

A detailed still life painting of a wooden window box. The box is filled with various objects: a human skull, a lit candle in a brass holder, several sheets of paper and documents, a bunch of white grapes, a red wax seal, and some fruit. The background is a dark, shadowed interior, and the window frame is made of light-colored wood. The overall composition is rich and detailed, typical of 17th-century Dutch still life painting.

EPISTEMOLOGY

A GUIDE

JOHN TURRI

WILEY Blackwell

Epistemology

Epistemology

A Guide

John Turri

WILEY Blackwell

This edition first published 2014

© 2014 John Wiley & Sons, Ltd

Registered Office

John Wiley & Sons, Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, UK

Editorial Offices

350 Main Street, Malden, MA 02148-5020, USA

9600 Garsington Road, Oxford, OX4 2DQ, UK

The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, UK

For details of our global editorial offices, for customer services, and for information about how to apply for permission to reuse the copyright material in this book please see our website at www.wiley.com/wiley-blackwell.

The right of John Turri to be identified as the author 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.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners.

The publisher is not associated with any product or vendor mentioned in this book.

Limit of Liability / Disclaimer of Warranty: While the publisher and author(s) have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. It is sold on the understanding that the publisher is not engaged in rendering professional services and neither the publisher nor the author shall be liable for damages arising herefrom. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Library of Congress Cataloging-in-Publication Data

Turri, John.

Epistemology : a guide / John Turri.

pages cm

Includes bibliographical references and index.

ISBN 978-1-4443-3369-5 (cloth) – ISBN 978-1-4443-3370-1 (pbk.) 1. Knowledge,

Theory of. I. Title.

BD143.T87 2014

121–dc23

2013016118

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

Cover image: Cornelis Norbertus Gysbrechts (fl. 1659-72), *Vanitas still life seen through a trompe l'oeil window*. © Private Collection / The Bridgeman Art Library.

Cover design by Nicki Averill Design.

Set in 10.5/13.5pt Palatino by SPi Publisher Services, Pondicherry, India

For Angelo

Contents

Preface	xi
Acknowledgments	xiii
1 The best case for skepticism about the external world? (Stroud, "The Problem of the External World")	1
2 Proving the external world exists (Or: Let's all give Moore a hand!) (Moore, "Proof of an External World")	6
3 Some ways of resisting skepticism (Moore, "Four Forms of Scepticism")	10
4 Plausibility and possibilities (Moore, "Certainty")	15
5 Skeptic on skeptic (Klein, "How a Pyrrhonian Skeptic Might Respond to Academic Skepticism")	19
6 Realism in epistemology (Williams, "Epistemological Realism")	24
7 Socratic questions and the foundation of empirical knowledge (Chisholm, "The Myth of the Given")	31
8–9 The foundation of empirical knowledge? (Sellars, "Does Empirical Knowledge Have a Foundation?" and "Epistemic Principles")	36

Contents

10	It's not a given that empirical knowledge has a foundation (BonJour, "Can Empirical Knowledge Have a Foundation?")	44
11	Interpretation, meaning and skepticism (Davidson, "A Coherence Theory of Truth and Knowledge")	49
12	Blending foundationalism and coherentism (Haack, "A Foundherentist Theory of Epistemic Justification")	54
13	Foundationalism, coherentism and supervenience (Sosa, "The Raft and the Pyramid")	60
14	Infinitism (Klein, "Human Knowledge and the Infinite Regress of Reasons")	67
15	The Gettier problem (Gettier, "Is Justified True Belief Knowledge?")	73
16	Some principles concerning knowledge and inference (Harman, <i>Thought</i> , Selections)	77
17	The essence of the Gettier problem (Zagzebski, "The Inescapability of Gettier Problems")	83
18	Knowledge is an unanalyzable mental state (Williamson, "A State of Mind")	85
19	Closure, contrast and semi-skepticism (Dretske, "Epistemic Operators")	92
20	Closure, contrast and anti-skepticism (Stine, "Skepticism, Relevant Alternatives, and Deductive Closure")	99
21	Keeping close track of knowledge (Nozick, "Knowledge and Skepticism")	103
22	Moore wins (Sosa, "How to Defeat Opposition to Moore")	111
23	The closure principle: dangers and defense (Vogel, "Are There Counter examples to the Closure Principle?")	116
24	Evidentialist epistemology (Feldman and Conee, "Evidentialism")	123
25	Non-defensive epistemology (Foley, "Skepticism and Rationality")	129
26	Reliabilism about justification (Goldman, "What Is Justified Belief?")	135

