
Epistemology

PRINCETON FOUNDATIONS OF CONTEMPORARY PHILOSOPHY

Scott Soames, *Series Editor*

Philosophical Logic by JOHN P. BURGESS

Philosophy of Language by SCOTT SOAMES

Philosophy of Law by ANDREI MARMOR

Truth by ALEXIS G. BURGESS and JOHN P. BURGESS

Philosophy of Physics: Space and Time by TIM MAUDLIN

Philosophy of Biology by PETER GODFREY-SMITH

Epistemology by ERNEST SOSA

EPISTEMOLOGY

Ernest Sosa

PRINCETON UNIVERSITY PRESS
PRINCETON AND OXFORD

Copyright © 2017 by Princeton University Press
Published by Princeton University Press, 41 William Street,
Princeton, New Jersey 08540
In the United Kingdom: Princeton University Press,
6 Oxford Street, Woodstock, Oxfordshire OX20 1TR

press.princeton.edu

All Rights Reserved

ISBN 978-0-691-13749-0

Library of Congress Control Number: 2016950917

British Library Cataloging-in-Publication Data is available

This book has been composed in Minion Pro and Archer

Printed on acid-free paper. ∞

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Contents

| | |
|---|-----|
| <i>Preface for Readers and Instructors</i> | vii |
| <i>Acknowledgments</i> | xi |
| CHAPTER ONE | |
| Descartes's Pyrrhonian Virtue Epistemology | 1 |
| CHAPTER TWO | |
| Dream Skepticism | 21 |
| CHAPTER THREE | |
| Regress Skepticism | 39 |
| CHAPTER FOUR | |
| Knowledge: What It Is and How We Might Have It | 55 |
| CHAPTER FIVE | |
| Knowledge as Action | 71 |
| CHAPTER SIX | |
| Varieties and Levels of Knowledge | 87 |
| CHAPTER SEVEN | |
| The Value of Human Knowledge | 106 |
| CHAPTER EIGHT | |
| Mind-World Relations: Action, Perception, Knowledge | 120 |
| CHAPTER NINE | |
| Two Forms of Virtue Epistemology | 140 |
| CHAPTER TEN | |
| Knowledge, Time, and Negligence | 157 |
| CHAPTER ELEVEN | |
| Virtue Theory against Situationism | 171 |
| CHAPTER TWELVE | |
| Virtue Epistemology and a Theory of Competence | 191 |

Contents

| | |
|-----------------------------|-----|
| CHAPTER THIRTEEN | |
| Knowledge and Justification | 207 |
| <i>Further Reading</i> | 223 |
| <i>Index</i> | 227 |

Preface for Readers and Instructors

A FAMILIAR STORY line for modern and contemporary epistemology gives pride of place mainly to perception, to inductive reason, and to a problem of the external world that is said to derive from Descartes through his epistemological writings and their influence on subsequent Western epistemology. This story line concerns the foundational evidence provided by sensory experience and describes how inductive reasoning, either enumerative or hypothetical, can lead us beyond the skeptic to a full view of the world around us in both its historical length and its scientific depth.¹

A subplot then welcomes reliable processes of belief formation and sustainment alongside logical inference. These reliable sources are said to include memory and testimony, along with perception. In the expanded picture, these three sources are given standing through their own reliability, with no need to reduce them to any inference based on the subjective given and its phenomenology. We are thus said to attain a broader and more plausible picture of human knowledge and its attainment.

Several available textbooks already lay out that account thoroughly and readably so as to provide an engaging introduction to epistemology so conceived. These provide a fine introduction to modern and contemporary epistemology, including its historical development, and important recent work.

As I argue in the first chapter, that story line misrepresents the father of modern epistemology even in its main plot. The true Cartesian problematic instead goes back to Pyrrhonian epistemology and focuses on skeptical issues more general than that of the external world, that of other minds, that of the past through memory, that of the future through induction, or that of the world's

¹Such a problem of empirical knowledge arises whether on an internally mental foundation or, going now beyond Descartes, on a perceptual external basis. Such foundationalism takes a similar approach to other standard problems of recent and contemporary epistemology: the epistemic problem of other minds, for instance, or that of memory or of testimony.

depth and general structure through inference to the best explanation. These are all important problems, and Descartes did tackle some of them as he came to grips with the more general epistemological problematic of main interest to him, one that also goes back to the schools of ancient Greece.

The first chapter lays out the case for understanding Descartes in the alternative way proposed. The next two chapters step back from such historical detail to a more general understanding of the skeptical problematic found within the Pyrrhonian or Cartesian framework.

More fully understanding that problematic requires closer attention to what sort of belief and what sort of knowledge are mainly of interest to that older tradition. I explore those issues in four central chapters, the fourth through the seventh. These lay out the core of the virtue-theoretic approach, which provides an answer to the *Theaetetus* query as to the nature of knowledge and also to the *Meno* problem as to its distinctive value. The key elements of that account are the *apt belief*, whose correctness manifests the competence of the believer, and the *fully apt belief*, which is guided to aptness by the believer's assessment of risk through a competent view of his or her own cognitive skill, shape, and situation.

In the eighth chapter that theoretical structure is seen to bear not only on human knowledge but also on perception and on action, so that all three mind-world relations are illuminated thereby.

With chapter 9 I return to epistemology more narrowly construed. First I take up a received division of virtue epistemology into two camps: one focused on reliability and competence and one on responsibility and character. I argue in this chapter that the dichotomy is false and misleading. A true epistemology will appropriately combine all four elements, and indeed will do so even in its approach to the most traditional issues of skepticism and of the nature and value of knowledge.

Chapter 10 takes up a problem from Saul Kripke, his dogmatism paradox. It concludes that, in order to address properly that fascinating puzzle, we need a notion of epistemic negligence that fits naturally within a virtue-theoretic framework like that developed in the earlier chapters.

The succeeding two chapters—the eleventh and twelfth—take up a challenge to virtue theory generally, one based on shocking but robust results obtained by social epistemologists in experiments that prompt a situationist attack on virtue theory. Chapter 11 examines those results and the arguments based on them and concludes that, though the challenge is interesting and important, it can be met with success. Chapter 12 then develops an account of competence, and of *epistemic* competence more specifically, in line with our general defense against the situationist attack.

The first twelve chapters thus focus on the nature, value, and extent of human knowledge and on the epistemic virtue or competence that constitutes such knowledge. However, some give pride of place in epistemology to justification rather than knowledge. How, then, is our virtue-theoretic approach to accommodate such justification? A problem is alleged here through a New Evil Demon argument that repurposes the celebrated Cartesian thought experiment. The conclusion is now that the Cartesian scenario's victim can be about as well justified as a normally situated believer, though this is impossible to explain by an externalist epistemology. Our final chapter, the thirteenth, offers a way out for the externalist virtue epistemologist, with implications for the perennial problematic of radical skepticism.

Although the approach to epistemology in this book differs from the familiar approaches of the past several decades, I return to the main epistemological tradition stretching from the Greeks to a German tradition that includes the epistemological writings of Hegel.² The main modern protagonist in this narrative is the philosopher often recognized as the founder of modern philosophy and of modern epistemology more specifically. Descartes is that, to be sure, but this must not be allowed to obscure the true content of his philosophy of knowledge, with its ancient sources. The present text tries to illuminate that older tradition and to develop a contemporary account with the same structure while giving it substance, not with theology, as does Descartes, but with science and common sense.

²See Michael N. Forster, *Hegel and Skepticism* (Cambridge, MA: Harvard University Press, 1989).

Acknowledgments

IN PURSUIT OF the objectives laid out in the preface, this textbook expounds a virtue epistemology developed over several decades. The main objective here is to make this view maximally accessible to the uninitiated. Drawing on earlier scattered writings, in the present text I have reformulated and reassembled the main pieces to that end. I try to present the view in a unified way accessible to advanced undergraduates and graduate students in a semester course. In pursuit of that objective, I have drawn selectively, from diverse earlier writings, formulations already as simple and lucid as I can manage, and supplemented these with new material aimed at interconnecting and unifying the whole.

I have received very helpful comments on various parts of the book from Rutgers graduate students, visitors, and postdocs in a reading course, in a seminar, and in an epistemology workshop. These have included Sara Aronowitz, Sarit Barzilai, Eric Bayruns, Robert Beddor, David Black, Rodrigo Borges, Laura Callahan, Eddy Chen, Megan Feeney, Will Fleisher, Daniel Forman, Bryan Frances, Simon Goldstein, Veronica Gomez, Jimmy Goodrich, Igal Kvart, Lisa Miracchi, Andrew Moon, Daniel Rubio, and Emily Sullivan.

Austin Baker, David Sosa, and Kurt Sylvan gave me written comments on the whole manuscript, as did Georgi Gardiner, who also gave me valuable help with the reading lists, both chapter by chapter and at the end of the book. I am grateful also to the superb PUP staff whose help was so helpful: Rob Tempio, Jenny Wolkowicki, and Marilyn Martin.

The virtue epistemology presented here was initially developed in some of my earlier books: one with Cambridge University Press: *Knowledge in Perspective* (1991); three with Oxford University Press: *A Virtue Epistemology* (2007), *Reflective Knowledge* (2009), and *Judgment and Agency* (2015); and one with Princeton University Press: *Knowing Full Well* (2011).

With permission in each case, this text draws (in part) from my previous work in the ways detailed below. The material drawn is then intermixed with and supplemented by additional new

Acknowledgments

material. The objective is a book that can serve as an advanced introduction, perhaps in combination with Descartes's *Meditations*, with supplementary readings, with an advanced standard introductory text, or with a collection of debates in epistemology. A very few further readings are suggested at the end of each chapter, providing other treatments of the issues in that chapter, mostly in close relation to my virtue-theoretic focus. Instructors might supplement a selection of chapters from this text with other materials, as suggested earlier, if they opt to provide a broader array of epistemological topics and options.

Chapter 1: "Descartes's Epistemology," in *Scepticism and Perceptual Justification*, ed. Dodd and Zardini, 13–33 (Oxford, UK: Oxford University Press, 2014).

Chapter 2: "Dreams and Skepticism," 2004 Presidential Address to the Eastern Division of the American Philosophical Association, in *Proceedings and Addresses of the American Philosophical Association* 79, no. 2 (November 2005): 7–18.

Chapter 3: Drawn in part from my contributions to symposia on my work in *Critica* 42 (2010): 77–93, *International Journal for the Study of Skepticism* 1 (2011): 138–49, *Philosophical Papers* 40 (2011): 341–58, *Synthese* 188 (2011): 309–21, and *Philosophical Studies* 153 (2011): 43–59.

Chapter 4: Drawn in part from "How to Defeat Opposition to Moore," *Philosophical Perspectives* 13 (1999): 141–53, and in part from "The Metaphysical Gettier Problem, and the X-Phi Critique," in *Explaining Knowledge: New Essays on the Gettier Problem*, ed. R. Borges, C. de Almeida, and P. Klein (Oxford, UK: Oxford University Press, forthcoming).

Chapters 5, 6, and 12: *Judgment and Agency* (Oxford, UK: Oxford University Press, 2015), chapters 3 and 4.

Chapter 7: "Knowing Full Well: The Normativity of Beliefs as Performances," *Philosophical Studies* 142 (2009): 5–15.

Chapter 8: "Mind-World Relations," *Episteme* 12 (2015): 155–66. The issue contains the proceedings of a conference to celebrate the tenth anniversary of the journal.

Chapter 9: "Virtue Epistemology: Character versus Competence," in *Current Controversies in Virtue Theory*, ed. Mark Alfano, 62–74 (London: Routledge, 2015).

Chapter 10: "Knowledge and Time: Kripke's Dogmatism Paradox and the Ethics of Belief," in *The Ethics of Belief*, ed. Jonathan Matheson and Rico Vitz, 77–89 (Oxford, UK: Oxford University Press, 2014).

Chapter 11: "Situations against Virtues: The Situationist Attack on Virtue Theory," in *Philosophy of the Social Sciences: Philosophical Theory and Scientific Practice*, ed. Chrysostomos Mantzavinos, 274–91 (Cambridge, UK: Cambridge University Press, 2009).

Chapter 12: "Virtue Theory against Situationism," in *Epistemic Situationism*, ed. Mark Alfano and Abrol Fairweather (Oxford, UK: Oxford University Press, forthcoming).

Chapter 13: "Knowledge, Competence, and Justification." This chapter, whose content is previously unpublished, aims to find a place for epistemic justification within virtue theory and offers a further response to radical skepticism on that basis.

Epistemology

Descartes's Pyrrhonian Virtue Epistemology

A. THE METHOD OF DOUBT AND ITS OBJECTIVES

What is Descartes up to in the *Meditations*? On one level at least, he is *not* engaged in a project of determining what he should believe, what it would be reasonable for him to believe. Consider, for example, the following two passages:

[When] it is a question of organizing our life, it would, of course, be foolish not to trust the senses, and the skeptics who neglected human affairs to the point where friends had to stop them falling off precipices deserved to be laughed at. Hence I pointed out in one passage that no sane person ever seriously doubts such things. But when our inquiry concerns what can be known with complete certainty by the human intellect, it is quite unreasonable to refuse to reject these things in all seriousness as doubtful and even as false; the purpose here is to come to recognize that certain other things which cannot be rejected in this way are thereby more certain and in reality better known to us.¹

My habitual opinions keep coming back, and, despite my wishes, they capture my belief, which is as it were bound over to them as a result of long occupation and the law of custom. I shall never get out of the habit of confidently assenting to these opinions, so long as I suppose them to be what in fact they are, namely highly probable opinions—opinions

¹ *Replies to the Fifth Set of Objections*, Cottingham, Stoothoff, and Murdoch [CSM], eds., *The Philosophical Writings of Descartes* (3 volumes), II: 243 (Cambridge University Press, 1984); emphasis added.

which, despite the fact that they are in a sense doubtful, as has just been shown, it is still much more reasonable to believe than to deny. In view of this, I think it will be a good plan to turn my will in completely the opposite direction and deceive myself, by pretending for a time that these former opinions are utterly false and imaginary. (First Meditation, CSM II: 15; emphasis added)

If we take Descartes at his word, then, no one sane *ever* seriously doubts his habitual opinions, which are *much* more reasonably believed than denied.

What else might be involved in the Cartesian method of radical doubt beyond *pretending* our customary opinions to be doubtful and even false? Let us examine the method more closely. Here, first, is a crucial passage:

[Those] who have never philosophized correctly have various opinions in their minds which they have begun to store up since childhood, and which they therefore have reason to believe may in many cases be false. They then attempt to separate the false beliefs from the others, so as to prevent their contaminating the rest and making the whole lot uncertain. Now the best way they can accomplish this is to reject all their beliefs together in one go, as if they were all uncertain and false. They can then go over each belief in turn and re-adopt only those which they recognize to be true and indubitable. Thus I was right to begin by rejecting all my beliefs. (*Replies to the Seventh Set of Objections*, CSM II: 324)

Near this passage Descartes invokes the famous apple-basket metaphor. Upon discovering some rot in one's basket, what is one to do? His answer: dump out all the apples and readmit only those that pass inspection. Only thus can we be sure that no rot will continue to spread undetected.

