I am human; you are human. Is it human to not know all? Is it also human to not fully understand the phenomenon of knowing?

Don't despair—yet. Chapter 17 considers an optimistic view of the attempt to understand knowledge, one that interprets all of us as having lots of knowledge. But like all substantive philosophical suggestions, this one is no less puzzling than those that preceded it. Maybe it is not so conceptually simple after all to accord lots of people lots of knowledge.

Should we therefore *deny* them knowledge? With that question, the door opens onto epistemology's famous *skeptical* questions. Skeptics deny us knowledge we thought we had (yes, they really do). We meet a series of them in Chapters 18 through 22.

For most people, though, it is hard to listen to skeptics without wanting to correct them, explaining to them why there is knowledge. Right now, do you, could you, believe—genuinely—that you know nothing? Chapters 23 through 25 present some classic attempts to defuse skepticism.

But skeptics never give up easily. They never have; they never will. Chapter 26 takes us back to modern skepticism's roots. It considers a particularly ancient and basic form of skepticism—one which doubts that people can even have beliefs, let alone knowledge. ("What? Can that be a serious suggestion?" Yes—and no. Wait—and see.)

It is appropriate, then, that the book begins and ends with questioning. It begins by asking what knowledge is and by calling on epistemology to answer that question. It ends (in Chapter 27) by asking whether we could ever adequately answer the question. What can epistemology hope to establish? Can we ever really understand knowledge? Surely we can *try* to understand it. But is that also the most that we can do?

0.3 Using This Book

There are several ways in which a teacher might use this book, since there are different courses for which it is suitable. A general epistemology course (touching on nonskeptical and on skeptical epistemology) could draw upon many, or even all, of the chapters, the choice reflecting a given teacher's preferences. A more specialized course concentrating on skeptical epistemology might use

Chapters 1, 2, 4, and some or all of 16 through 27. Similarly, a course focusing on nonskeptical epistemology could call on some or all of Chapters 1 through 16 and 27.

Note that many chapters contain more than one important subtopic or issue. Note, too, that each chapter begins and ends with a pertinent puzzle. These puzzles could be used in tutorial discussions or for essay questions. Note, finally, that the readings listed at the end of each chapter are not intended to be exhaustive. They are meant only to guide students toward further reading (which will guide them further still—to more reading, which will guide them even further, ever onward). Still, I hope to have listed enough readings from which to choose at least most of the material needed for some engaging epistemology courses.

Some of the chapters will also be useful for many epistemology segments in general introductory philosophy courses. Such courses often include discussion of such topics as truth (Chapter 2), belief (Chapter 3), justification (Chapter 4), Gettier cases (Chapter 5), Descartes's dreaming argument (Chapter 19), his evil demon argument (Chapter 18), and Hume's inductive skepticism (Chapter 20).

0.4 Acknowledgments

I have greatly appreciated the help given to me by Spencer Carr at Westview Press. He has been an accessible, enthusiastic, and sensitive editor. The support staff at Westview has been excellent, too. I have also valued the work of three anonymous reviewers. They made many helpful suggestions and criticisms pertaining to different stages of the manuscript. My wife, Parveen Seehra, did so as well. Her sharp questioning made my task harder, and I am grateful to her for it. This book, like my life, is richer and more enjoyable for Parveen's existence.

Stephen Cade Hetherington
Sydney, Australia

Contents

<i>Preface</i>	xiii
1 Introducing Epistemology	1
1.1 What Is Epistemology? 1	
1.2 Why Do Epistemology? 3	
1.3 Avoiding Dogmatism, 5	
Further Reading, 6	
2 Truth	7
2.1 Knowledge Entails Truth, 7	
2.2 Objective Truth, 7	
2.3 Semantic Paradox, 9	
2.4 Rejecting Objective Truth, 10	
2.5 Finding the Objective Truths, 11	
Further Reading, 12	
3 Belief	13
3.1 Does Knowledge Entail Belief? 13	
3.2 Is Saying Believing? 14	
3.3 Using Beliefs to Find Beliefs, 15	
3.4 Again, Does Knowledge Entail Belief? 16	
Further Reading, 18	
4 Justification	19
4.1 Does Knowledge Entail Justification? 19	
4.2 Pragmatic Justification, 22	
4.3 Justification and Truth, 22	
4.4 Justification and Objectivity, 23	
Further Reading, 24	
5 The Gettier Problem	25
5.1 Gettier Cases, 25	
5.2 Does Knowledge Matter? 27	