27	Reliabilism: a level assessment (Vogel, "Reliabilism Leveled")	141
28	Against externalism (BonJour, "Externalist Theories of Empirical Knowledge")	146
29	Against internalism (Goldman, "Internalism Exposed")	151
30	A skeptical take on externalism (Fumerton, "Externalism and Skepticism")	156
31	A friendly take on internalism (Feldman and Conee, "Internalism Defended")	159
32	Warrant (Plantinga, "Warrant: A First Approximation")	164
33	Intellectual virtues (Zagzebski, <i>Virtues of the Mind</i>)	169
34	Virtue epistemology (Greco, "Virtues and Vices of Virtue Epistemology")	172
35	Knowledge, luck and virtue (Pritchard, "Cognitive Responsibility and the Epistemic Virtues")	176
36	Epistemic value and cognitive achievement (Sosa, "The Place of Truth in Epistemology")	181
37	Giving up on knowledge (Kvanvig, "Why Should Inquiring Minds Want to Know?")	187
38	Giving up on (exact) truth (Elgin, "True Enough")	192
39	Naturalized epistemology advertised (Quine, "Epistemology Naturalized")	196
40	Naturalized epistemology criticized (Kim, "What is 'Naturalized Epistemology'?")	203
41	Naturalized epistemology radicalized (Antony, "Quine as Feminist")	207
42	A priori justification and unrevisability (Putnam, "There is at Least One A Priori Truth")	211
43	A priori justification and revisability (Casullo, "Revisability, Reliabilism, and A Priori Knowledge")	215
44	Philosophical method and empirical science (Bealer, "A Priori Knowledge and the Scope of Philosophy")	219
45	Experimental epistemology (Weinberg, Nichols and Stich, "Normativity and Epistemic Intuitions")	226

Contents

46	Natural kinds, intuitions and method in epistemology (Kornblith, "Investigating Knowledge Itself")	230
47	Contextualism and skeptical puzzles (DeRose, "Solving the Skeptical Problem")	235
48	Contextualism and infallibilist intuitions (Lewis, "Elusive Knowledge")	240
49	Contextualism and intuitional instability (Cohen, "Contextualist Solutions to Epistemological Problems")	244
50	Knowledge and action (Stanley, " <i>Knowledge and Practical Interests, Selections</i> ")	247
51	Rationality and action (Fantl and McGrath, "Evidence, Pragmatics, and Justification")	252
52	One invariantist's scorecard (Hawthorne, "Sensitive Moderate Invariantism")	258
53	A relativist theory of knowledge attributions (MacFarlane, "The Assessment Sensitivity of Knowledge Attributions")	264
54	Rationality and trust (Baker, "Trust and Rationality")	270
55	Testimony and gullibility (Fricker, "Against Gullibility")	273
56	Some reflections on how epistemic sources work (Burge, "Content Preservation")	277
57	Testimony and knowledge (Lackey, "Testimonial Knowledge and Transmission")	282
58	Memory and knowledge (Huemer, "The Problem of Memory Knowledge")	286
59	Perception and knowledge (McDowell, "Criteria, Defeasibility, and Knowledge")	291
60	Skills and knowledge (Reynolds, "Knowing How to Believe with Justification")	295
	Index	299

Preface

Epistemology is the philosophical discipline that studies the evaluative dimensions of cognition, their metaphysical bases, and the language we use to ascribe cognitive states. This book introduces you to a wide range of topics in epistemology, including skepticism, knowledge, justification, understanding, evidence, epistemic value, virtue epistemology, contextualism, invariantism, naturalism, testimony, perception, memory, and the a priori. It is designed and written especially to accompany the second edition of *Epistemology: An Anthology* (Blackwell Publishing, 2008), edited by Ernest Sosa, Jaegwon Kim, Jeremy Fantl, and Matthew McGrath. You can profit from this book even if you don't read it alongside the anthology, but you'll profit from it more if you do.