The apples are beliefs or opinions, among them old familiar ones, stored since childhood. Once the beliefs in our basket are found to contain the rot of error, we are to dump them *all*. But

how do we understand this metaphor? What is it to “remove” a belief from the basket? What is it to *reject* a belief?

In a familiar view, to reject a belief is to give it up, to withhold or suspend judgment on its content.² In any case, it would be to replace believing with *not*-believing. The ground by the empty basket would then be free of *believings*, because in the present view to dump a believing is to destroy it. Strewn there would be found *believables*, contents earlier believed. Concerning all of those contents, the subject would now be suspending judgment. That is the view.

That view is highly problematic for several reasons. For one thing, rejecting all our beliefs that way would entail believing nothing, replacing belief universally with unbelief. What would that require? Could one bring up each content separately, replacing acceptance of it with suspension? Surely not. The contents would need to be handled in manageable clusters, for collective suspension in one fell swoop.

Suppose, accordingly, that we identify the beliefs in question indirectly, as for example “beliefs I hold” or “old and customary opinions learned since childhood.” If we pick them out only so generally, however, no mental operation available to us would seem to result in the desired universal suspension. It is doubtful that we can suspend judgment *de re* on *each* content thus picked out—just as, say, “long-held opinions”—simply by taking it *de dicto* that these contents are all doubtful or supposing that they are all false.

And there is a further reason why Descartes’s “rejection” cannot plausibly amount to suspension or withholding. Recall how the process is supposed to go. The beliefs dumped out of the basket must undergo inspection. Only those that pass will be re-admitted. But the relevant inspection will have to involve some

²Of course one might also reject a particular belief not by suspending judgment but by believing the opposite. But this could not be a sensible way for Descartes to reject *all* of his beliefs, since it would be to acquire an opposing, equally numerous set of beliefs, one that would be even more doubtful than the rejected set. Far more sensible would be the less committal sort of “rejection” that requires only suspension. But we shall find even this to be highly problematic.

process of reasoning. It is through such reasoning that we will determine whether a certain condition is satisfied, which will earn that belief readmission. And how could we possibly perform any such reasoning while deprived of beliefs? Note well: the reasoning in question cannot be just conditional. The desired conclusion is that the belief under examination passes inspection. Thus would we obtain the assertoric basis for a practical syllogism that warrants readmission. An asserted conclusion (an asserted categorical, not something unasserted or conditional) requires asserted premises, however, explicit or implicit. That is to say, we can attain epistemic status for such an assertoric conclusion through a bit of reasoning only if our reasoning has premises with such assertoric status of their own.

We have found three weighty reasons to think that Descartes has no intention of “rejecting” his beliefs by abandoning them all, replacing the attitude of belief with that of suspension. First, we have seen his outright statement that *no one sane would ever do such a thing*. Second, doing so *de re*, for each belief in turn, lies beyond our psychological capabilities. Third, if he *were* to accomplish such universal suspension, he would necessarily block his own project! His project requires inspecting the “rejected” beliefs so as to determine whether they deserve readmission. And this inspection, this determination, must be done through reasoning that, in turn, would seem to require beliefs.

Given how problematic it is to understand rejection as abandonment, let us set aside that view of rejection and explore an alternative.

Descartes’s project is, I submit, on the second order. Taking a belief out of the basket is declining to endorse it epistemically in a certain way. (This is the “rejection” that he performs while “pretending” that the rejected beliefs are false. This pretense is itself on the second order; it picks out clusters of beliefs under a certain description and under that description pretends, about them generally, that they are false.) Whether one had earlier endorsed them or not, one now declines to do so. But what *is* the relevant way in which Descartes declines to endorse his “rejected” beliefs? Recall the special importance of the status of *certainty*, whereby one is in no doubt whatsoever that one’s belief is true. *Here is per-*

haps the key to how we should understand endorsement. Proper Cartesian endorsement of a belief requires one to have no reason, not the slightest, for any doubt about its truth. This is, then, the proper endorsement of a belief as *doubtless* true.

Cartesian rejection, when proper, would thus involve forbearing from such endorsement: that is, from endorsement of a belief as doubtless true, not just as true. To dump a belief from one's basket of beliefs is to forbear from thus endorsing it. One may or may not have previously endorsed it. One may have failed to do so much as *consider* whether to endorse it. In any case, the belief is dumped when one now forbears to endorse it. And now the dumping of a belief, its relevant "rejection," seems compatible with undiminished confidence in its content. So we would surmount two of the three key problems encountered earlier. If our account is correct, Descartes need not reduce his confidence in order to engage in his project of Cartesian doubt. Nor must he be deprived of beliefs in terms of which to conduct the inspection. On our account, Descartes retains sufficient confidence on the first order that his first-order beliefs can all remain in place, even when on the second order he forbears consciously endorsing those first-order beliefs. By retaining his first-order action-guiding animal confidence, he can sanely go on about his everyday business, adroitly avoiding deadly jumps off high cliffs, and he can continue to engage in first-order reasoning in terms of those retained beliefs.

However, we still face the third of our problems. How can Descartes access his beliefs individually *de re* so as to reject (forbear endorsing) them or, eventually, so as to endorse them? The answer is that his project requires no such distributed access to his beliefs separately, one by one. He explicitly notes how hopeless that would be.³ The relevant *rejection* and the correlative *endorsement* must be *under a description, de dicto*. We must be able to pick out beliefs in clusters so as to reject them or endorse them as "those that satisfy condition C" for some given condition. Thus, for example, if we find that doubt inevitably clouds any belief based (directly or indirectly) on perception, then we may be able to dump all "beliefs

³See the second paragraph of Meditation One (CSM II: 15).

based essentially on perception” by forbearing to endorse them under that description. That is perhaps how the project is supposed to go. But we must next consider a further twist.

B. CREDENCE VERSUS JUDGMENT

Recall Descartes’s insistence that no one *ever* seriously doubts the deliverances of the senses and that his habitual opinions are highly probable opinions that it is much more reasonable to affirm than to deny. But how, then, could he ever “reject” such opinions as doubtful, or false, even when engaged in the project of determining the extent of possible human certainty?

Well, there *is* something he clearly *can* do. He can “feign” anything he likes while still harboring his old and customary opinions with undiminished assurance. In the second passage considered above (from Meditation One), *feigning* is what he explicitly proposes to do (from the Latin *fingem*, French *feignant*). Moreover, he can still use his belief that not-*p* in reasoning he performs even while making believe (feigning) that *p*. Thus, at the movie theater I can appropriately forbear shouting a warning even when I make believe that I see someone about to be hit over the head from behind. Here I seem to rely through implicit reasoning on an assumption that no one within earshot really needs any such warning. And this action-guiding reasoning can be perfectly appropriate despite my concurrent feigning of the contrary. Make-belief is one thing, real belief quite another.

That does, however, bring up a further question. Why should Descartes have thought that pretending that not-*p* would help him resist the temptation to continue to believe that *p* while endorsing one’s belief? It helps here to distinguish between two attitudes that might be called “belief.” One is an implicit confidence that suffices to guide our action, including action on practical options, such as whether to shout a warning. The other is an act of judgment made freely and voluntarily or a disposition to so judge upon considering the relevant question. In his philosophical meditation, Descartes is clearly concerned with the second of these. He emphatically distinguishes two faculties. There is *first*

a faculty of understanding, whose deliverances, received passively, are "perceptions" with some degree of clarity and distinctness. And there is *second* a faculty of judgment based on the subject's free will.

A possible explanation thus opens up as to why Descartes may have thought that by pretending that *p* one might be helped to avoid believing that not-*p*. At the theater one might feign (through visual imagination) that someone is about to be hit with a hatchet. Surely one would *not* then also freely *judge* that no-one is about to be hit. In particular, one is unlikely to judge consciously that the scene before one is unreal. The "suspension of disbelief" involved in such imagination tends to block one's consciously disbelieving by *affirming* the opposite of what one imagines (the two of which may even fail to cohere).

Note, however, that this can leave one's underlying subconscious credence still in place with undiminished confidence. One certainly does not lose one's confidence that one is sitting in a darkened theater viewing a screen (and not seeing a gory murder instead). Despite making it harder to *judge* that not-*p*, moreover, pretending that *p* does not constitute an *insurmountable* obstacle. That might thus be how Descartes thought pretense would help in his project. It would counteract our normal automatic tendency to judge in line with our stored credences, but it would not make it *impossible* for us to so judge. However, we would now be more free to judge in line with true reason and not just custom.

Accordingly, we can also see how our everyday guiding attitudes, such as the appearances of the Pyrrhonists, can remain in place below the surface of consciousness and do their guiding, even if one forbears endorsing them, and also suspends conscious assent. One can sustain highly confident credence that *p* even while suspending any conscious endorsement of that attitude, and even while suspending any correlated conscious judgment that *p*.

Consider the Cartesian "perceptions" that can have various degrees of clarity and distinctness. These are not just *sensory* perceptions. Indeed, the most clear and distinct of them include a priori intuitions involving rational rather than sensory awareness. These are rather *seemings*, including not only sensory seemings but also a priori seemings. Moreover, we should focus not just on *initial*

seemings that might enter into conflicts to be resolved through pondering or deliberation. We should focus rather on *resultant* seemings, *credences* involving some degree of confidence, representable through the unit interval.

Based on the testimony of a friend, for example, it might initially seem that *p* until contrary testimony from a more trusted and trustworthy friend shifts the balance, so that now, both things considered, it seems that *not-p*. If no other factor is rationally relevant, then, one's resultant seeming is a seeming that *not-p*. This is a "credence," then, or, in other words, it is a confidence, however slight, an all-things-considered seeming that *not-p*. Credences may thus be viewed as simply resultant seemings, with a degree of confidence representable through the unit interval.

Such seemings will, then, come with some degree of clarity and distinctness, but their apparent degree is not necessarily their real degree. In order to qualify as *really* sufficiently clear and distinct, such seemings must satisfy epistemic requirements, and we might incorrectly take a seeming to be thus clear and distinct even when it falls short.

As did the Pyrrhonian skeptics long before him, Descartes believes that we can guide our lives practically through such confident-enough seemings or appearances, ordinary opinions that it would be laughable to put in serious doubt as one navigates an ordinary day. Such beliefs are *never* put in serious doubt. Some gain the status of sufficient clarity and distinctness, moreover, *not* through direct, unaided intuition but only indirectly, through deductive reasoning. It would appear, then, that we reason through such "perceptions," resultant seemings, Pyrrhonian "appearances." We rationally make our choices based on what seem to be the pertinent potential risks and rewards, costs and benefits, *all things considered*. We rationally take into account even resultant seemings of relatively low intensity or degree of confidence.

Consider the Pythagorean theorem, which is not immediately and obviously true, in the way of simple logical truths such as the principle that nothing is self-diverse or that no proposition is both true and false. That theorem is nonetheless provable with great simplicity and persuasiveness. It *becomes* clearly true with the help of that proof. In a more humble example, someone carefully mea-

sure a one-inch line that looks like this: <———>, and believes accordingly. Suppose a second line is then unveiled right below that one, looking like this: >———<. If not wise to the Müller-Lyer illusion, he might then conclude immediately that the second line was longer than one inch. Here he will have put together two strong resultant seemings and will have arrived thereby at his highly confident additional resultant seeming, to the effect that the second line is longer than one inch. Each of the three seemings is a resultant seeming; our subject is not aware that an illusion is in play and hence arrives as he does at his all-things-considered credences (or all-things-considered seemings).

Just as did the Pyrrhonists, Descartes can continue to inquire, and to guide his daily life, through the use of such credences (including propositional perceptions) even once they have been put in doubt.⁴ To put them in doubt is *not* to disable them from functioning in the guidance of action or even in the reasoning required for inquiry. Compatibly with a belief's retention of its animal or *cognitio* status and its ability to provide the guidance that we expect of such beliefs, the believer might nonetheless refuse to endorse his belief when it is brought to consciousness for rational inspection.

C. THE PROJECT OF THE *MEDITATIONS*

Consider how it goes in the early meditations leading up to the *cogito* passages. Descartes argues that *cogito* propositions at long last give us what we want: contents that we *can* believe with proper endorsement. These offer absolute safety from deception. In arguing for this, he must, of course, make use of certain premises. These are the premises in the reasoning that shows *cogito* beliefs to pass Cartesian inspection. Among these premises is the assumption <If I think that I am, then I am>. Take a skeptic who puts in doubt simple truths of arithmetic and geometry. Take one in doubt not

⁴There is more on Pyrrhonism in section D of chapter 9. A fuller interpretation is offered in chapter 10 of my *Judgment and Agency* (Oxford University Press, 2015).

only as to what 3 and 2 add up to, how many are the sides of a square, or whether there really are any shapes at all. No such skeptic is likely to grant us without question knowledge of the following: *if I think that I am, then I am*. The reasoning by which Descartes means to underwrite his certainty of the *cogito* can thus be seen to have a certain limitation: namely, that it relies on a premise also subject to skeptical doubt.

Descartes aims to establish that our beliefs with certain contents or with certain sorts of contents would be *bound* to be correct (and to do so without blatant bootstrapping or any other type of vicious circularity). However, his reasoning turns out to be *open to skeptical challenge*. *Provided* I know <If I think that I am, then I am>, I can thereby underwrite that I could not possibly go wrong in affirming that I exist, which protects my affirmation from radical deception. But the radical skeptic of Meditation One has put in doubt even the simplest a priori truths of arithmetic and geometry. No such skeptic will allow Descartes to just help himself to the premise that if he thinks he exists, then he does exist. This conditional seems no less dubitable an a priori claim than the simple a priori truths that Descartes's skeptic already puts explicitly in doubt in the early meditations.

Accordingly, Descartes will need to consider whether *such assumptions*—the ones he needs in order to underwrite even cogito thoughts—can themselves be upgraded in the sort of way he tries to upgrade cogito thoughts (with the limited success we have observed). Taking his cue from Meditation Two, he needs some way to legitimate such assumptions, to endorse them properly. This, I submit, is what sets up the project in the rest of the meditations. Descartes goes in search of reasoning that will satisfy certain specifications:

- (1) that it raise key beliefs to the required superlative level, even those that are now in some slight metaphysical doubt;
- (2) that it do so while avoiding the *blatant* bootstrapping of assuming as a premise the very conclusion to be argued for; and
- (3) that the beliefs so raised include the ones that enable him to endorse cogito propositions.

This project he pursues through the rational theology prominent in the later meditations. It is through such reasoning that he thinks he can upgrade his relevant beliefs. He can show them to have the required status because he can reach, through proper reasoning, the conclusion that his clear and distinct perceptions will reliably enough provide deliverances he can properly trust. And this reasoning will avoid *blatant* bootstrapping.

Let us back up a bit. What has put Descartes's beliefs in such slight doubt? Recall the skeptical scenarios of Meditation One: the dream scenario, for example, and the evil demon. In some of these we retain a normal set of beliefs about the world around us based on perceptual evidence, as is normal for such beliefs. Although we are there radically deceived, it is hard to see how we can possibly rule out those scenarios. If we cannot do so, however, then we can hardly be *certain* in our beliefs. Those are scenarios wherein our beliefs are false and hence *not* known to be true. Unless we can rule out that we are now so deceived, therefore, we cannot be sure we really know that our present beliefs are true.