5.3 Doing More to Have Knowledge, 28	
Further Reading, 30	
6 Surface Reliability	31
6.1 Believing Reliably, 31	
6.2 Reliability and Gettier Cases, 32	
6.3 Reliability and Guesses, 33	
6.4 Restricting Reliabilism, 33	
Further Reading, 36	
7 Underlying Reliability	37
7.1 Reliability and Belief-Forming Methods, 37	
7.2 Reliability and Good Belief-Forming Methods, 38	
7.3 Reliability and the Senses, 40	
7.4 Reliable Methods and Gettier Cases, 42	
Further Reading, 43	
8 Causality	45
8.1 The Basic Idea of a Causal Account of Knowledge, 45	
8.2 Causality and Gettier Cases, 46	
8.3 Appropriate Causal Connections, 47	
8.4 Causality and A Priori Knowledge, 49	
Further Reading, 50	
9 Defeasibility	51
9.1 The Basic Idea of Defeasibility, 51	
9.2 Defeasibility and Gettier Cases, 53	
9.3 Which Facts Are Undercover Facts? 54	
Further Reading, 55	
10 Social Defeasibility	57
10.1 Realistic Defeasibility, 57	
10.2 Social Defeasibility, 58	
10.3 Social Defeasibility and Gettier Cases, 59	
10.4 Restricted Social Defeasibility, 60	
Further Reading, 62	
11 False Evidence	63
11.1 Avoiding All Falsity, 63	
11.2 Avoiding Significant Falsity, 64	
11.3 Why Avoid Falsity? 65	

11.4	Avoiding More Falsity by Thinking Less, 66	
	Further Reading, 68	
12	Induction	69
12.1	Induction and Gettier Cases, 69	
12.2	What Is Good Inductive Thinking? 70	
12.3	The Paradox of the Ravens, 71	
12.4	The Problem of Grue, 72	
	Further Reading, 74	
13	A Priori Knowledge	75
13.1	The Basic Idea of A Priori Knowledge, 75	
13.2	A Priori Knowledge and Necessary Truths, 76	
13.3	Kripke's Puzzles, 78	
13.4	A Priori Knowledge and Gettier Cases, 80	
	Further Reading, 80	
14	Externalism	81
14.1	The Basic Idea of Externalism, 81	
14.2	Externalism and Dogmatism, 83	
14.3	Externalism and Pointless Evidence, 85	
	Further Reading, 87	
15	Internalism	89
15.1	The Basic Idea of Internalism, 89	
15.2	Internalism and Dogmatism, 90	
15.3	Internalism and Being Too Open-Minded, 91	
15.4	Internalism and Good Inquiry, 93	
	Further Reading, 95	
16	Vagueness	97
16.1	Is Knowledge Understandable? 97	
16.2	Vagueness and Knowledge, 98	
16.3	Slippery Slopes and Reliability, 99	
16.4	Generalizing the Puzzle, 101	
16.5	Escaping the Slippery Slope, 101	
	Further Reading, 103	
17	Fallibilism	105
17.1	Fallibilism and the Slippery Slope, 105	
17.2	Fallibilism and the Lottery Paradox, 108	

17.3 Rejecting Fallibilism, 109	
Further Reading, 110	
18 Infallibility Skepticism	111
18.1 The Basic Idea of Skepticism, 111	
18.2 Skepticism and Infallibility, 112	
18.3 The Evil Demon, 113	
Further Reading, 116	
19 External World Skepticism	117
19.1 The Cogito, 117	
19.2 A Less Powerful Evil Demon, 118	
19.3 The Dreaming Argument, 120	
19.4 Possible Responses to the Dreaming Argument, 122	
Further Reading, 123	
20 Inductive Skepticism	125
20.1 The Basic Idea of Inductive Skepticism, 125	
20.2 The Humean Skeptical Argument, 126	
20.3 The Humean Skeptical Solution, 128	
Further Reading, 131	
21 Rule Skepticism	133
21.1 Knowing One's Own Mind, 133	
21.2 The Basic Idea of Rule Skepticism, 134	
21.3 The Rule Skeptical Argument, 135	
Further Reading, 138	
22 Regress Skepticism	139
22.1 Inferential Justification, 139	
22.2 Epistemic Regress and Skepticism, 140	
22.3 Possible Responses to the Regress, 144	
Further Reading, 145	
23 Foundationalism	147
23.1 The Basic Idea of Foundationalism, 147	
23.2 Cartesian Foundations, 148	
23.3 Sensory Foundations, 150	
Further Reading, 153	
24 Contextualism	155
24.1 Real Doubts, 155	