This book's design is modular. It consists of 60 sections corresponding to the anthology's selections. Each is intended to be brief, accessible even to the beginner, and, to the extent possible, independently intelligible. You can glean the main aim and argument of any selection in the anthology by studying the relevant section in this book. I have focused especially on providing examples and clarifying key concepts and methodological points that are essential

Preface

to the main arguments, but which the beginner couldn't reasonably be expected to be familiar with.

This book's sections can't be entirely independent, however, because the topics discussed across the selections aren't entirely independent. Indeed, many selections were chosen precisely because they complement one another. So while understanding this or that selection may be your only goal – and for some purposes this would be eminently reasonable – you'd then be missing out on some interesting and important connections running through multiple selections, both within and across the anthology's divisions. Of course, this book can't cover all the ideas, arguments, and connections in the anthology, because that would defeat the goal of brevity. It would also make it much more difficult to help you distinguish the central from the peripheral, which is crucial to achieving proficiency in any field.

Since brevity and accessibility demand selectivity, I should be clear about the general principle I've employed in trying to meet that demand. I have focused on the main theses and arguments found in the selections and sought to emphasize areas where the authors are – or, with a little imagination, easily could be – in conversation with one another. The hope is that this promotes not only brevity and accessibility but also integrity within and continuity across the various sections.

Acknowledgments

I would like to thank Ernest Sosa, Jeremy Fantl, and Matt McGrath for supporting this project. I thank Jeff Dean and Jen Bray for their support and for helping to guide the project to completion. Thanks also to Wesley Buckwalter, Trystan Goetze, and many students in epistemology courses at Huron University College and the University of Waterloo for evaluating parts of the manuscript. Thanks to Peter Blouw for work on the index. I owe my family an enormous debt of gratitude for supporting (and tolerating!) me while I worked (too) many long hours on this project, with all the sacrifice that entails: thank you Vivian, Sarah, Geno, Mom, Dad, Rich, Doug, Kelly, Mallory, and Julia. Finally, my greatest debt is to my older son, Angelo, who between the ages of 9 and 12 read every section of this book multiple times and offered valuable comments and suggestions at every stage. He is not only my son but also my greatest student and a true friend. I dedicate this book to him.

§ 1

The best case for skepticism about the external world? (Stroud, “The Problem of the External World”)

We’re all intimately familiar with what goes on in our own minds. We make plans, form opinions, experience pleasure and pain, and so on. It’s also natural to suppose that we know a lot about what goes on *outside* our own minds too, about the world around us, based on the information we get through our senses of sight, hearing, smell, taste, and touch. Natural as that idea may be, it’s surprisingly easy to get yourself into a skeptical frame of mind about the possibility of such knowledge. Can we really know anything about the world outside our own minds?

Barry Stroud aims to understand the attraction of skepticism about the external world, why knowledge of the external world based on sense experience poses a philosophical problem. To accomplish this, he focuses intensely on the argument presented at the

Stroud, Barry, “The Problem of the External World,” Chapter 1 in *The Significance of Philosophical Skepticism* (Oxford: Clarendon Press, 1984). © 1984 by Barry Stroud.

Epistemology: A Guide, First Edition. John Turri.

© 2014 John Wiley & Sons, Ltd. Published 2014 by John Wiley & Sons, Ltd.

beginning of Rene Descartes's *Meditations on First Philosophy* (published originally in 1641), the most influential work of one of the most influential philosophers of all time.