That is one way of constructing a dream scenario, but there is also a second way with similar skeptical import. According to the first way, in the dream scenario we dream that *p* while it is false that *p*. According to the second way, we dream that *p* whether it is true or false that *p*. That this second case has for Descartes similar skeptical import is suggested by his fourth skeptical scenario. In that scenario there is no God to create us or sustain us; we emerge through "fate or chance or a continuous chain of events, or by some other means [other than Divine agency]" (First Meditation, CSM II: 14). Under that supposition, there is no metaphysical grounding for our assured competence. With respect to any question we take up, we might or might not be well enough constituted and well enough situated that we would not go wrong while properly using our faculties (our epistemic competences or abilities).

Note the strength of what Descartes requires, as suggested by the status he gives to that scenario: In order to attain *true certainty* on a question whether *p*, we *must* be so constituted that we *can* not go wrong (given adequate care and attention). Absent a powerful and benevolent enough creator and sustainer, however,

we will *not* necessarily be so constituted. Descartes accordingly requires not just the aptness of one's belief but its *superlative* aptness, which also includes its "security." A belief is thus secure only if the competence manifest in its truth is so safely in place that it cannot possibly have been missing.⁵

D. A DEEPER PROBLEM

Our proposal meets the three objections to the naïve view of the way Descartes addresses his epistemological problematic: (1) that he emphatically affirms that no one sane would *ever* put ordinary beliefs in serious doubt so as to reject them, (2) that it is hard to see how he could manage to put his vast corpus of ordinary beliefs in serious doubt so as to reject them individually and seriatim, and (3) that if the rejection involved is withholding—or, equivalently, suspension—of belief (and disbelief), then he deprives himself of the wherewithal required for the inspection to which rejected beliefs must be subjected before they can be properly readmitted into his body of beliefs.

Our proposal distinguishes between animal beliefs that can continue to guide us subconsciously in the everyday and the consciously reflective beliefs that are not needed for such animal guidance. What would be insane is the abandonment of the beliefs needed for guidance. The judgments involved in conscious reflection can be suspended, however, with no need to abandon the corresponding animal beliefs. These judgments do require the rational endorsement that is made problematic by the skeptic. So much for objection a.

As for b, it helps again to distinguish the vast storehouse of implicit animal beliefs from the conscious reflective beliefs that

⁵Of course Descartes does allow us a measure of freedom that makes it possible for us to go wrong even if thus endowed. Where we cannot possibly go wrong is in our understanding, in our ability to perceive with sufficient clarity and distinctness what is thus perceivable (so long as we avoid unfortunate inattention, passion, and other such disablers). It is this ability that is *securely* our God-given endowment.

constitutively involve judgment. There are implicit judgments, or *judgmental* beliefs, such as when we say of someone asleep that “in his judgment” we ought to pursue a certain course of action. We are not saying that he is at that moment, while asleep, performing a certain act of judgment. Rather we are saying that he is disposed to so judge occurrently if he considers the question and endeavors to answer it correctly, disposed to affirm accordingly at least to himself, *in the privacy of his own mind*. Suppose our episodic judgments to be largely ones over which we exercise voluntary, free control, *as Descartes emphatically believed*.⁶ In that case, the corresponding dispositional judgments will be in effect freely upheld policies to answer corresponding questions affirmatively. And it is not at all implausible that *these* beliefs, these judgmental policies of response upheld by the will, could be modified with a general act of will. The act of will involved would be quite like deciding to abandon our policy to signal our turns as we drive. In one fell swoop, we would affect the policy as it concerns each of the corners where we turn as we drive home every evening. Through a general dispensation we change each of those more specific policies by changing the overall governing policy. Similarly, we could try to change our general disposition to respond affirmatively to a great variety of questions as we took them up. Skeptical reasoning could surely affect our beliefs that way. As a result of the conscious reasoning, we might try to give up our policies to respond affirmatively to such questions as “Is there such a thing as snow?” “Is it white?” “Is there such a thing as the sky?” “Is it blue?” And so on. So, once we focus on the sorts of beliefs that held primary interest for Descartes, as for the Pyrrhonists, we can much more plausibly consider a universal abandonment of our beliefs, that is, of our judgmental beliefs: of our dispositions to respond affirmatively, freely, and voluntarily.

⁶His faculty of judgment was actively volitional by contrast with the faculty of the understanding, whereby he explains how error is possible despite God's perfect epistemic benevolence.

E. FOUR KEY CONCEPTS OF CARTESIAN EPISTEMOLOGY

We have touched already on *certainty*, *doubt*, and *endorsement*, but we have yet to consider Cartesian *error*. What we uncover about this crucial concept will also bear on the other three.

Error. Ordinarily we take error to consist in falsity. An erroneous belief or opinion is just a false one. How far this is from Descartes's own view may be seen in the following two passages:

But there was something . . . which I used to assert, and which through habitual belief I thought I perceived clearly, although I did not in fact do so. This was that there were things outside me which were the sources of my ideas and which resembled them in all respects. Here was my mistake; *or at any rate, if my judgment was true, it was not thanks to the strength of my perception.*⁷ (Third Meditation, CSM II: 25; emphasis added)

If . . . I simply refrain from making a judgment in cases where I do not perceive the truth with sufficient clarity and distinctness, then it is clear that I am behaving correctly and avoiding error. But if in such cases I either affirm or deny, then I am not using my free will correctly. If I go for the alternative which is false, then obviously I shall be in error; if I take the other side, then *it is by pure chance that I arrive at the truth*, and I shall still be at fault since it is clear by the natural light that the perception of the intellect should always precede the determination of the will. In this incorrect use of free will may be found the privation which constitutes the essence of error. (Fourth Meditation, CSM II: 41; emphasis added)

Falsity is *sufficient* for error but not *necessary*. One can still be in error with a true belief so long as its truth is not attributable to

⁷In the Latin: "Atque hoc erat, in quo vel fallebar, vel certe, si verum judicabam, id non ex vi meae perceptionis contingebat." In the French: "Et c'était en cela que je me trompais; ou, si peut-être je jugeais selon la vérité, ce n'était aucune connaissance que j'eusse, qui fût cause de la vérité de mon jugement."

one's clear and distinct perception. One's belief is then true by accident:

It is also certain that when we assent to some piece of reasoning when our perception of it is lacking, then either we go wrong, or, if we do stumble on the truth, *it is merely by accident*, so that we cannot be sure that we are not in error. (CSM I: 207; emphasis added)⁸

In the crucial second paragraph of the Third Meditation we find a further clue. By that point we have finally reached a true certainty, *sum res cogitans*. Having wondered aloud what could possibly yield such certainty, Descartes answers his own question. "As far as I can see, certainty here derives from clear and distinct enough perception."⁹ Perception of such clarity and distinctness is said to yield certainty, however, only if nothing could ever be so clearly and distinctly perceived and yet be false. It is such clarity and distinctness, then, that will properly account for the correctness of one's perception, with no chance of falsity, so that it will perfectly explain why the corresponding judgment must be true. It must be true because it corresponds to a perception by that subject so clear and distinct that it could not possibly be false.¹⁰

The *essence* of error is said to reside in a judgment that does *not* manifest the sort of competence required, one whose manifestations would leave little enough to chance. (See the earlier passage from Meditation Four.) So even when we judge with truth, as Descartes emphasizes, we can be in error if our judgment fails to be true "thanks to our perception," with its required level of clarity and distinctness. When our hitting the mark of truth is *not* thus explained by a competence that leaves nothing to chance, our judgment is still in error despite being true.

⁸ A defense of Descartes's exact words here would require some verbal acrobatics. For it would seem that a judgment could fail to be infallible without descending to being right *merely* by accident. But we can perhaps take his formulation to be a forgivable exaggeration in line with his focus on certainty.

⁹ As he notes in the third paragraph of Meditation Three; see CSM II: 24.

¹⁰ Here it is easy to get confused. What accounts for the truth of the *judgment* is not necessarily what accounts for the truth of its *content*.

A judgment might be not only true but indeed *necessarily* true, while still in error. Suppose one believes that the square of 2 squared is 2 to the fourth power. One hence multiplies four 2s, concluding thereby that the square of 2 squared is 16. Suppose one arrives at how many 2s to multiply, however, by *adding* the exponents. Only because adding these two exponents ($2 + 2$) yields the same as multiplying them (2×2) does one here get the right result. If the exponents had been in any way different, one would have arrived at the wrong result. It is no thanks to competence that one hits the mark of truth. Yet one's judgment could not possibly be false, since the square of 2 squared could not possibly be anything other than 16.

Certainty. To attain absolute certainty, then, is to hit the mark of truth in one's judgment or belief thanks entirely to (the quality of) one's perception, which could not possibly lead one astray. To attain such certainty is thus superlatively to avoid error. One hits the mark of truth, true enough, and moreover does so thanks to (the quality) of one's perception. But one does even more than that, since one's perception is of such high epistemic quality that it leaves no room for error in the way ordinary, sensory perception leaves room for error even when highly reliable.

Doubt. To doubt a certain content is to forbear endorsing belief of it, and in Descartes's quest for certainty one is to forbear unless one can endorse one's belief as certainly apt. No matter the intensity of one's credence, one still entertains some doubt concerning that belief so long as one forbears endorsing it as certainly apt. A reason for doubt is, accordingly, a reason to forbear endorsing.

There is more than one way to adopt such a meta-attitude toward a credence that you hold. You might adopt it under a description, whereby you pick out the credence as one that satisfies a certain condition: "credence whose source is perception," as it might be. Alternatively, your meta-attitude might instead target a belief whose content is fully present and on display as the focus of your attention. If this is how you forbear from endorsing your belief that p as certain, your forbearing will bring with it your suspending judgment on the question whether p : you will judge

on that question neither affirmatively nor negatively; you will neither affirm nor deny.

Endorsement. And so we come to this important concept. In keeping with the foregoing thoughts, to endorse a belief is to regard it as correct, and in Descartes's project that requires one's endorsing it as *certainly* apt, as one that hits the mark of truth thanks entirely to the subject's clear and distinct perception, where this in turn amounts to an infallible competence. If one endorses the belief that *p* while aware of its content *directly present to our minds*, one will also judge affirmatively on the question whether *p*. For Descartes this requires certainty, moreover, so that the belief must manifest infallible competence.

F. THE CARTESIAN PROJECT

The Cartesian epistemological project is, at a minimum, one of examining human epistemic competence, our actual modes of acquiring and sustaining beliefs. Descartes considers how defensible are our actual modes, and also which are the best of those available to us.

One way in which a first-order credence—whether continuing or newly acquired—might benefit from such a project is by the subject's picking it out specifically, with its content on full display, and endorsing it while thus picked out. Such endorsement is fully proper only if the subject knows the competence involved to be, first, sufficiently reliable and, second, manifest in his holding of that first-order credence.

A normal human could not upgrade many of his credences up to that level, at least not *through conscious reflection*, at any given time. There is a limit to the scope of our concurrent attention.

If we moderate our ambition, we can widen the reach of certainty, however, by allowing the second-order endorsement to be implicit and to remain implicitly stored in memory. We require the judgment to be made or sustained competently enough through the sufficient competence of the faculty that prompts it. We may even require that the subject have an appropriate second-

order account of how that faculty is so reliable. However, we must *not* require, for this more realistic level of upgrade and for the corresponding endorsement, that the specific belief be picked out separately and consciously. It suffices, first, that the subject have some implicit awareness of it as a belief that manifests the competence in question while, second, his sustaining of that belief is positively influenced by that awareness.

G. CONCLUDING REMARKS

In the view here defended, Descartes uses his principle of clarity and distinctness in order to raise his first-order judgments to the *scientia* level. He must assure himself that judgments rationally based on clear and distinct perceptions avoid error. *Error* is what one must avoid, not just falsity. So he seeks not just truth but also *aptness*. And aptness requires a good enough competence, one that is reliable enough. You are to assure yourself that you attain such aptness, which is required for confidence that you avoid error, and attain certainty. But this assurance is forthcoming only with assurance that the operative source of your judgment is indeed a reliable-enough competence. And this includes your present judgments as well as those you made in your past or will make in your future. This raises an issue of circularity, since it is hard to see how you could really assure yourself of how reliable your competence is in the absence of any first-order premises. So we face an issue of vicious circularity and the notorious Cartesian Circle. Moreover, it is a circle that also affects contemporary virtue epistemology when it postulates a level of reflective knowledge above that of animal knowledge. This is what one would expect if one perceived the parallel between the two epistemological distinctions: that between the animal and the reflective on the contemporary scene and that between *cognitio* and *scientia* in Cartesian epistemology. Virtue epistemology, whether Cartesian or contemporary, must address this allegedly vicious circle, which it can arguably do with success.¹¹

¹¹ Epistemically vicious circularity is the theme of my *Reflective Knowledge* (Oxford University Press, 2009), with its defense of virtuous circularity.

It was already hard for Descartes's contemporaries to make sense of his epistemological project with its distinctive quest for certainty. This comes out most famously and clearly in the *Objections and Replies*. There is no promising way to do so except by (1) ascending to the second order, as we have done; (2) distinguishing credence from judgment; (3) highlighting the fact that some propositions are indubitable because, even when considered consciously and reflectively, they demand our assent; and (4) addressing the problem of the circle through the distinction between *cognitio* and *scientia*. But proceeding in this fourfold way raises the difficult interpretative and philosophical questions that we have taken up.

In all important structural respects, Cartesian virtue epistemology takes the same view as a virtue epistemology on the contemporary scene, *virtue perspectivism*.¹² The structure of this view does not require the theological content that Descartes gives to his own version. The role of theology can be played instead by science, by common sense, or by the two combined. Although we have seen in this chapter how the two versions of virtue epistemology are closely akin, the full extent of the kinship remains to be detailed. But it should become increasingly clear as the contemporary view is developed to include more explicitly the epistemic agency that played so central a role in Cartesian epistemology. This is a development of the contemporary view that is now underway.¹³

FURTHER READING

Broughton, Janet. *Descartes's Method of Doubt*. Princeton, NJ: Princeton University Press, 2001.

Frankfurt, Harry. *Demons, Dreamers, and Madmen*. Indianapolis, IN: Bobbs-Merrill, 1970.

Newman, Lex. "Descartes's Epistemology." In *The Stanford Encyclopedia of Philosophy* (Winter 2014 edition), ed. Edward N. Zalta. Stanford,

¹²This is the view I defend in *Knowledge in Perspective* (Cambridge University Press, 1991) and in a series of later publications.

¹³In *my Judgment and Agency* (Oxford University Press, 2015).

Chapter One

CA: Metaphysics Research Lab, Center for the Study of Language and Information. Available at <http://plato.stanford.edu/archives/win2014/entries/descartes-epistemology/>.

Van Cleve, James. "Foundationalism, Epistemic Principles, and the Cartesian Circle." *Philosophical Review* 88 (1979): 55–91.

Williams, Bernard. *Descartes: The Project of Pure Inquiry*. Atlantic Highlands, New Jersey: Humanities Press, 1978.