24.2 Contextual Doubts, 157	
24.3 Contextualism and Epistemic Regress, 159	
Further Reading, 161	
25 Coherentism	163
25.1 The Basic Idea of Coherentism, 163	
25.2 Coherentism and Epistemic Regress, 164	
25.3 Coherentism and Consistency, 165	
25.4 Coherentism and Fiction, 166	
Further Reading, 169	
26 Pyrrhonian Skepticism	171
26.1 The Basic Idea of Pyrrhonism, 171	
26.2 The Pyrrhonist Argument, 172	
26.3 Pyrrhonism and Belief, 174	
26.4 Pyrrhonism and Intellectual Tranquillity, 175	
Further Reading, 177	
27 Questioning Epistemology	179
27.1 The Problem of the Criterion, 179	
27.2 Applying Pyrrhonism to Epistemology, 182	
27.3 Puzzling Theories, 183	
Further Reading, 184	
<i>References</i>	185
<i>About the Book and Author</i>	189
<i>Index</i>	191

1

Introducing Epistemology

If you were to think of buying a car, and you had not already learned much about cars, wouldn't it be rash not to learn something about cars before making your choice? And if you do seek knowledge about cars, but you have not already learned much about knowledge, isn't it rash not to learn something about knowledge? If so, you've come to the right place. Read on! Isn't knowledge at least as important to the world as cars are? (Although cars can transport knowledge, they cannot even exist unless there is knowledge!)

1.1 What Is Epistemology?

Epistemology is the theory of knowledge; “theory of knowledge” is what “epistemology” means. But in practice it is many theories of knowledge, there being little agreement among epistemologists as to which theories to believe. Still, there is agreement as to which theories to discuss. I will present, for your consideration, many of those theories—their claims, concepts, arguments. But I will try not to favor any one theory in particular. That would misrepresent the general state of epistemological debate. Epistemology is still a debate, not yet a set of dogmas—and your contributions are welcome. As you read, try to decide which, if any, of epistemology’s theories you accept (and why).

Naturally, the first concept for epistemologists to consider is *knowledge*. Here are two basic questions about it:

A What is knowledge?

B Is there any knowledge? (Does someone—anyone—have some?)

A zoologist might wonder “What is a cobra?” and “Are there any cobras? Where are they?” Epistemologists approach knowledge with similar care. And here are some ways in which you might respond to those basic questions.

Maybe you answer B first, and you say that there is indeed some knowledge. It would then be natural for you to provide examples of knowledge. You say that

there is knowledge. Then where is it? You might say that you know you are alive, or that you know what your cat's name is, or that you know that the moon is closer to Earth than Uranus is, or any number of other things. Then again, if you answer B with no, you might say why such examples are not really knowledge and why people are misled into thinking of them as knowledge. (This answer would make you a *skeptic*; we meet skeptical arguments later, starting in Chapter 18.)

If you do answer B with yes plus some examples, your next thought might be, "What do those examples have in common? Why are they knowledge?" (A zoologist dumps some writhing snakes at your feet: "Cobras," she says. "Oh yeah," you say composedly, ever the inquirer: "What makes them cobras?") You thus confront A. Having answered B by saying that there is such a thing as knowledge, you now seek to understand what type of thing it is. How is knowledge different from . . . well, anything else—such as a wish or a frog? What makes knowledge knowledge? (What makes a cobra a cobra?)

And with this question, you have begun doing epistemology (zoology of knowledge?). For at the heart of epistemology is question A. If you ignore A, trying to answer only B, you merely *list* examples of putative knowledge. That activity is not clearly epistemological. How can it provide a theory of knowledge? A list is not a theory. A list of cases of supposed knowledge is not an explanation, an understanding, of knowledge. And epistemology's aim is to explain and understand knowledge, to theorize about it (to good effect). By providing data to study, a list of cases of supposed knowledge can be part of an epistemological effort, a first step toward doing epistemology. But epistemology is a journey, not a step. The step of listing possible cases of knowledge is not epistemological unless it is followed by further steps, including attempts to answer A. To do epistemology is to theorize. And only question A clearly calls on you to theorize; B does not. Can you know that you are alive (and include this on your list, in response to B), without ever theorizing about what it is to have such knowledge (hence, without ever doing epistemology, in response to A)?

The choice of which question, A or B, to first try to answer presents you with what epistemologists call the *problem of the criterion*. If you think you can start by answering B, you are a *particularist*. Your first epistemological move would be to list particular cases of knowledge—and then, presumably, to try to understand them. But if you think you can start by answering A, you are a *generalist*, or *methodist*. Your first epistemological move would be to describe general ways, or methods, of knowing—and then, presumably, to try to apply them.

Why is the problem of the criterion a problem? There is no reason why it should be immediately apparent that it is a problem; not until section 27.1 do I attempt to explain why it is. In the meantime, I suggest that we begin our inquiries by adopting this particularist hypothesis:

UNO You have some knowledge.