Let's note a couple important points before we proceed. First, people who reflect on knowledge nearly unanimously agree that knowledge requires truth, or as it's sometimes put, that knowledge is factive.¹ This means that we can know something only if it is true or a fact. We cannot know a falsehood. (We can of course *believe* a falsehood, but that's a different matter.) Now if you think that knowledge is not factive, I recommend a simple solution: everywhere we here speak of "knowledge," understand it to mean "knowledge of the truth," and every time we claim or ask whether someone "knows that so-and-so," understand it to mean "*knows it's true* that such-and-such." Second, something can be possible without being real or actual. Indeed lots of things are possible that aren't actual. For instance, it's possible for winged horses to exist, even though none actually do. Likewise for wizards, dragons, phlogiston, the luminiferous ether, etc. With those points in mind, let's proceed.

Imagine Descartes at work in his study on a cold night, sitting a few feet from a comforting fire. Unsure for the moment how his narrative should best proceed, he takes a break and turns his attention to the fire. He sees its colorful flames flitting and flickering; he hears it crackling and popping; he feels its heat emanating; he smells the fragrant wood burning. In light of all this, Descartes of course believes he's near a fire. But do these sense experiences enable him to know he's near a fire?

It's hard to imagine Descartes's senses putting him in a better position to gain knowledge of the external world. He is as well-positioned as any of us could ever hope to be. So if the answer to our question at the end of the previous paragraph is "No," then it seems very likely that we never know anything about the external world, at least by way of our senses.

¹ For some recent controversy over the "truth requirement" on knowledge, see Allan Hazlett, "Factive Presupposition and the Truth Condition on Knowledge," *Acta Analytica* 27.4 (2012): 461–478, and John Turri, "Mythology of the Factive," *Logos & Episteme* 2.1 (2011): 143–152.

The way things look, sound, smell, and feel make it appear to Descartes as though he's near a fire, and it's this appearance that he trusts when judging that he's near a fire. But of course things might appear *exactly* the same in a perfectly realistic dream. And a perfectly realistic dream is a genuine possibility. It's certainly possible for him to have all those sensations despite merely dreaming that he's near a fire. Indeed, any sensory experience might be a mere component of a perfectly realistic dream. Thus sense experience, being equally compatible with dreaming or waking, could never enable him to know that he is awake rather than merely dreaming.

As Descartes recognizes, if he's merely dreaming that he's near a fire, then he certainly doesn't know that he's near a fire. And he also recognizes it is at least possible that he's merely dreaming. So he knows that a certain genuine possibility, the *dream-possibility* (as Stroud calls it), is incompatible with his knowing that he's near a fire. So in order to know that he's near a fire, he must know that the dream-possibility is false.

Notice that, on this way of thinking, in order for the dream-possibility to potentially threaten Descartes's knowledge of the fire, he doesn't need to know, or even so much as believe, that it is actually true. No, the dream-possibility threatens simply because Descartes recognizes that it is possibly true, and that if it were actually true, he wouldn't know that he's near a fire.

Could Descartes ever come to know that the dream-possibility is false? Sense experience itself won't enable such knowledge because, as we've already said, any sense experience is perfectly compatible with the dream-possibility. But isn't there some test he could perform to determine whether he is merely dreaming? Unfortunately not, because in order for him to learn from the test, he'd need to know that he wasn't merely dreaming that he was performing the test!

If you're wondering why he couldn't then just perform a second test to determine whether he's merely dreaming that he performed the first test, consider: he could equally well be dreaming that he's performing the second test. The same is true for a third test he might perform to determine whether he's merely dreaming that he performed the second test. And so on. No matter how many tests he

performs, the same problem recurs. And since it's not possible to perform an infinite series of tests, we find no relief in this direction.