Dream Skepticism

A. PRELIMINARIES ON SKEPTICISM IN GENERAL

If a skeptic challenges us to prove *ex nihilo* that we see a hand or a fire, we should decline. Nor should we take the bait when he offers us as premises just obvious *a priori* truths and facts about our own current subjective states. Even given such premises, pure reason (deductive or inductive) will not yield nearly enough of what we believe ordinarily.¹ Our knowledge of hands and fires is not explicable based on proper reasoning from the agent's reflectively accessible subjective states (from the "given"). Far from accepting such skeptical challenges, we should decline the pre-suppositions upon which they are launched. We should reject the sheer assumption that only reasoning from the foundational given will provide or explain whatever knowledge we may enjoy. Are there skeptics more menacing than that?

Academic skeptics are more assertive, if only by claiming that we know nothing, either in general or in some large department

¹This skepticism is based in part on how implausible it seems that proper scientific inference and confirmation could be carried out in the privacy of one's own mind. Consider old models of inductive reasoning, either formal or informal, such as enumerative induction or inference to the best explanation. The worry is not only about the unpersuasiveness of such models when applied to the private given (though that is worrisome enough). More relevantly, those models seem already implausible on their own as accounts of proper scientific procedure. Even a sophisticated Bayesianism must face problems like those in John Earman's *Bayes or Bust* (MIT Press, 1992). In response, one might turn to the kinds of social factors that Thomas Kuhn and others had emphasized in preceding decades (as does Earman). But these are factors that will have to be checked in when we enter the inner sanctum of the lone thinker's mind.

of our supposed knowledge: the external world, for example, other minds, or morality.²

Granted, any attempt to refute such a *fully* global claim *will* beg the question.

A claim might be irrefutable, however, without being true. Consider the global claim that we know nothing at all. To refute this, one must adduce some premise, doing which is implicitly to claim knowledge of its truth, thus begging the question. Because the context of dialectic prohibits such question-begging, our global skeptic is dialectically irrefutable. But he is not thereby shown to be right. Besides, if the very making of a claim commits the global skeptic also to knowledge of what is claimed, then he contradicts himself.

Such varieties of skepticism are less problematic than one that spots a commitment ostensibly at the heart of common sense and argues on its basis that there is very little we can know, either in general or in some main department.³ Let us next turn to that more interesting challenge.

B. SKEPTICISM, SENSITIVITY, AND SAFETY

Our skeptic posits that a belief constitutes knowledge only if *sensitive*, only if it satisfies the following condition: *had the proposition believed been false, it would not have been believed by the believer*. In order for you to know that you see a hand, your belief must be *sensitive* in that, *had you not seen a hand, you would not have believed that you saw one* (even based on your visual basis, on your visual experience as of a hand). Thus, your speedometer reading may happen to be correct when it is stuck on a certain

²Pyrrhonian skepticism, by contrast, requires suspension of belief, on question after question, rather than either affirmation or denial.

³Let's be clear on what the project is. The objective is to show how we *could* become doxastically justified despite the arguments of skeptics, even if we put aside the skeptics who make unreasonable demands, such as the demand that in order to be justified you must have an infallible basis, or the like. If a philosopher (such as Descartes) argues that we actually *are not* justified, he does so in effect by arguing that we *could not be* justified, from which it follows modally that we *are not*.

reading. Correct it may be, but still not a source of knowledge even so. Why? The advocate of a sensitivity condition for knowledge has a ready answer: *Even if your speed had been different, you would still have believed the same, on the same speedometer basis.* So your speedometer-based belief is not sensitive. Our skeptic alleges that this sensitivity requirement is among our core common-sense commitments. And if he is right, then the skeptic is in a good position. Belief that one is not radically misled cannot be sensitive. Indeed, skeptical scenarios are framed to secure precisely this result. If you were now a brain in a vat being fed experiences as if you enjoyed normal perception, that would not stop you from believing that you were *not* radically deceived.⁴

In a later chapter we shall consider more closely that requirement of sensitivity. Here we shall move beyond it quickly. We begin by noting that subjunctive conditionals do not validly contrapose. For example, *if water flowed out of my kitchen faucet, it would be false that water flowed while the main house valve had been closed for hours with the faucet open.* But consider the contrapositive of this conditional: namely, *that if water flowed out the faucet while the main valve had been closed for hours with the faucet open, then water would not flow.* This contrapositive would be obviously false, showing contraposition to be invalid for such conditionals.

That suggests a requirement of safety rather than sensitivity, as follows:

Basis-relative safety. A belief cannot constitute knowledge if the believer might too easily have so believed on the same basis while his belief was false. Alternatively, in order to know, one must believe on a basis such that one's so believing on that basis must have a strong enough tendency to be right. (Here the basis might just possibly be the null basis, as a limiting case.)

⁴Various responses to such skepticism have been developed over recent decades, including closure-denying tracking approaches and contextualist semantic ascent. A rich and subtle dialectic can be found in an extensive and still growing literature. No such sensitivity-influenced approach is fully satisfactory, however, though that is not something we can go into in the space available here. See the guide to further reading for more on these issues.

Safety does not serve the skeptic as well as sensitivity. Belief that one is not radically deceived is insensitive, and it is even insensitive relative to its deep experiential bases. But a belief can be safe while insensitive. Scenarios of radical deception are outlandish, remote possibilities *not* too easily liable to occur. Therefore, a belief that one is not radically deceived is safe while insensitive: *not* too easily might one have been radically deceived. Belief that one is not radically deceived would tend to be correct. The possibility of radical deception is so outlandish that one's belief to the contrary would tend to be correct.

C. WHY THE DREAM SCENARIO IS SPECIAL*

1. Our line of reasoning is effective against radical scenarios, such as that of the brain in a vat; Descartes's evil demon, the Matrix; and so on. Only *the dream scenario* stands apart. Dreams being so common, the possibility that one dreams is not outlandish. Therefore, we cannot defend the safety of our belief that we are awake by alleging that it would not be at all easy for us to go wrong in so believing. After all, too easily for comfort might we have been *not awake* but only *dreaming*. This is the dream scenario wielded by the skeptic.⁵

Which dream scenario? The dream scenario that life—*all of life*—is but a dream is hardly less outlandish than a brain-in-a-vat or evil-demon scenario. So that is *not* the dream scenario of special interest to us. The more relevant dream scenario is the one posed by Descartes when he wonders, as he sits before a fire, whether he is *then* just dreaming. Our dream scenario is one that arises for any arbitrary case in which we consider whether our

⁵The epistemological problem of dreams appears early, as when Socrates asks Theaetetus: "How can you determine whether at this moment we are sleeping, and all our thoughts are a dream; or whether we are awake, and talking to one another in the waking state?" (*Theaetetus* 158b–d).

*Here I mark with numbers or letters subsections that have a certain unity of theme or thought, even within a more broadly unified line of argument. I return to this device when it seems helpful, in this chapter and occasionally in later chapters.

ostensible perceptual knowledge—of a fire, say, or a hand—is real *on that occasion*.

Such a dream scenario has a distinctive importance by comparison with the familiar radical scenarios. Unlike the others, it is a close possibility, especially if bedtime is near or one is already abed. At that point the possibility that one is dreaming is not a remote possibility. Yet we are still properly reluctant to grant the skeptic that we no longer know that we're planning the next day's agenda (for example) or that the room is dark.

For another thing, our dream scenario threatens our perceptual beliefs *directly*.⁶ The skeptic might argue less directly as follows: "If you could really know through perception that you see your hand, you could competently deduce that you are not then a brain in a vat and could thus know the latter that way." And he could then triumphantly conclude: "Since you cannot know that conclusion that way, therefore you do not know the premise either." That, too, is a considerable, threatening argument. But our earlier reasoning does not depend on the assumption that you always know what you competently deduce from what you already know, and it especially does *not* depend on the assumption that what you competently deduce from the known is something that you *thereby* know.

The threat to our perceptual beliefs that is posed by the nearby possibility that one dreams involves *danger* or *risk*, the danger or risk of believing falsely. This does not depend on the endangered subject's awareness that he may be at risk, nor does it depend on his evidence concerning such risk, or even on whatever evidence may be *available* to him. As one strides across a minefield, one can be in great danger, even with no inkling of that fact or any available evidence for so believing.

2. The danger to one's attaining ordinary knowledge does require a certain orthodox conception of dreams, according to which beliefs

⁶The dream scenarios might also be used to argue that, just as the modal proximity of enough possible encounters with fake barns creates a problem for the belief that one perceives a barn, so the modal proximity of enough possible dreams creates a problem for the belief that one perceives an external reality. Here *X* is modally proximate to our actual situation in proportion to how easily it might have happened, to how close (as opposed to remote) a possibility it is, relative to our actual situation.

and experiences in our dreams are hosted not only in the dream but also in actuality, while we dream. Only thus would our ordinary perceptual beliefs be threatened by the possibility that in a realistic dream we might believe the same on the same experiential basis.

Is that really how we should conceive of our dreams? Are dreams made up of conscious states just like those of waking life except for how they fit their surroundings? The orthodox answer is in the affirmative. Dream states and waking states are thought to be intrinsically alike, though different in their causes and effects.

The dream argument stands out, then, because the dream possibility is too close for comfort. If while dreaming we have real beliefs based on real phenomenal experiences, then a normal perceptual judgment could always be matched by a subjectively similar, similarly based judgment made while we dream. Too easily, then, might we now be dreaming when we form perceptual beliefs. On the orthodox conception, a dreaming subject might then through his dream form such a belief not just in the dream but in reality, albeit while dreaming. No doubt it would be a false belief, based on illusory phenomenal experience. Any given perceptual belief, or one intrinsically just like it, might thus too easily have been false though formed on the same experiential basis. This possibility, too close for comfort, threatens perceptual belief more than any radical scenario.

Fortunately, the orthodox conception is not beyond question. A lot rides epistemically on just how dreams are constituted.

D. WHAT DREAMS ARE MADE OF

1. Do the characters in my dreams have beliefs and intentions? They do in general, but do I myself also have them as protagonist in my dream? Unquestionably I do believe and intend things *in my dream*.⁷ In my dream I am conscious, I assent to this or that,

⁷Here I distinguish between first-person participation in the dream and third-person participation, as when one sees oneself do something as if in a movie or on

I judge or choose.⁸ This all happens *in the dream*, of course, but does it thereby *really* happen, albeit while I dream? When, if ever, is that so?

When something happens *in my dream*, reality tends not to follow suit. When in my dream I am chased by a lion, this poses no threat to my skin. No physical proposition about the layout of the world around me is true in actuality just because it is true in my dream. What about mental propositions about how it is in my own mind? Must any such proposition be true in actuality whenever it is true in my dream? No, even if *in my dream* I believe that a lion is after me, and even if *in my dream* I intend to keep running, *in actuality* I have no such belief or intention. What is in question is the *inference* from <In my dream I believe (or intend) such and such> to <In actuality I so believe (or intend)>.

My exposition relies heavily on distinguishing between two expressions: “in my dream” and “while I dream.” From the fact that *in my dream* something happens it does not follow that it happens *while I dream*. From the fact that in my dream I am chased by a lion it does not follow that while I dream I am chased. Moreover, from the fact that while I dream something happens it does not follow that it happens in my dream. From the fact that while I dream it rains and thunders it does not follow that in my dream it rains and thunders.

2. At any given time, nearly all one’s beliefs remain latent. A belief might be manifest when formed, or it might occasionally rise to consciousness from storage. To make one’s belief explicit one must *judge* or *assent* or *avow*, at least to oneself.⁹ The same is true of one’s intentions, few of which rise to consciousness at any given time. One does, of course retain countless beliefs and intentions while asleep and dreaming. Among these are intentions recently

a TV screen. One can figure in one’s movie as a victim of a recent knockout and will not thereby undergo any present experience.

⁸Let’s use “affirmation” here for conscious assent to propositional content and “volition” for conscious assent to a possible course of action (including simple actions and even, as a limiting case, those that are basic and instantaneous).

⁹However, as will emerge, one might judge or avow something that one does not believe, and even something that one disbelieves.

formed—to stop by the library the next day, for example—and beliefs recently acquired—that the weather will be fine in the morning, say. If so, then what one knows as one dreams is that one is at home for the night; one lay down in the knowledge that one would be there that long, and this knowledge has not been lost. *Staying home after dinner until the morning* is what one intended through most of the day, even as one thought about other things, as one had dinner, and so on. That was still one's intention as one lay down, and there is no reason to suppose that it was lost as one fell asleep. One does not lose one's intentions for the coming morning. One retains intentions as to what one *will* do upon awakening. One also retains, as one drifts off to sleep, the belief that one is at home for the night. This one knows as one goes to bed, and there is no reason to think that one loses that knowledge, given all the knowledge and intention that one retains subconsciously. It is hard to see how one could then concurrently believe that one is being chased by a lion in an open field, rather than at home.¹⁰

Granted this for our guiding states of belief and intention, with their crucial functional profiles, perhaps conscious episodes are different. Perhaps one does really undergo these while dreaming whenever one does so in one's dream, conscious judgments and choices included. Conscious assent to a proposition does not guarantee that it is really *functionally* believed, nor does conscious assent to a course of action guarantee the corresponding intention. One might even consciously assent to the opposite of what one functionally believes or intends. Actions speak louder than words this way, and louder than conscious assents, too. A deep-seated prejudice might be disavowed sincerely while still surviving, firmly entrenched. Similarly, a belief might survive in storage while consciously disavowed in a dream. Conscious affirmations and volitions might thus contradict stored beliefs and intentions,

¹⁰ A sleepwalker who steps outside, or even one who just walks near the door, does not know that he is at home, *unlike* the normal sleeper who believes himself to be at home for the night and who knew as he dozed off that he would be home for the night, a knowledge that he does not lose by falling asleep any more than he loses his knowledge of math or geography or history or his knowledge of what he will be doing in the day ahead.

and dreams may provide just a special case of that general phenomenon. The fact that one retains stored beliefs and intentions while dreaming seems thus compatible with real affirmations and volitions to the contrary, made not only in one's dream but thereby also in reality, while dreaming.

3. What then of propositions about your own *current* conscious states, whether conscious experiences or conscious assents? Even if while dreaming you do not *functionally believe* that you are then chased by a lion, perhaps you do still consciously *affirm* it. If in a dream one is in a certain conscious state, is one then *actually* in that state, while dreaming? If in my dream I make a conscious choice, do I thereby really make that choice, while dreaming? In a dream you may covet thy neighbor's wife, in the dream a sultry object of desire. Do you then violate the Biblical injunction? If in the dream you go so far as to succumb, are you then subject to blame? Having sinned in your heart, not only in your dream but in actuality, you could hardly escape discredit. Is one then blameworthy for choices made in a dream? That has near-zero plausibility, about as little as blaming a storyteller for his misdeeds as the protagonist in a story spun for a child. (One might blame him for telling such a story to such an audience, but that is different; one does not thereby blame him for doing what he does *in the story*.)¹¹

¹¹ Compare Augustine in book 10, chapter XXX, of his *Confessions*: "Verily Thou enjoimest me continency from the lust of the flesh, the lust of the eyes, and the ambition of the world. Thou enjoimest continency from concubinage; and for wedlock itself. Thou hast counselled something better than what Thou hast permitted. And since Thou gavest it, it was done, even before I became a dispenser of Thy Sacrament. But there yet live in my memory (whereof I have much spoken) the images of such things as my ill custom there fixed; which haunt me, strengthless when I am awake: but in sleep, not only so as to give pleasure, but even to obtain assent, and what is very like reality. Yea, so far prevails the illusion of the image, in my soul and in my flesh, that, when asleep, false visions persuade to that which when waking, the true cannot. Am I not then myself, O Lord my God? And yet there is so much difference betwixt myself and myself, within that moment wherein I pass from waking to sleeping, or return from sleeping to waking! Where is reason then, which, awake, resisteth such suggestions? And should the things themselves be urged on it, it remaineth unshaken. Is it clasped up with the eyes? Is it lulled asleep with the senses of the body? And whence is it that often even in sleep we resist, and mindful of our purpose, and abiding most chastely in it, yield no assent

If while dreaming one does *actually* assent to misdeeds, even to crimes, does its being just a dream protect one from discredit? That seems implausible. If sudden paralysis prevents you from carrying out some deplorable intentions, this does not protect you from discredit, from the full weight of the Biblical injunction. How, then, can you be protected by the disengagement of your brain from the physical causal order? How can you be protected by the disengagement of your inner mental life, as in a dream?