With UNO (the book's first hypothesis), we assume for argument's sake that the answer to B is yes. Then we will try to answer A by devising an account or description of whatever knowledge we assume is covered by UNO. Finally, we should be able to return to B, with a newly acquired description of knowledge, so as to think about whether UNO should be retained after all. We might decide, "Oh, since that's what knowledge is, I don't have any after all," or, alternatively, "So that's what knowledge is like? Then many of us have lots of it! UNO is true of each of us."

To think about what knowledge is does not entail concluding that anyone definitely has it. We are asking what knowledge is like *if* there is any. We begin this book by hypothesizing, for argument's sake and to provide us with possible examples to study, that there is knowledge (specifically, that you have some). That is, we shall assume—via UNO—that there is some, and then we shall spend our time testing UNO. Will it survive? Will it fall?

1.2 Why Do Epistemology?

But wait. Why does anyone ever do epistemology in the first place? Why theorize about knowledge at all?

Well, you assume that you have knowledge: UNO (in your mouth, INO: "I have some knowledge"?) is your probable answer to B. Might that answer be wrong, though? And if it is, might you have been misled by a false theory of knowledge? Perhaps you mistakenly think that you have knowledge, because you are mistaken as to what knowledge is. Now, that sounds like a rather serious mistake to make, since UNO seems like one of the more important claims that could ever be made about you. If it might be false, and if you can ascertain this only by reflecting on exactly what knowledge is, then epistemology beckons. For epistemology is where such reflections will lead you. To reflect on the nature of knowledge is to do epistemology.

To this reasoning, though, you might respond as follows:

*Why would I treat as a mere hypothesis the proposition that I have knowledge?
That would connote hesitation on my part about accepting that proposition.
But what could be more obvious than that I have knowledge?*

Thus, the most basic of all epistemological puzzles asks why we should do epistemology in the first place (and hence why we should investigate what seems not to need investigation).

Of course, one problem with your confidence about having knowledge is that, although it seems obvious to you that you have some, it might not seem so clear to the rest of us. We might have less confidence in your abilities than you do. For example, you might think that you know a lot about baseball. We might think

that you do not. Isn't this something on which we need not take your word? ("You say you know that the world is flat? Okay, your word's good enough for me! If you say that you know, I must believe you.") Epistemology could have a point for us, if only because its application to you—you and your knowledge—can have a point. We can legitimately ask whether you know.

What would you do in the meantime, while we were thinking about whether you really do have knowledge? Could you simply ignore us and our epistemological reflections on you, blithely assuming that you have knowledge? If so, you would be answering B with a cheery "Of course," and then ignoring A. Yet mightn't this attitude be an intellectual equivalent of an ostrich's sticking his head in the sand? By answering B with yes, you say that you have knowledge. But, by ignoring A, you imply that you have no need to justify that confidence by explaining why you are a knower. How can you be so cavalier? Do you know that you know? Might it only *seem* obvious to you that you do not need to do epistemology in order to understand your claims to know? Mightn't you be wrong to assume that you have knowledge? Mightn't UNO be false?

I grant that it is "obvious" to you that you know—obvious in that you already, and naturally, accord yourself knowledge. But does the fact that it *seems* obvious to you make it true? Surely not all of your beliefs are true. How do you know that this one is? Are you ever wrong? I am willing to bet that you are. In fact, I am willing to bet that on each day of your life you have at least one false belief. See if you can prove me wrong! (And in your effort to prove me wrong, mightn't you find yourself doing epistemology as you reflect on your beliefs and on how to avoid having false ones?) Why, then, must you be right in thinking that you have knowledge in this particular case? If you are wrong sometimes, how can it be obvious that this is not one of those times? To try to answer this question, and hence to explain why you are right to claim knowledge, is to begin doing epistemology. For you would be attempting to say what knowledge is, and thus why you have some. You would be trying to answer A, hence to expand on your answer of yes to B. And you would no longer be acting like an ostrich about UNO's epistemological implications.

Furthermore, once you admit that you sometimes make mistakes in your thinking, at times ending up with false beliefs, isn't there another reason for you to pursue epistemology? You might be unable, without calling on epistemology, to *correct* your intellectual mistakes, your false beliefs. Don't you wish to decide which of your beliefs are false (so that, if possible, you can discard them)? Epistemology might help you to eliminate your intellectual mistakes by making clearer what is involved in gaining knowledge. If it can do that, it might help you to avoid ever being in the position of (1) admitting, with normal humility, that you make mistakes in your intellectual efforts, while (2) having no specific belief of which you say "That is mistaken." That is, epistemology might save you from a version of the so-called *preface paradox*.