Let's encapsulate the preceding line of thought in the following argument, broken up into two parts to enhance clarity. The main argument goes like this:

1. If Descartes doesn't know that he's near a fire, then we never know anything about the external world. (Premise)
2. Descartes doesn't know that he's near a fire. (Premise)
3. So we never know anything about the external world. (From 1 and 2)

The argument is logically valid: if its premises are true, then its conclusion must be true too. That leaves us to ask whether its premises are true. 1 is at least very plausible, and Stroud is willing to grant it. That leaves only 2 to seriously question. The following supplementary argument supports 2:

- a. Descartes knows that the dream-possibility is incompatible with his knowing that he's near a fire. (Premise)
- b. If Descartes knows that a possibility is incompatible with his knowing some specific claim, then in order for him to know the specific claim, he must know that the possibility in question is false. (Premise)
- c. So in order for Descartes to know that he's near a fire, he must know that the dream-possibility is false. (From a and b)
- d. But Descartes couldn't know that the dream-possibility is false. (Premise)
- e. So Descartes doesn't know that he's near a fire. (From c and d)

Notice that (e) is exactly the same as 2.

Should we accept this argument? Stroud wonders whether we can seriously entertain the skeptical conclusion expressed by 3, because it's allegedly either absurd or even unintelligible. But merely rejecting it as absurd or unintelligible deprives us of the opportunity to learn something potentially important about

knowledge (or at least about our concept of knowledge). Accordingly, he challenges those of us inclined to reject the conclusion to *locate the argument's flaw*. Whatever it is, it isn't obvious.

Stroud suggests that (c) is false. Yet (c) follows from (a) and (b), so rejecting (c) requires us to reject at least one of (a) and (b). (a) is obviously true, which leaves (b).

The problem is that (b) is arguably "embodied" in our ordinary procedures for "making and assessing knowledge-claims." Consider for instance a bird watcher who judges a certain bird to be a goldfinch. We ask her why she thinks it's a goldfinch. "Because it's yellow," she says. "But for all you've said," we respond, "it's possible that it's a canary – canaries are yellow too." We don't think she knows it's a goldfinch, because she knows very well that canaries aren't goldfinches, and yet she doesn't know it's not a canary. She must rule out this relevant possibility, *the canary-possibility*, in order to know it's a goldfinch.

The question then becomes whether the dream-possibility is in all relevant respects similar to the canary-possibility, so that when we insist that the bird watcher must rule out the canary-possibility, we thereby commit ourselves to insisting that Descartes must rule out the dream-possibility. Does Descartes have to rule out the dream-possibility in order to know there's a fire nearby, as the bird watcher must rule out the canary-possibility in order to know that she's looking at a goldfinch? If not, why not? Each subject knows the possibility in question is incompatible with his or her knowing the claim in question. So what *could* be the difference?

A plausible explanation of the difference, should there be any, would go a long way toward resolving "the problem of the external world." Therein lies the challenge, and potential reward, of confronting philosophical skepticism.

References

- Allan Hazlett, "Factive Presupposition and the Truth Condition on Knowledge," *Acta Analytica* 27.4 (2012): 461–478.
- John Turri, "Mythology of the Factive," *Logos & Episteme* 2.1 (2011): 143–152.

§ 2

Proving the external world exists (Or: Let's all give Moore a hand!) (Moore, "Proof of an External World")

Suppose we disagree about the number of books on the desk. You say there are at least two. I disagree. And it's no mere verbal disagreement – we're referring to the same desk, and mean the same thing by "book" and "at least two," etc. How might you prove your point?

Here's one way. You walk over, point to one book sitting on the desk, and then point to another, all while saying, "Here's one book on the desk, and here's another. So there are at least two books on the desk." I couldn't rightly criticize the proof. I'd have to concede the point. What else could I possibly be looking for in a proof? Your premises ("here's one book the desk, and here's another") are different from your conclusion ("there are at least two books on the

Moore, G. E., "Proof of an External World," extracted from pp. 147–70 in Thomas Baldwin (ed.), *G. E. Moore: Selected Writings* (London & New York: Routledge, 1993). © 1993 by Thomas Baldwin.

Epistemology: A Guide, First Edition. John Turri.

© 2014 John Wiley & Sons, Ltd. Published 2014 by John Wiley & Sons, Ltd.

desk”), in which case you didn’t simply *beg the question*. Your conclusion follows straightforwardly from your premises, and you know that it does. And you obviously know the premises – after all, you aren’t blind, you’re looking right at the books, and you’ve correctly verbally identified them. Without question, your proof perfectly settles the matter in your favor.