4. Is dreaming perhaps like being drunk or drugged? These disabling conditions lighten responsibility. Perhaps when dreaming you do make conscious choices, while your disabling state lightens your responsibility. Is *that* why we don't blame people for sins in their dreams? No, it is not that one is *less* responsible for what happens in one's dream. Rather, one is not responsible in the slightest.¹²

Dreams seem more like imaginings, stories, or even daydreams, all fictions of a sort, or quasi-fictions. Even when in a dream one makes a conscious choice, one need not do so in actuality. Nor does one necessarily affirm in reality whatever one consciously affirms in a dream.

5. What, then, of current *phenomenal* experiences? Does their presence in a dream entail their real presence in the conscious life of the dreamer, albeit while he dreams? Here at least, it may be

to such enticements? And yet so much difference there is, that when it happeneth otherwise, upon waking we return to peace of conscience: and by this very difference discover that we did not, what yet we be sorry that in some way it was done in us^o (E. B. Pusey translation).

¹²But here it is important to distinguish between being responsible for *dreaming that one commits a certain sin* and being responsible for *committing that sin*. One might suspect (however implausibly) that one bears some at least indirect responsibility for dreaming *such* a dream, while denying that one bears any responsibility whatsoever for committing that sin, especially when one *in fact* commits no such sin. Nor does it seem plausible that one bears any responsibility for the corresponding intentions and choices in the dream, which aligns with the view that these too are as imaginary as when they appear in a story that one spins. (And of course here again one might still suspect that one might still bear some indirect responsibility for *dreaming* that one harbors such intentions or makes such choices.)

thought, we can plausibly draw the line. But consider the consequences. With respect to such experiences it is supposedly just as if a lion were after me. Yet I may form neither the belief that this is so nor the intention to escape. Am I not now deserving of discredit? Even if such a belief and such an intention are formed *in the dream*, they are not thereby formed *in actuality*, despite the actual experiences that would seem to require them in anyone rational. If the phenomenal experiences in dreams *are* real experiences, while dream beliefs are not real beliefs, then every night we are guilty of massive irrationality or epistemic vice.

E. DREAMS AND SKEPTICISM

1. Let us now explore what follows for philosophy from the view of dreaming as imagining. If that is the right model, then traditional formulations of radical skepticism, Descartes's included, are not radical enough. The possibility that we dream now threatens not only our supposed perceptual knowledge but even our supposed introspective knowledge, our supposed takings of the given. It is now in doubt not only whether we see a fire but even whether we *think* we see a fire or *experience* as if we see it. How so?

With my hand in view, I may ask: Do I now *think* I see a hand? Well, might it not be just a dream? Might I not be only *dreaming* that I think I see a hand? If I am only dreaming, then I do not *really* think I see a hand, after all.

If I *ask* whether I think I see a hand, however, I cannot thereby be dreaming that I think I see a hand. If in my dream I ask myself a question and answer it with a choice or an affirmation, the asking plausibly belongs with the choice or the affirmation. If the latter belongs only in the dream, not in reality, the asking will also have its place in that same dream. So, again, if I really ask whether I think I see a hand, I cannot thereby be only dreaming that I think I see a hand. Is this not privileged access after all, protection from the possibility that it is just a dream?

2. Fair enough. But compare my question whether I *see* a hand. If I really *ask* whether I see a hand, I cannot *thereby* be dreaming

about the hand and my seeing it. So we seem to have equally privileged access to the fact that we see a hand, at least equally privileged in respect of *protection from the dream argument*.

What might possibly make the cogito *especially* privileged? What could give it a status not shared by such perception of the hand? One advantage, at least, it turns out *not* to enjoy: it enjoys no special protection from the possibility that one is only dreaming.

"The cogito has got to be different nonetheless from our knowledge that we see a hand," comes the reply. In response, we might defend the cogito by retreating to a thinner concept of thinking, in which even dreaming and imagining are themselves forms of "thinking." On the *thicker* notion of thinking, if I imagine that *p*, hypothesize that *p*, or dream that *p*, I do not *thereby* think that *p*; I may not even think that *p* at all. On the *thinner* notion of thinking, by contrast, in imagining that *p* one *does* thereby think that *p*. And the same is now true of dreaming. On the thinner notion, in dreaming that *p*, one does thereby think that *p*. More idiomatically, let's say rather this: in dreaming or imagining that *p*, one *hosts the thought that p*. So, "thinking that *p*" in the thinner sense would amount to "hosting the thought that *p*"; one can host that thought even by just asking oneself whether *p*.

Compare (a) one's affirming that one affirms something with (b) one's having (the thought) that one has a thought. The latter is also a self-verifying (thin) thought. But it has, in addition, something missing from the former: namely, being dream-proof. If one were now dreaming, one would be affirming nothing. But one would still host the thought that one was having a thought.

So not only is my present thought that I host a thought guaranteed to be right; in addition, I would not so much as *seem* to have it without having it, not even if I were dreaming. Compare my *affirming*, my *affirmation*, that I am affirming something. This, too, is guaranteed to be right. But, by contrast with the thinner thought, here I *could* be fooled by a mere appearance. I might right now be dreaming that I am affirming something, while in fact affirming nothing. So things might in a way seem subjectively just as they do now, although I am *just* dreaming: thoughts might be crossing my mind, and I might thus host them, without my really affirming anything.

However, the more defensible thinner thought falls short crucially in the dialectic against the skeptic. It is not the sort of thought that suffices to constitute knowledge. Knowledge requires something thicker than merely hosting a thought. Accordingly, the move from thick thought to thin thought is not a way to save the cogito after all.

3. Consciously and affirmatively thinking that I think does have a special status: one could not go wrong in so thinking. It can thus attain a high degree of reliability and epistemic status. It attains this status through its *being* a conscious state of thinking that I think. Moreover, this status is not removed, or even much diminished, by the threat of an impostor state, one subjectively very much like it. A vivid and realistic dream is, of course, subjectively very much like its corresponding reality.¹⁵ Perhaps it is only *in my dream* that I now affirmatively think that I think. Despite being subjectively much like the state of thinking that one thinks, in dreaming one does *not* think; one does not so much as think that one thinks. That is to say, even if *in one's dream* one affirmatively thinks that one thinks, this does not entail that in reality one so thinks that one thinks, while dreaming.

Those two states can thus be hard to distinguish subjectively, though in only one is the subject justified in thinking such and such. Of course the two states are constitutively different. One is an apparent state of thinking one thinks, doing so (thinking one thinks) *only in a dream*, so that it is really only a state of *dreaming* that one thinks one thinks. By contrast, the other is a state of thinking one thinks, doing so (thinking one thinks) *in actuality*. Only the latter yields justification for one's thought that one thinks. The former not only yields no such justification; in it there *is* no such thought—this despite the fact that, by hypothesis, the two states

¹⁵Much as someone with a powerful visual imagination can picture a scene so vividly that the imagined scene and the one seen earlier are very much alike in content, despite the fact that the two conscious states do not share any actual sensory experiences.

are indistinguishable, as indistinguishable as is reality from a realistic enough dream.¹⁴

4. Have we here found a way to defend our perceptual knowledge from the skeptic's dream argument? Even if we might just as easily be dreaming that we see a hand, this does not entail that we might now be astray in our perceptual beliefs. For, even if we might be dreaming, it does not follow that we might be thinking we see a hand on this same experiential basis, without seeing any hand. After all, in dreaming there is no real thinking and no real experiencing. So even if I had now been dreaming, which might easily enough have happened, I would not thereby have been thinking that I saw a hand, based on a corresponding phenomenal experience.¹⁵

In conclusion, we can argue not only that *I think, therefore I am* but also that *I think, therefore I am awake* (or at least not just dreaming). And that disposes of the special threat allegedly posed by dreams to the safety of our beliefs. But there are other varieties of skepticism to consider, to which we turn in the next chapter.

¹⁴This is not to say that there are no important intrinsic and relational differences between a realistic dream and a correlative stretch of waking life. It is only to say that a dream can be very realistic and that in the dream we take the goings-on to be certainly real, which leads naturally to the thought that "this," referring to the contents of one's present waking consciousness, insofar as one takes notice of them, could all be (the contents of) a dream. One could, of course, protest that though *in the dream* one is taken in, this does not show that one's *waking* consciousness could mislead in that way. The topic is far from exhausted, however; my *A Virtue Epistemology*, chapter 1, contains a fuller treatment.

¹⁵Rational intuition has a similar epistemic structure: one could be intuitively justified in believing that *p* even though in another situation not distinguishable in any relevant subjectively accessible respect one still would not be intuitively justified in believing that *p*. Though subtly different from our conclusion about the *cogito* and dreams, the two are closely related. In both cases, one is apparently in relevantly indistinguishable situations, though epistemically justified in only one of them. For supporting reasoning, see my "Skepticism and the Internal/External Divide," in *The Blackwell Guide to Epistemology*, ed. J. Greco and E. Sosa, 145–58 (Cambridge, MA: Blackwell 1999).

APPENDIX (MORE ADVANCED MATERIAL)

Two objections deserve our attention, finally, and each will receive a reply. Each reply will appeal to one main fact: to the fact that, in epistemic meditation, the focus is on *knowledge* and *inquiry* (as in the case of the ancient Pyrrhonists, who already engaged in *skepsis*). It is *epistemic* rationality that is then relevant, which requires no defense of our intellectual procedure *by comparison with playing tennis instead*. Indeed, on occasion it might well be more rational for us to opt for tennis, and perhaps (in part for that reason) *irrational overall* to affirm that p. But even then it might still be perfectly *rational epistemically* to so affirm.¹⁶

First objection:

We do not come to believe the contents of our dreams when we dream, so we do not affirm our dream events while dreaming. But we do engage in another activity that is in some ways similar to affirmation: we come to imagine. I shall call this activity "quasi-affirmation." Quasi-affirmation is not affirmation, but it is in many ways similar to affirmation, just as imagination is in many ways similar to belief. From an internal point of view, for the dreamer quasi-affirmation is importantly like affirmation and indistinguishable from it. One cannot, then, rationally affirm that one is awake, ignoring the possibility that one is quasi-affirming something false.¹⁷

Reply:

When I affirm that p, there are infinitely many possibilities that I perforce "ignore" in the sense of not adverting to them, not taking

¹⁶I am trying to distinguish between two assessments of rationality. It might be irrational to consider a question at all if one should instead be focused on a game of tennis, which was (perhaps obviously) an objectively preferable way for one to spend the afternoon. But this would bear not at all on the *epistemic* rationality of one's answering the question addressed as one does.

¹⁷Jonathan Ichikawa, in his "Scepticism and the Imagination Model of Dreaming," *Philosophical Quarterly* (2008): 519–27; quote on 523.

them consciously or in any way *into account*. Among these are all sorts of possibilities as to what my situation might be or might have been. What distinguishes the possibility that I might be quasi-affirming that I am awake when I affirm that I am awake? Why does *it*, in particular, deserve special attention. Well, I might say, "That possibility is, let us suppose, content-indistinguishable from my situation when I affirm that I am awake." Actually, more strictly, what matters as I consider whether to affirm at a certain time is whether my situation is *at that time* content-indistinguishable from my quasi-affirming (at the moment I make my decision on how to proceed or at the moment that I do intentionally so proceed). And let us suppose that indeed it will be thus indistinguishable.

Why, again, does that deserve special attention? How would such indistinguishability have a distinctive bearing on whether I should affirm at that time that I am awake? How would it matter *epistemically*? Presumably it would matter because it would be a way of going wrong, of having a false belief. The phenomenal indistinguishability of the two situations would bear that way. It would make the following two cases indistinguishable to one's internal perspective: (a) *the case of affirming that one is awake while one is indeed awake* and (b) *the case of quasi-affirming that one is awake while one is only dreaming very realistically that one is awake*. But why should that matter? Why, if such indistinguishability will *not* matter for the correctness of one's belief when one does affirm that one is awake? Why, *especially* if one *knows* that such indistinguishability *could not matter that way*?

What you know in such a case is that if you affirm that you are awake, that will be experientially content-indistinguishable from quasi-affirming that you are awake (although you are only dreaming). On our shared assumption that the imagination model is true and known by us to be true, however, you *also* know that if you do affirm, then it will be *false* that you only quasi-affirm while dreaming. So, why should the noted indistinguishability stop you from rationally affirming in that case? After all, you know that if you *do* affirm, then you *do not* merely quasi-affirm. Indeed, a way of *preventing* yourself from merely quasi-affirming in a certain act is to go ahead and affirm in that act.

Second objection:

Dream-shadow believing . . . is like believing, except that it happens only in dreams. If dream-shadow believing is possible, then if I dream-shadow believe that I am dreaming, necessarily I am correct; if I dream-shadow believe that I am awake, necessarily I am wrong. The first is self-verifying, the second self-defeating. *The skeptic can now ask: Should I try to form the belief that I am awake or instead the dream-shadow belief that I am dreaming?*—and to this question, Sosa’s argument gives no answer.¹⁸

Reply:

It is true that our argument gives no answer to that question. But compare the following questions: “Should I try to form the belief that I am awake, or should I play tennis instead?” “Which of *these* options is it rational for me to choose?” Our argument does not answer these questions either. And that is because these are questions of how to live, of how to proceed in one’s life at this juncture, whereas the theory of knowledge gives no answer to such questions. The theory of knowledge is more properly concerned with questions of knowledge and inquiry. Belief is thus within its purview, whereas imagination is not.

At least, imagination is not relevant in the same way. It may still be relevant in some ways. It may be relevant instrumentally, for example, in the sort of way a good diet and getting enough sleep might be. True, imagination is much more closely associated with epistemology than are those other factors. But in the following respect it is relevantly similar nonetheless. By imagining we do not *thereby* know. Imagining can be very important in presenting hypotheses to consider. However, unlike believing, it is not *constitutive* of knowledge. So it is not relevantly an alternative to believing in the endeavor of answering knowingly a question whether *p*, as we inquire into that question.

¹⁸This objection is quoted from Eric Schwitzgebel (in his blog, *The Splintered Mind*, for December 5, 2013). Available at <http://schwitzsplinters.blogspot.com/2013/12/dream-skepticism-and-phenomenal-shadow.html>.