Immanuel Kant, perhaps the most influential of all modern philosophers, once said it was “a scandal to philosophy” that the existence of real, mind-independent external objects “must be accepted merely on *faith*” rather than a “satisfactory proof.”¹ Kant thought he had rescued philosophy from scandal by giving such a proof, indeed, the *only* possible such proof. Setting aside the merits of Kant’s own proof, G.E. Moore denied that Kant’s was the only possible such proof. A much simpler and fully convincing proof is readily available.

Moore offers his proof by saying, “Here’s one hand, and here’s another. So external objects exist,” as he gestures and holds his hands up before us. This proof, Moore says, is “perfectly rigorous.” It meets the three criteria we noted earlier when discussing your proof about the number of books on the desk. Its premises are different from its conclusion; its conclusion follows, as Moore knows, straightforwardly from its premises; and finally, Moore obviously knows the premises.

Might a satisfactory proof require more than meeting those three criteria? Not if our ordinary practice is any indication. As with your earlier proof about the books, we “constantly take proofs of this sort as absolutely conclusive.”

Note an interesting connection with Stroud’s discussion from §1. We wondered whether Descartes was right to claim that he must know the dream-possibility is false in order to know that he’s near a fire. And Stroud worried that Descartes was indeed right about that, because such a requirement might be “nothing more than an instance of a general procedure we recognize and insist on in making and assessing knowledge-claims in everyday” life, which procedure helps to define

¹ Quoted by G.E. Moore at the beginning of his “Proof of an External World.” The quote is alluded to but does not explicitly appear in the excerpt included in the anthology.

out concept of knowledge. Here Moore appeals to our everyday procedures for offering and evaluating *proofs*.

Moore anticipates that some will say his so-called proof fails. One type of critic insists that in order for Moore's proof to really succeed, he must also prove his premises – prove that *here is one hand* and that *here is another*. If this critic is right, then the three criteria we earlier identified aren't sufficient for a conclusive proof after all. At least sometimes, a conclusive proof requires more. Moore rejects this, and explicitly disavows any intention to prove his premises. He doubts it could be done, because proving them requires proving that he's not merely dreaming that he has hands. And even though he has "conclusive reasons" (or "conclusive evidence") that he's not merely dreaming, he cannot articulate that evidence to us, which he of course must do in order to offer a proof.

One is reminded of a scene in Robert Louis Stevenson's *The Strange Case of Dr. Jekyll and Mr. Hyde*. A shaken and ill Jekyll tells his friend Utterson that Mr. Hyde "will never more be heard of." When Utterson suggests – rightly, it turns out, as the story subsequently unfolds – that Jekyll's assertion might not be entirely warranted, Jekyll replies, "I have grounds for certainty that I cannot share with anyone."² What does Jekyll mean by "cannot" here? He might mean that he cannot prudently share it with anyone. Revealing his relation to Hyde would be disastrous to Jekyll personally, so prudence forbids it. And, indeed, this is likely the first interpretation to occur to the reader. But great writers imbue their work with many layers of meaning, and in light of Moore's discussion, one wonders whether there's more to Jekyll's last claim than first meets the eye. Perhaps Jekyll (also) meant that he is simply incapable of sharing at least some of his reasons. This is made all the more plausible later when Jekyll writes that his transformational experiences were marked by "indescribably new" sensations.³ But this shouldn't be

² In the Section "Incident of the Letter."

³ In the section "Henry Jekyll's Full Statement of the Case."

all that surprising; at times all of us have experiences that are "beyond words."

Another type of critic claims that Moore fails to meet one of the everyday criteria for successful proofs: Moore doesn't know that his premises are true. But why think that? Is it not, as Moore says, simply absurd to suggest that he didn't know he was gesturing toward his hands as he performed his proof? The critic thinks not. To know something, she claims, requires that you be able to prove it – no knowledge without proof, as it were. And as Moore himself admits, the critic continues, Moore cannot prove his premises, so he doesn't know them. But Moore rejects this as well, claiming, "I can know things, which I cannot prove."

Have you ever known something that you wanted to convince someone of, but found yourself saying, "If only I could prove it," or more despairingly, "But I just can't prove it!"? In the episode "Identity Crisis" of the television show *CSI*, Gil Grissom and Catherine Willows are confounded by a computer fingerprint analysis which fails to confirm that Judge Douglas Mason is in fact Grissom's nemesis Paul Millander, the long-sought serial killer. Grissom says to Catherine, "I don't care what the computer says – that guy is Paul Millander." Catherine replies, "Yeah, we know that. How do we prove it?" They of course go on to prove it. But that's not the important point. You can imagine them finding out that they simply cannot prove it. Indeed for a moment viewers are left to think as much when they learn that a repository storing relevant vital records burned down long ago. Yet this doesn't lead us to think, as we watch the episode, that Grissom and Catherine don't really know that Mason is Millander. Through a combination of memory and visual recognition, Grissom knew Mason was Millander the moment he set eyes on him in Mason's courtroom. Catherine learned it based on Grissom's testimony. But at that point in the story they weren't able to prove it. At least, that's how the screenwriters portrayed it, and viewers didn't detect any incoherence in the plot or dialog. This suggests that Moore is right when he says that knowledge does not require proof.

§ 3

Some ways of resisting skepticism (Moore, “Four Forms of Scepticism”)

Recall the *dream-possibility* from §1: it is genuinely possible for all your sense experiences to be part of a perfectly realistic dream, rather than accurate portrayals of the world around you. The dream-possibility featured centrally in Stroud’s understanding of Descartes’s skeptical reasoning. Descartes also discussed another skeptical possibility: the *demon-possibility*. You think your sensations are caused by ordinary objects in the world around you. But the demon-possibility says it’s genuinely possible for all your sensations to be part of an elaborate deception created by a supremely powerful evil demon.

The demon-possibility impressed Bertrand Russell, one of G.E. Moore’s friends and perhaps the most famous Anglo-American philosopher of the twentieth century. This section focuses on Moore’s response to Russell’s discussion of the demon-possibility.

Moore, G. E., “Four Forms of Scepticism,” pp. 220–2 in *Philosophical Papers* (New York: Collier Books, 1962). © 1962 by G. E. Moore.

Epistemology: A Guide, First Edition. John Turri.

© 2014 John Wiley & Sons, Ltd. Published 2014 by John Wiley & Sons, Ltd.

Russell worried that we can't know that the demon-possibility is false. Moore says we can and do know it's false.

Russell and Moore framed the demon-possibility differently from how I did just a moment ago. They speak of a "malicious demon" producing "percepts" in me, in which case these percepts are not produced by the sort of external objects we think normally surround us. But this is a mere verbal difference. We may understand *percepts* to be nothing other than the *sense experiences* or *sensations* we've discussed so far, such as the visual experience of a flame flickering, the sound of a fire crackling, the feeling of heat emanating from the fire, the smell of wood burning, and so on.

After some preliminary work to decide how best to understand what Russell means by "logical possibility," Moore ultimately settles on the following as the most charitable interpretation of Russell's argument. (The numbering used here does not correspond to Moore's numbering.)

1. If the truth of the demon-possibility is compatible with everything we know immediately, then we can't know for certain that the demon-possibility is false. (Premise)
2. The truth of the demon-possibility is compatible with everything we know immediately. (Premise)
3. So we cannot know for certain that the demon-possibility is false. (From 1 and 2)

Notice the qualifier "for certain" after "know." Sometimes Moore drops the qualification. But usually he doesn't. Up till now we've talked about *knowledge*, not *certain knowledge*. I leave it up to you to consider carefully whether there are relevant differences between *knows* and *knows for certain*, and how it might affect the debate.

Before evaluating the argument, let's clarify what we mean by "knowing *immediately*." Some of our beliefs are based on reasoning. Some aren't. Of those that aren't, perhaps some have no basis at all: we *just believe* them. But not all our beliefs are like that. Some that aren't based on reasoning are still based on *something*. For example, your belief that there's a text in front of you is based *on sight*.