So, suppose that our question is *whether we are awake and not just dreaming*. Our three options here are *belief, disbelief, and suspension of belief*. Our inquiry is directed at choosing among *these* options. Imagining (quasi-affirming and dream-shadow believing) are not directly relevant. And, *if given those three options*, it looks as if we cannot relevantly go wrong if we opt for the first of them. In this respect our question is just like the Cartesian questions as to whether one exists and whether one thinks at a given time *t*. Finally, the affirmative answer to *our* question is similarly well founded, or so it is on the imagination model of dreaming.

FURTHER READING

- Blumenfeld, David, and Jean Beer Blumenfeld. "Can I Know I Am Not Dreaming?" In *Descartes: Critical and Interpretive Essays*, ed. Michael Hooker, 234–55. Baltimore: Johns Hopkins University Press, 1978.
- Ichikawa, Jonathan Jenkins. "Skepticism and the Imagination Model of Dreaming." *Philosophical Quarterly* 58 (2008): 519–27.

Regress Skepticism

SAFETY-BASED AND SENSITIVITY-BASED skepticism deserve answers. Fortunately, as we have seen, answers are available.¹ But that leaves standing a traditionally important variety of skepticism.

A. SKEPTICISM AND CIRCULARITY

1. Two main strands of thought can be distinguished in the dialectic on philosophical skepticism. I focus here on what I take to be the more fundamental of these two. Once we see how to handle that one, it should have clear implications for how to handle the other one, and for how the two responses should be interwoven for a more satisfactory epistemology.

The lesser strand concerns the problem of the external world and how that relates to a particular sort of externalism in epistemology, naturalist externalism, which has taken various forms in the epistemology of recent years, including causal, tracking, process, and virtue varieties.² This is the worry that we put aside here, in the expectation that the position to be defended concerning the other worry will help with this one as well. Here we consider this lesser worry no further, not in its own right.

¹Ordinary perceptual beliefs might retain their status as knowledge, moreover, even if dreams *did* mislead us into false perceptual beliefs. This much will become clear when we turn at greater length to the question of what exactly is required for knowing that *p*, starting with chapter 4.

²See the following entries in the *Stanford Encyclopedia of Philosophy*: “epistemology,” by M. Steup; “knowledge,” by J. Ichikawa and M. Steup; “reliabilist epistemology,” by B. Beddor and A. Goldman; and “virtue epistemology,” by J. Greco and J. Turri. Available at <http://plato.stanford.edu/contents.html>.

The more important worry concerns a much more general epistemological and skeptical problematic found in Pyrrhonism and also important in Descartes's epistemology. This will be our main focus.

2. Consider, to begin, the following passage near the beginning of Descartes's Meditation Three:

I am certain that I am a thinking thing. But do I not therefore also know what is required for me to be certain of anything? Surely in this first instance of knowledge, there is nothing but a certain clear and distinct perception of what I affirm. Yet this would hardly be enough to render me certain of the truth of a thing, if it could ever happen that something that I perceived so clearly and distinctly were false. And thus I now seem able to posit as a general rule that everything I very clearly and distinctly perceive is true.

Note the sequence here. Descartes begins with the attribution to himself of certainty that he is a thinking being. And he then reasons as follows:

- (a) I am certain that I am a thinking being.
- (b) Nothing gives me this certainty but the clear and distinct perception that I am a thinking being.
- (c) But such clear and distinct perception could provide such certainty only if nothing could ever be that clear and distinct while false.
- (d) Hence, nothing could ever be that clear and distinct while false.

Here premise c is an unequivocal commitment to a strong reliabilism.³ With that premise Descartes makes evident his belief in

³This is so in a generic understanding of "reliabilism" as a view that makes the epistemic normative status of a belief dependent on its likelihood of being true given some feature of it. Of course, if that feature is just its truth, then that form of reliabilism is laughably inadequate. More considerable alternatives would include its being a manifestation of an epistemic competence and its being the yield of a certain psychological diachronic process. On this understanding of reliabilism, Descartes would qualify given his focus on judgments that accord with the subject's clear and distinct perception. So his reliabilist ground for the epistemically

a strong connection between normative epistemic status (such as *certainty*) and truth-reliability.⁴

Immediately thereafter he still wonders how he can be sure that his acceptance of premise d is really true, which sets off the theological inquiries leading to his positive perspective on clear and distinct perception: that is, on the direct clear and distinct perception constituted by intuition and on the indirect clear and distinct perception provided by deduction from the intuitive.

Not only is Descartes a reliabilist, however. He is also a virtue epistemologist among whose main epistemological concepts is that of the apt belief or judgment. And he is, moreover, a bilevel virtue epistemologist, with an animal level of *cognitio* and a reflective level of *scientia*.⁵

3. Descartes does, of course, tackle the external-world epistemological problematic. That is not his deepest skeptical and

normative status of a judgment is the judgment's corresponding to the subject's clear and distinct perception of its content.

⁴Objection: *Descartes is importantly different from contemporary "strong" reliabilists in not trying to give sufficient conditions for justification in reliabilist terms. "Strong reliabilism" could mean the reliabilism that tries to give such necessary and sufficient conditions without appealing to anything internalist.* Reply: I don't think the process reliabilist would allow just any process. Any relevant process will need to be a cognitive process (though it can also be distal). In addition, there are other requirements that reliabilists impose, such as that there be no other countervailing process. Also, a kind of "internal" factor is perhaps brought in when we are told that the relevant process must be a *cognitive* process, a psychological process of some sort.

Moreover, it is not clear that Descartes is not an externalist, too. After all, he wants to distinguish real intuitions and real deductively proved conclusions (real direct clarity and distinctness, direct and intuitive or indirect and deductive) from pseudo-intuitions. And when the subject matter is external (as in the case of the existence of God), then we have a kind of externalism. So I myself think that he was a complete, infallibilist, strong reliabilist.

⁵This was argued in chapter 1, and it is argued elsewhere in further detail: first, in "Descartes and Virtue Epistemology," in *Reason, Metaphysics, and Mind: New Essays on the Philosophy of Alvin Plantinga*, a volume in honor of Alvin Plantinga edited by Kelly Clark and Michael Rea, (Oxford, UK: Oxford University Press, 2013), and second, in "Descartes's Epistemology," in the proceedings of an Arché conference, *Scepticism and Perceptual Justification*, edited by Dylan Dodd and Elia Zardini (Oxford, UK: Oxford University Press, 2014).

epistemological concern, however, because the skeptical doubt that he takes up most seriously extends far beyond external-world concerns about fires and dreaming to cover even the simplest truths of arithmetic and geometry. What follows focuses on this deeper problematic, which concerns not just sense perception but even intuition and deduction, the faculties of main interest to Descartes in his quest for certainty.

So let us join Descartes in his radical internalist virtue reliabilism of intuition and deduction. After he has reflected in the early meditations and reviewed various radical scenarios, by the beginning of his Meditation Three he is prepared to self-attribute certainty that he is a thinking being, which he does in the passage quoted above. And he notes that, as far as he can see, what gives him this certainty is simply clarity and distinctness.

4. Descartes believes in a faculty of intuition that delivers certainty through direct clear and distinct perception of basic truths independent of any inference. But he is willing to put even this faculty in question as part of his general question as to what might assure him that nothing could possibly reach a certain level of clarity and distinctness—the level reached by his knowledge that he is a thinking being—without being true. And now what could stop him from reasoning as follows?

Here in this first instance is something I intuit to be true and *is* true.

Here in this second instance is something I intuit to be true and *is* true. Et cetera.

The best explanation for all this success is that intuition is a supremely reliable source of truth.

(In each instance here, the truth claim derives simply from intuition.)

Descartes did, after all, endorse inference to the best explanation.⁶ And he might thus have adopted this easy inference in which each

⁶See his *Principles of Philosophy*, Principle 205. (This would, of course, require going beyond intuition and deduction strictly conceived; but there are other

of the premises is delivered (at least in essential part) through an exercise of intuition, reaching thereby the conclusion that intuition is reliable. Thus he would have been able to arrive at easy knowledge of the reliability of his intuition, of his clear and distinct direct perception of truth.

That is one way in which a problem of bootstrapping affects even a Cartesian epistemologist.

5. Moreover, a further issue of bootstrapping is faced by Descartes. As we have seen, he seems to reason as follows:

- (a) If I (now) think, then (necessarily) I exist.
- (b) I (now) think that I exist.

Therefore,

- (c) My (present) thought that I exist cannot possibly be false (and I cannot possibly be deceived in so thinking).⁷

If so, why could he not have reasoned differently in Meditation One? I refer to his worry that he might be going wrong even about the simplest questions of geometry or arithmetic, such as

ways in which Descartes is willing to be thus flexible, as when at the end of the *Meditations* he reasons that it would be laughable to be fooled by dreams, which evidently differ from waking life with respect to a kind of diachronic coherence that we can presumably track with the aid of our memory.)

⁷Objection: *But it is not at all clear that Descartes does reason that way. After all, he says in the Replies to the Second Set of Objections that his knowledge that he exists is not derived from the general proposition that whatever thinks exists. Rather, the Meditator recognizes [the cogito] as something self-evident by a simple intuition of the mind. The general proposition is itself derived from the cogito, rather than the other way around. (See CSM 100.) If the cogito is thus self-evident through sheer intuition, moreover, then it is not derived inferentially from the conditional claim that if one thinks, then one exists.* Reply: That is all compatible with my suggestion, however, because the Meditator might (a) recognize the *cogito* (as self-evident), and (b) even recognize that it is self-evident, by directly intuiting each of these (the *cogito* itself and also its being self-evident), while yet being able to also (c) reason his way to a meta-belief that it is infallible, where this latter also bears on its epistemic status, by further defending it reflectively. In effect, the Meditator's *cognitio* that he exists could be just intuitive (and might attain the highest level of certainty as *cognitio*), even though his *scientia* that he exists must go beyond that by relying also on the sort of reflective defense provided by the reasoning early in the *Meditations*.

what 3 and 2 add up to. Why could he not have reasoned as follows, already in that meditation, without need of any deep theological ponderings:

(a') $3 + 2 = 5$ (necessarily so).

(b') I (now) think that $3 + 2 = 5$.

Therefore,

(c') My (present) thought that $3 + 2 = 5$ cannot possibly be false (and I cannot possibly be deceived in so thinking).

Yet this primed reasoning, a'-b'-c', seems viciously circular, unlike the earlier a-b-c. The primed argument amounts to "blatant bootstrapping," which adds no significant epistemic support for its conclusion. By contrast, Descartes's cogito reasoning (as in the unprimed a-b-c) does not seem *similarly* suspect. Descartes may favor cogito thoughts in his Meditation Two because those are thoughts that he can defend against skeptical doubts.⁸ He can defend them by appeal to the general consideration that they cannot be believed incorrectly: their very content guarantees that when affirmed they must be true. (That they are also indubitable is something he stresses in the *Replies to the Second Set of Objections*.⁹) And this is a premise we can affirm *without reasoning from the truth (or the necessary truth) of what is affirmed*, which distinguishes Descartes's cogito reasoning from the blatant bootstrapping of our earlier primed reasoning.

That gives us a way to understand how Descartes can favor the cogito as he does in the *Meditations* and why he could not plausibly have defended his acceptance of simple math propositions in the way he defends his acceptance of the cogito. The difference is

⁸Objection: *But how could the first premise, premise a, be defended against skeptical doubts? And, if it can't be so defended, how could the conclusion be defended? On the other hand, if the conclusion can be defended against skeptical doubts but the premise cannot, doesn't that give us reason to think that the certainty of the conclusion is not actually grounded in the premise or the argument?* Reply: As argued earlier, there can be a defense of the cogito belief by appeal to given premises, even if the epistemic status that it derives from this defense is not that of a first-order conclusion based on first-order inference from those premises.

⁹CSM 104.

supposed to be that the defense of the simple math propositions in accordance with the argument forms given earlier would involve blatant bootstrapping. It would involve assuming as a premise the truth or the necessary truth of the conclusion reached in the argument. This is not true of his defense of the cogito.¹⁰

One might thus (with Descartes) entertain the question whether one's armchair faculties are reliable enough to deliver certainty. And here the well-known antinomies and paradoxes of reason give rise to the worry that armchair reason itself is subject to cognitive illusion. So one might hope to attain some perspective on one's own faculties of reason—such as intuition and deduction—some perspective that might assure us that they are indeed reliable enough to deserve our trust. And that is precisely Descartes's situation as the *Meditations* unfold. How, now, might one go about seeking such assurance? Descartes does not hesitate to put his best and most basic faculties to work in the development (through rational theology) of such a reassuring perspective.

Call such first-order knowledge *animal knowledge*. This would be knowledge unsupported by higher-order attitudes on the part of the subject. Call the higher-level knowledge gained through a reassuring perspective *reflective knowledge*. By acknowledging *both* kinds of knowledge we subscribe to a "bilevel epistemology." This bilevel approach allows the use of our basic foundational faculties in attaining a second-order assuring perspective. So we can use the animal knowledge that we attain through the exercise of such faculties; we can use such animal knowledge in the (proper, coherence-aimed) elaboration of the endorsing perspective. This endorsing perspective would be a proper awareness of

¹⁰ Compare the following reasoning: (a1) *I exist*; (b1) *I (now) think that I exist*; therefore, (c1) *my (present) thought that I exist is true*. This reasoning, too, seems viciously circular, as does the primed reasoning in the text. It, too, seems a case of blatant bootstrapping. But the initial unprimed reasoning in the text seems better, as does the following unmodalized correlate of it: (a2) *Whatever thinks exists*; (b2) *I (now) think that I exist*; therefore, (c2) *my (present) thought that I exist is true*. Neither the original a-b-c reasoning, nor its correlate a2-b2-c2, appeals to a premise that simply asserts as true (or as necessarily true) the very content of the thought said to be true (or necessarily true) in the conclusion.

our competences through whose exercise we can gain our first-order knowledge.

That whole approach has been rejected as follows:

Either the coherence of a belief with other beliefs is *by itself* intellectually satisfying, or it is intellectually satisfying only if accompanied by the prior animal-level justification. . . . [The latter? Faced with that suggestion] . . . the skeptic is bewildered. The skeptic imagines someone who has the very same beliefs that we have about appearances but who couples those beliefs with . . . skeptical hypotheses. . . . The skeptic wants to know why the beliefs we have are any more likely to be true than beliefs embedded in one of these alternative belief systems. . . . Any answer that appeals to the purely external virtues of animal-level justification enjoyed by one set of beliefs will strike the skeptic as blatantly question begging and unable to give us the relevant intellectually satisfying justification.¹¹

But this accusation of *petitio*, of begging the question, would be no less in order against the Cartesian approach than against the broader bilevel, animal or reflective contemporary approach. For these are really at bottom the *same* approach, the same bilevel approach. It is just that Descartes restricts himself to armchair faculties, whereas we more broadly do not. But the structural issues are exactly the same. If the circularity-wielding opponents of the Cartesian virtue-epistemological approach are wrong, then the *petitio*-wielding opponents of the broader animal or reflective virtue-epistemological approach are wrong for the same reasons.