Compare that to my belief that my son must be home. I don't see him. But I see his jacket hanging in the closet. And I know that he wouldn't leave home without his jacket in this weather. From this I infer that he's home. Call a belief based on reasoning *inferential*. Call a belief based on nothing *baseless*. Call a belief that is neither baseless nor inferential *immediate*. Likewise, call knowledge based on nothing (if such a thing is possible) *baseless*. Call knowledge based on reasoning *inferential*. And call knowledge that is neither baseless nor inferential *immediate*.

Now we can evaluate the argument. The first thing to note is that it's logically valid: if its premises are true, then its conclusion must be true too. So if Moore rejects the conclusion, which he does, he should reject at least one of the premises.

Moore accepts 2. He thinks we know immediately things about our own minds, but *not* things about the external world. I might know that this is a pencil, or that I have hands, or that Michael Jackson died in June 2009, or that the evil demon is not deceiving me. But I don't know any of these things immediately. So assuming that our earlier threefold distinction among baseless, inferential, and immediate knowledge exhausts the different ways we could know things, it follows that our knowledge of the external world must be either baseless or inferential. And Moore indeed agrees that it's inferential. (He doesn't seem to have seriously considered in this context whether it might be baseless; see §20 for a view that advocates baselessness for some knowledge.)

We find many different types of inference or argument. (I won't distinguish between inferences and arguments.) But for our purposes we can divide good arguments into two main types: deductive and nondeductive. Presumably, only good inferences produce knowledge. A good deductive argument is *logically valid*. We've seen this term used already several times in this book. A logically valid argument has the following important property: if its premises are true, then its conclusion must be true too. Put otherwise, the truth of its premises absolutely guarantees the truth of its conclusion. A good nondeductive argument is *logically strong*. We've not yet seen this term used. A logically strong argument is such that if its premises are

true, then its conclusion is at least probably true. Put otherwise, the truth of its premises makes its conclusion at least likely. A good deductive argument is not the least bit risky. But a good nondeductive argument is still risky.

Does Moore think our inferential knowledge of the external world is deductive or nondeductive? Nondeductive. He says that such knowledge is based on “analogical or inductive arguments.” Analogical and inductive arguments are nondeductive.

How do these arguments about the external world go? Their premises are things that we know immediately about our own mind, for example, premises about what sort of sensory experiences we seem to be having. Their conclusions are the things we know about the external world. The truth of the premises doesn’t absolutely guarantee the truth of the conclusion, or else we would have a good *deductive* argument, not merely a good nondeductive argument. (Think carefully about why that’s so.)

This brings us to the heart of the disagreement between Russell and Moore. Russell thought that only good deductive inferences could enable inferential knowledge that is also certain. Moore thought that at least some nondeductive inferences could do so as well. That’s one main reason why Russell accepted 1, whereas Moore rejected it.

Moore had at least one thing to say in his defense on this point. Moore and Russell both might agree, for example, that

4. We know for certain that this is a pencil *only if* nondeductive inference enables certain knowledge.

With that in mind, we must ask ourselves: what’s more plausible, that we *do* know for certain that this is a pencil, or that nondeductive inference *doesn’t* enable certain knowledge? Moore chooses the former. He says it’s much more plausible that we do know for certain that this is a pencil. So, he reasons, it’s rational to conclude that nondeductive inference enables certain knowledge.

Moore indicates that he’s willing to reason similarly about other things too. For instance, if you convinced him that nondeductive

inference couldn't enable certain knowledge after all, then he'd reject 4. He'd conclude that either (i) we do after all have a good *deductive* argument from what we know immediately about our minds to conclusions about the external world, or (ii) we can after all know things about the external world *immediately*.

More generally, suppose we give Moore a choice between two things:

- A. The claim that we do have certain knowledge about the external world.
- B. Any theoretical claim about what knowledge requires, which implies we that *don't* have certain knowledge of the external world.

He'll always say A is more plausible and reasonable. We'll see more of this general argumentative strategy from Moore in the next section.