One might hold out for an intellectual reassurance that will *not* be subject to the easy-knowledge (circularity, begging the question) sorts of objections. One might thus hold out for a "legitimizing account of our knowledge in general." This would demand an account of our knowledge in general that lays out justifying reasons for every bit of our knowledge without circularity or infinite

¹¹ Richard Fumerton, "Reflective Knowledge and Intellectual Assurance," *International Journal for the Study of Skepticism* 1 (2011): 113–23; quote on 122 (emphasis added).

regress. Because it quickly becomes obvious that this is beyond human limitations, and indeed beyond the Almighty himself, the demand is obviously “incoherent.” This means not just that it is unfulfillable but that it is unfulfillable with metaphysical necessity. A further incoherence involves the knowledge that the aim is obviously impossible, together with the demand or desire for its realization nonetheless. These two attitudes do not cohere properly. One cannot coherently desire that *p* while knowing that it is obviously *impossible* that *p*.

We need not insist on the terminology of “coherence,” however, because the targeted combination is in any case *bad*. It is surely bad to want something that not even God could attain, especially when the impossibility becomes obvious. There is now a strong reason to try to escape that bad situation. But there is not much to be done about the impossibility. One must instead try to get rid of the desire.

Conceivably the desire in question, or even its combination with the impossibility, might have excellent consequences. If so, this might then trump our reason for getting rid of it: the reason, namely, that this desire will be *inevitably* frustrated, obviously so. Because we can foresee no such excellent consequences, however, nothing like that speaks in favor of retaining any desire for a legitimating account of our knowledge, and it still seems clear that we should try to get rid of that desire (if we have it) and put our time to better use.

6. To abandon the search for a general “legitimating” account is not to give up altogether on the project of developing a general epistemology. We must be clear, however, on the sort of epistemology desired. It cannot be a *methodology*, a set of rules or principles by following which we would attain *any* knowledge that is attainable. According to such a “methodology,” it would be by following such instructions that we would always gain the status of justification required for knowledge.

That seems little better than the search for a legitimating account, and indeed the two ill-conceived objectives are closely related. Following a rule would require that we ascertain the obtaining of the conditions laid down in the antecedent of the rule.

But this would require that we be justified in believing that those conditions obtain. So there must have been some nonposterior application of a rule. And once again there is no way to follow any such set of rules without vicious circularity or regress so as to attain *all* of our justified belief.

If we still seek a general epistemology, then, it cannot be a general methodology of that sort. But it might still be a general explanatory account of our knowledge, an account that would be able to explain, for any bit of knowledge, how it is that the constitutive belief comes to have the epistemic normative status required for knowledge. If we call that normative status “(epistemic) justification,” the objective would be a general explanatory account of such justification.

Once we are clear on this, consider now the circularity that such a successful account would inevitably involve. A successful account would be one through which we epistemologists would succeed in understanding human knowledge in general. We would presumably have a set of basic competences, some subset of which we would use in coming to know our general account of how we attain our knowledge. This general account would then have to give us an understanding of such competences and their reliability. So the competences by means of which we would have come to know our account would be among the competences that our account invokes, and so they would be among the competences that our account must claim to be reliable. There is no vice in this sort of circularity.

Compare the circularity that is involved in coming to understand how our faculty of vision gives us reliable access to our surroundings. In arriving at the relevant theory about the specifics of our faculty of vision we will presumably be using our eyes to gather relevant data based on which we will come to know about the optic nerve, the structure of our eyes, and so on, so as to have a specific theory that explains how it is that vision gives us reliable access to the shapes and colors of objects around us. In reliably arriving at that theory we will have exercised the very faculty whose reliability is explained by the theory. But no problem of vicious circularity is here discernible.

A quite general theory of knowledge would be like this theory of visual knowledge in the relevant respects. We would exercise

various epistemic faculties in arriving at a theory concerning our nature and surroundings based on which we could see those very faculties to be reliable. The theory would presumably go into more or less detail in explaining the constitution of these faculties and of how they fit us to gain access to relevant domains of facts. Here we would just have a more general example of the same phenomenon of circularity that we found in the specific case of visual knowledge. And here again no vice in the circularity is clearly inevitable for the general theory by means of which we would gain the relevant general understanding. Moreover, we are not only armchair thinkers but also knowers of the world around us. And our knowledge of that world is as dependent on our sensory perception as is our knowledge of the abstract realm of reason dependent on our direct intuition. Any telling charge of vicious circularity applicable to our epistemological reliance on perception would apply equally to Descartes's epistemological reliance on intuition. If we are denied a philosophically desirable assurance by that circularity, then so would Descartes be denied such assurance by a precisely parallel circularity. Moreover, this would have been so *even if every point of theology he relied on had been perfectly true and even if he had correctly intuited every truth that he thought himself correctly to intuit as he meditated*. If, in such circumstances, Descartes *would still have fallen short* of proper philosophical assurance, it is hard to see what this desideratum amounts to or why we should really care to attain it.

B. CONCLUSION

In conclusion, let us draw a distinction among competences, between those that are and those that are not fundamental. Thus, our visual ability to tell that the bells are tolling can gain support from our auditory ability to tell that they are tolling, and vice versa. Although it is quite doubtful, *suppose* that we could gain from circular, blatant bootstrapping *some* support for our ability to tell the contents of our fuel tank by just reading our fuel gauge. *Even so*, we would still be far better off if we could test the gauge independently, by direct perception of the contents of the tank or through some other means.

Some practices or faculties or competences, by contrast, are *not* checkable independently. Which are these? They are practices or faculties or competences that are “fundamental.” Is sense perception one such? Not clearly. Again, hearing is testable through vision, and vice versa. Plausibly, then, sense perception, in *each* of its modes, *can* perhaps after all be checked independently. In another way, however, it cannot. This may be seen as follows. Assume that we have a set of perceptual faculties each of which is checkable through one or more of the others, although none is checkable in any other way. In that case, perception as a whole is not checkable independently of perception (of some sort or other).

We can hardly object to the use of perception on that basis, however, nor can we plausibly hold on that basis that there is no way for us to properly judge, concerning any variety of perception, whether or not it is reliable. After all, there is bound to be some total set of faculties or practices none of which is checkable except through the use of one or more of the others. This total set, which we can call “fundamental human cognition,” is then *not* checkable independently from “fundamental human cognition.” But we can hardly object *on that basis* to the use of any fundamental human cognition, nor can we plausibly conclude that no fundamental human cognitive practice or faculty could ever be properly judged reliable.

It would, of course, be an accomplishment to reduce perceptual cognition to more fundamental cognition, say to armchair cognition, including introspection and pure reason. However, as Descartes saw so clearly, armchair cognition is not properly endorsable except through the operation of armchair cognition (perhaps in a priori theology). So it is not reducible to something more fundamental. The objection to a mode of cognition, then, cannot possibly be that it is not endorsable except through its own operation. What, then, would be the proper response to (a) our perceived inability to endorse even armchair cognition that way, along with (b) our perceived inability to endorse empirical, scientific cognition through the operation of armchair cognition? One possible response would be to reject empirical, scientific cognition as groundless. But once we can see no hope of further grounding

for armchair cognition itself, once we suffer an anti-Cartesian disillusionment, it seems just arbitrary to hold onto armchair cognition while rejecting empirical, scientific cognition wholesale. It is crucial, however, to be clear on why we see no further grounding for armchair cognition itself. It is *not* because the Cartesian Circle is inevitably circular, *viciously* so. It is only because the theological reasoning on which Descartes rests his case is unacceptable, and we see little hope for any acceptable substitute if we are bound to the armchair. Given this, we might still reasonably hold out for an endorsing perspective that we can access through a more inclusive fundamental human cognition, a perspective that will enable us to endorse inclusive human cognition, both armchair and empirical. This perspective is, of course, inaccessible through a single flash of insight. If it is to a sufficient extent accessible at all, it will be so only through the patient study of human cognition that joins together the efforts of the special sciences and, so we can hope and expect, the insights of a reasonable epistemology.

APPENDIX (MORE ADVANCED MATERIAL)

Objection:

Isn't there a fundamental skeptical problem that is local to perception (and memory, given the Russell hypothesis that the world was actually created recently with all the ostensible causal traces in place, and induction, given Hume's reasoning)? I have in mind underdetermination-based arguments for skepticism. There is no strong analogue of these arguments in domains where it is more plausible to collapse appearance and reality (e.g., the conscious mental realm). This is not regress skepticism.

Of course, these arguments do assume that justification requires evidence. Otherwise it needn't matter that there is evidential underdetermination. Is it denied that evidence is thus required? Is this why the problem is being set aside?

That seems a bit too fast. Not just any externalist will think that we can set it aside. The only one who will think that is one who denies that good evidence is even necessary for justification. To affirm that it is necessary is compatible with denying that it is suffi-

cient. That second denial is the core anti-evidentialist claim, not the stronger claim that evidence is not so much as necessary.

Reply:

For Descartes, a high enough level of pure adroitness is enough to provide the certainty enjoyed by the cogito. For consider his procedure. He goes through the various skeptical scenarios and in the end rules that the cogito passes muster. And what gives it that status? Well, in part it is the infallible correctness of affirming it. And this is a status that the affirmations of the atheist mathematician will share. Without theology, Descartes thinks, the mathematician will not be able to endorse reflectively his affirmations and his whole proof procedure. But the mathematician does have some considerable positive epistemic status for his affirmations as he goes through the proof. This is, for Descartes, the status provided by sheer *clarity and distinctness*.

But what is involved in the clarity and distinctness of *sum* or of *cogito*? Descartes does not seem to attend in his early meditations to any *prejudgment seeming* that will back up the epistemic status of *cogito* propositions, such as <cogito> or <sum> or <since I think, I am>.¹² All he needs is that such *cogito* beliefs be inherently infallible and that the competence to attain truth by such affirmation be an infallible competence. *This* seems to be all that he defends in his sequence of rebuttals of the various skeptical scenarios, except only for the fourth, where he ascends to the second order in considering the *origin* of his beliefs and competences. All that he defends in those three earlier scenarios is that he could not possibly go wrong in affirming the relevant cogito propositional contents, such as that he thinks and that he exists. He never adverts to any prejudgment seeming or experience.

So all that Descartes has at the end of those meditations is that cogito beliefs have a certain epistemic status. What status, exactly? Well, at a minimum, the status of infallibility. But, arguably, also the status of a minimal sort of proper reflective endorsability. Either way, there is no appeal to prejudgment reasons or evidence. No such rational basis plays a role that I can discern

in the early meditations, as Descartes reflects on his sequence of skeptical scenarios. He does, of course, already, even at the start of Meditation Three, have a rational basis for reflective endorsement. That is the result won in his struggle with the skeptic in the earlier meditations. And of course he is going to add to *that* reflective basis for endorsement. He is going to do that with the theology to follow in his later meditations.

Consider the sort of view we can reasonably take on the epistemology of the simplest logic, arithmetic, and geometry and on the epistemology of judgments like <It has been more than five seconds since I awoke>. In my view, the seemings involved here are strong enough to yield resultant seemings of great intensity, which lead to judgments with a high level of epistemic standing. And those seemings also seem evaluable epistemically in a positive way *requiring* no pre-seeming rational basis. Here I have in mind the initial seemings, not the resultant ones, since the latter are presumably based on the former.¹⁵ Furthermore, no epistemic divide seems to separate these cases from cases of perceptual attitudes based on perceptual experiences (based, therefore, on the sorts of experiences that are absent in the cases of math, logic, etc.). The reason there is no such divide is that the perceptual experiences in question are themselves evaluable, being states that can and do function with epistemic propriety. They are functional epistemic states that can manifest proper epistemic function or can fail to do so.

Anyhow, what we find with *cogito* propositions is now something a bit stronger, something that goes further along a preferable line. How so? Well, if we reason along with Descartes, it looks as if we need not invoke any prejudgment seemings at all. We can focus just on the act of affirmation itself; we can, with Descartes, take note of the fact that skeptical scenarios do not threaten the

¹⁵ An *initial* seeming is the attraction to assent that derives from some particular factor, such as the testimony of a friend or a visual appearance, as in a perceptual illusion, whereas the *resultant* seeming is the attraction to assent on that same question that derives from the operation of *all* the relevant factors that bear for one at that time on that question. In the Müller-Lyer case that would include not just the visual appearance but also the careful measurement whose result is retained in memory, the testimony of others, and so on.

infallibility of these affirmations, given just their content. We can thus see that there is a competence at work in so affirming, an infallible one at that. And we can now even reflectively endorse the functioning of that competence, in general and in specific cases, to the effect of a reflectively endorsed judgment that is apt and now also fully apt.

FURTHER READING

- Alston, William P. "Epistemic Circularity." *Philosophy and Phenomenological Research* 47 (1986): 1–30. Reprinted in Alston, *Epistemic Justification: Essays in the Theory of Knowledge*, 319–49. Ithaca, NY: Cornell University Press, 1989.
- Bergmann, Michael. "Epistemic Circularity: Malignant and Benign." *Philosophy and Phenomenological Research* 69 (2004): 709–27.
- Cohen, Stewart. "Basic Knowledge and the Problem of Easy Knowledge." *Philosophy and Phenomenological Research* 65 (2002): 309–29.
- Fumerton, Richard. *Metaepistemology and Skepticism*. Lanham, MD: Rowman and Littlefield, 1995), chapter 6.
- Vogel, J. "Epistemic Bootstrapping." *Journal of Philosophy* 105 (2008): 518–39.
- . "Reliabilism Leveled." *Journal of Philosophy* 97 (2000): 602–23.

Knowledge: What It Is and How We Might Have It

ACCORDING TO THE foregoing chapters, we can and do enjoy knowledge despite the protestations of skeptics. But what is this knowledge? At a minimum, how might we determine the conditions necessary and sufficient for its attainment?

A. WHAT IS KNOWLEDGE?

1. A gambler rolling dice might be persuaded that the dice will come up 7 or 11 on his next toss, and he might be right. But he still would not know. What *more* would he need in order to know, besides a belief that is true?

A reason, perhaps, a good reason?

At noon you may think it is noon based on reading a clock, but suppose the clock happened to stop at noon days ago. Even if you have excellent reason to trust that clock, you could hardly know the time that way.

You may believe that someone here owns a Ford because you have excellent reason to think that Nogot here owns one. And someone here does indeed own a Ford, namely Havit, not Nogot.

Out for a ride in the country, you see a barn façade and of course take it that the structure you see is a real barn. However, nearly all the façades in the area are *mere* façades, part of a movie set.¹

In all three of those cases, you have a true belief, and you have (by hypothesis) excellent reason for so believing, but you do not

¹ According to the dictionary, a façade can be a *mere* façade with no building behind it, but it can also be just a side of a real building.

know.² Even a true belief based on excellent reason may thus fall short of knowledge.

2. Perhaps we need not a true belief that has a good rational basis but rather one that is *sensitively* true.³ We thus return to an idea broached in chapter 1, to which we aimed to return. Here, then, is the new proposal:

Sensitivity Condition: In order for a true belief to be a case of knowledge, it must be *sensitive*: it must be one that the believer *would not* hold if its content were false rather than true.

This requirement is violated by the true believers in our four examples above. All four of those beliefs are *insensitive*. The gambler would (or anyhow might well) still believe that the dice would come up 7 or 11 even if it were false that they would do so. The clock reader would (or might well) still believe that it was noon (from reading the stopped clock) even if he were reading the clock at a time other than noon (say at 11:50). You would (or might well) still believe that someone here owned a Ford, even if no one did (say, because not even Havit owned one or because Havit was not here after all). And you would (or might well) still believe that the structure you saw was a real barn even when that was not so (and you saw a *mere* barn façade).

What is more, there is a further case in which one fails to know despite getting it right with good justification, as one expects cor-

²For decades the three cases were lumped together, but we now distinguish the first two cases, distinguished as cases of “intervening luck,” from the third, a case of “environmental luck.” And it has now been argued that this difference makes an important epistemic difference. See, for example, my *A Virtue Epistemology* (Oxford, UK: Oxford University Press, 2007), Lectures 2 and 5.

³We thus turn directly to the sort of modal approach to knowledge that prepares the way for virtue epistemology. But there is a vast literature on the Gettier problematic that explores several other approaches. A historical account may be found in the article on the analysis of knowledge by Jonathan Jenkins Ichikawa and Matthias Steup in the *Stanford Encyclopedia of Philosophy*. Available at <http://plato.stanford.edu/contents.html>. I turn to this form of virtue epistemology in my *Knowledge in Perspective* (Cambridge, UK: Cambridge University Press, 1991), 277: “We have reached the view that knowledge is true belief out of intellectual virtue, belief that turns out right by reason of the virtue and not just by coincidence.”

rectly and justifiably to lose a lottery (one with astronomically long odds). Here again the subject plausibly believes based on a fine rational basis. Here again a true belief is widely assumed to fall short despite being based on such a rational basis. Moreover, here again the sensitivity condition ostensibly delivers the right result, because our subject's belief that his ticket will lose is insensitive: it would still be held by the believer, on the same probabilistic basis, even if, against all odds, as a matter of contingent fact the ticket would win.

3. We shall find that our sensitivity condition is open to serious objections. Fortunately, its undeniable intuitive attractiveness can be explained through an easily confused but far preferable notion of "safety." This was suggested already in chapter 2, and now we can deepen our treatment. The denouement of our reflections will be to show how replacing sensitivity with safety makes it possible to defend plain common sense against the spurious advantages over it claimed by alternative accounts. What follows supports that replacement by showing how the sensitivity requirement runs against serious problems.

First of all, suppose we require sensitivity for knowledge and we agree with the following: *that any conclusion reached through competent deduction from something known to be true is also known to be true.* Given that, consider notorious skeptical scenarios such as these:

One's envatted brain is fed stimulations that present a common-sense world of a familiar sort.

Or, alternatively:

An evil demon feeds one directly a natural course of common-sense experience.

Consider in some such case one's belief in the ordinary surroundings that would normally correspond to such a course of experience. Those are the beliefs one would then have regardless of the surroundings so long as the course of experience remained the same. Because of how they originate in experiences with such sources, those experience-based beliefs fail to be sensitive.

Take next the belief that one is *not* in any such skeptical scenario as that of the brain in a vat or that of the evil demon's victim. Take indeed *any* such belief, to the effect that one is not thus epistemically victimized. Any such belief would then be *insensitive* and thus disqualified by the sensitivity condition. Any belief that you are not victimized fails to constitute knowledge even when you have competently deduced that conclusion. If you were a victim, you would nonetheless believe that you were *not* a victim. Even if you had competently deduced that you were not a victim, you would still fail to know that you had not suffered that fate. For your belief that you had not suffered that fate would still be insensitive, even if you had competently deduced it from various ostensible pieces of knowledge, such as that you saw your hands or perceived a fire.

Here now is a most repugnant implication. Consider first this principle:

(DK) *Competent deduction transmits knowledge: anything believed as a conclusion based on competent deduction from known premises is also thereby known.*

Given that principle, we cannot know such things as that we see a hand or perceive a fire. For if we *could* know such things, then from the premise that we do see the hand or the fire we could competently deduce that we were not relevantly victimized, and, according to principle DK, we would be able to know that we were not so victimized. But, as we have just seen, our sensitivity condition precludes our being able to know that conclusion.

So we have a *reductio*: from (a) *the sensitivity condition*, (b) principle DK, and (c) the assumption that we are able to make such simple deductions, it follows that we cannot know the world around us.

However, this argument is *itself* best seen as a broader *reductio*, one that reduces to absurdity that sensitivity condition.⁴

⁴However, this ostensible *reductio* is by no means the only problem for sensitivity, because it is already troubling enough that it would rule out our ever knowing ourselves not to be misled. (This problem does not require closure as a premise.)

4. *Better Safe Than Sensitive.* The problems we have found to affect sensitivity do not affect a closely related “safety.” A belief is sensitive iff had it been false, S would not have held it, whereas a belief is *safe* iff S would not hold it without its being true.⁵

The requirement that a belief must be sensitive if it is to be “knowledge” is found to be *prima facie* plausible: as we have seen, it is found intuitively that the failure of a belief to be “knowledge” may be explained through the fact that the belief would remain in place even if false. The problem for this way of arguing is that an alternative explanation is equally adequate for undisputed cases. According to this alternative explanation, it is safety that “knowledge” requires, but this requirement is violated in the ordinary cases cited, wherein the subject fails to know.

One fails to know in those cases, it is now said, because one’s belief is not safe. Suppose this generalizes to all uncontentious cases adduced in favor of a sensitivity requirement. Suppose that in all such cases the condition required could just as well be safety as sensitivity. And suppose that, as argued earlier, the problems for sensitivity briefly noted do not affect safety. If so, then one cannot differentially favor sensitivity as the right requirement, in support of the skeptic’s main premise.

Here is the striking result: if we opt for safety as the right requirement, then a common-sense stance is defensible from the *sensitivity-based* objections, which enables us to escape that route to skepticism. That is to say, one does satisfy the requirement that one’s belief of not-H (that one was not so radically deceived) be safe: after all, not easily *would* one believe that not-H without its being true (which is not to say that not possibly *could* one

⁵For short, in symbols: S’s belief B(p) is sensitive iff $\neg p \rightarrow \neg B(p)$, whereas S’s belief is safe iff $B(p) \rightarrow p$. These are not equivalent, because subjunctive conditionals do not contrapose. Thus the following combination is possible (where the “main valve” is the valve through which water comes into the house from the public pipes outside): (a) it is true that if water flowed from the kitchen faucet, then it would not be so that water flowed from the kitchen faucet while the main valve was closed; (b) it is false that if it were so that water flowed from the kitchen faucet while the main valve was closed, then water would not flow from the kitchen faucet. Since this combination of truth values is possible, subjunctive conditionals do not generally contrapose.

believe that not-H without its being true). In the actual world, and for quite a distance away from the actual world, up to quite remote possible worlds, our belief that we are not radically deceived matches the fact as to whether we are or are not radically deceived.

5. *A Defense of Common Sense.* Granted, a skeptic's requirement of sensitivity is quite plausible, as we have seen. However, one *need not* explain plausibility in terms of *truth*. Many false things are plausible; we can explain why they are plausible without having to consider them true, and we can explain why other things are plausible by appeal to such illusion. We are said to face illusions at every turn, from the humble perceptual and cognitive illusions of interest to psychologists to the more momentous illusions alleged by Freud and Marx. In all such cases, illusion may be said to explain plausibility.⁶

Consider, moreover, the need to explain how the skeptic's premise—that one does not know oneself not to be radically misled, and so on—is as plausible as it is. That requirement must be balanced by an equally relevant and stringent requirement: namely, that one explain how that premise is as *implausible* as it is. To many of us it just does not seem so uniformly plausible that no-one can ever know that they are not envatted while being fed experiences as if they were normally situated. It just does not seem so uniformly plausible that one cannot be said correctly to know that one is not at this very moment being fed experiences while envatted. So the explanatory requirement is in fact rather more complex than might seem at first. And, given the distribution of intuitions here, the skeptic would still owe us an explanation.

Interestingly, our distinction between sensitivity and safety may help us meet the more complex explanatory demand, compatibly

⁶One might, however, prefer to view illusion as misbegotten plausibility, so that the plausibility is *constitutive* of the illusion, which therefore cannot explain it really; still, in all such cases of illusion it may be explained why something strikes us as plausible despite being false.

with our preferred common-sense stance. Here, now, is a sketch of our preferred explanation.

- (a) It is safety that is required for knowledge (and for its correct attribution), not sensitivity. It is required that $B(p) \rightarrow p$ and not that $\sim p \rightarrow \sim B(p)$.
- (b) Take our belief that we are not radically deceived as in a skeptical scenario. Because that belief *is* safe, the skeptic cannot argue for his distinctive premise by alleging that here we violate the safety requirement.
- (c) Safety and sensitivity, being mutual contrapositives, are easily confused, so it is easy to confuse the correct requirement of safety (for knowledge and its correct attribution) with a requirement of sensitivity. It is easy to overlook that subjunctive conditionals do not contrapose validly.
- (d) Those who find the skeptic's distinctive premise plausible *on the basis of sensitivity considerations* may thus be confusing sensitivity with safety and may, on that basis, assess as correct affirmations of that premise. After all, the requirement of safety is well supported by the sorts of considerations adduced generally by the sensitivity-based opposition to G.E. Moore's common-sense standpoint. Because sensitivity is so similar to safety, so easy to confuse with it, it is no surprise that one would find sensitivity so plausible, enough so to mislead one into assessing as correct affirmations of that premise.
- (e) The plausibility of the skeptic's sensitivity premise is thus explained compatibly with its falsity, which fits the stance of those who agree with Moore.

Thus may a Moorean epistemology defend itself against "sensitivity-based" objections. Such objections presuppose that in order to constitute knowledge a belief must first be "sensitive." We reject that requirement, and thereby support our preferred common-sense alternative.

Of course all we really need in order to explain the plausibility of the skeptic's premise is that it clearly enough follows from something plausible enough. And the sensitivity requirement fulfills that role well enough independently of whether it is confused with

a safety requirement. But that still leaves the question of why the requirement of sensitivity is so plausible if it is just false. And here there is still a role for safety if it can function as a plausible enough requirement, one both true and defensible through reflection and one that appeals to us simply through our ability to discern the true from the false in such a priori matters. Compatibly with that, some of us may be misled into accepting the requirement of sensitivity because it is so easily confused with the correct requirement, that of safety, thus succumbing to cognitive illusion.

Next we step back in order to consider the methodology we have implicitly followed up to now. In doing so, we shall also turn to the full account of knowledge that was once the consensus view in the analytic epistemology tradition. I mean the thesis that knowledge is tantamount to justified true belief. This was, in fact, the specific thesis targeted by Edmund Gettier with his famous refutation. We will consider the thesis, but our main focus will be on the methodology whose employment led to the conclusive refutation of that account.

B. THE METAPHYSICAL GETTIER PROBLEM

1. The English schema “S knows that p” has fascinated philosophers and linguists, who have sought an account of its meaning. Others have focused—alternatively, or in addition—on the concept (or concepts) expressed by that schema. Philosophers have also been interested in knowledge itself, however, in a way that goes beyond semantics and conceptual analysis.

2. Consider the following thesis:

JTB *Necessarily: S knows that p if and only if S has a justified true belief that p.*

This thesis on its face is about neither language nor concepts (as psychological entities). Taken at face value, it is a necessary biconditional about people or subjects (reading the scope of the modal operator as short for its universal closure on the variable “S”).

If JTB is the target, then a Gettier counterexample postulates a case C about which it seems intuitively plausible that it both:

- (a) is possible and
- (b) entails that someone has justified true belief without having knowledge.

This would refute JTB by means of the principle that if X both is possible and entails Y, then Y is possible. The example would thus show it to be possible for someone to have justified true belief that p without knowing that p, which contradicts JTB.⁷

3. Timothy Williamson questions whether the usual Gettier case entails that someone has justified true belief while lacking knowledge. He questions whether, necessarily, if the Gettier case is actual then someone has justified true belief that p without knowing that p. In his view, Gettier examples lack enough content to make this sufficiently plausible. In particular, it is not really so clear that if the example *as described* is actual, it follows necessarily that someone does have justified true belief while failing to know what he thus believes. Nor is it really clear how to expand the example so as to yield that result without just begging the question trivially. Moreover, Williamson's reasoning would seem to apply not just to Gettier examples but to counterexamples in philosophy more generally.

Williamson targets philosophical theses (including purported "analyses") in the form of modal claims such as JTB. His own approach is to develop a methodology for assessing such claims in terms of subjunctive conditionals, with no reasoning distinctive of philosophy. The reasoning is rather of a sort used constantly in everyday assessment of such conditionals.⁸

Here I will not try to evaluate Williamson's proposal, which is developed with much interest and detail.⁹ I will let it stand unchallenged as a proposal for the epistemic assessment of modal claims

⁷ It is not enough that *in the example* someone can have JTB without K. We must also commit to the example's being possible (possibly actual). Think of Escher drawings or Lewis Carroll stories.

⁸ Williamson, Timothy, *The Philosophy of Philosophy* (London: Blackwell, 2007).

⁹ Compare, however, the critical discussions by Jonathan Ichikawa ("Knowing the Intuition and Knowing the Counterfactual," *Philosophical Studies* 145 [2009]: 435–43); Anna-Sara Malmgren ("Rationalism and the Content of Intuitive Judgments," *Mind* 120 [2011]: 263–327); and Georgi Gardiner ("Normalcy and the Contents of Philosophical Judgements," *Inquiry* 58 [2015]: 700–740).

such as JTB. I will ask instead: Must philosophical theses take the form of such simple conditional or biconditional modal claims? There *is*, of course, an analytic project that *does* seek conditions individually necessary and jointly sufficient for the applicability of the concept under analysis. These individually necessary and jointly sufficient conditions must avoid vicious circularity by not invoking explicitly in the *analysans* the very concept under analysis, the *analysandum*. That project we here put aside. We turn instead to a project of “philosophical analysis” whose analyses do not most deeply take quite that form.

4. Crucial to our approach will be the distinction between the following two claims:¹⁰

| | |
|------------------|---|
| JTBm(odal) | <i>Necessarily:</i> S knows that p if and only if S has a justified true belief that p. |
| JTBe(xplanatory) | <i>Necessarily:</i> Whenever S knows that p, S knows that p <i>by (in virtue of)</i> having a justified true belief that p, and whenever S has a justified true belief that p, S <i>thereby</i> knows that p. ¹¹ |

Williamson’s attack is plausibly on target against Gettier counterexamples to JTBm, but it would be off target against Gettier counterexamples to JTBe.

Here is why. The attack targets the assumption that, necessarily, if the Gettier example *as described* is actual, then someone justifiably believes a truth that he does not know. This is rejected as insufficiently plausible on its face. Take an example where S competently deduces a truth from a justifiably believed falsehood. What gives us the right to suppose that he is then justified in believing that conclusion? Might he not have acquired a defeater as he drew his conclusion?

¹⁰Here “JTBm” is short for “JTBmodal,” and “JTBe” is short for “JTBe explanatory.” The reason for this terminological choice will soon emerge in the main text.

¹¹This aims to identify something that would always ground knowledge. According to this JTB view, a general and always operative metaphysical explanation for how anyone ever gets to know that p is *by* having a justified true belief that p.

That much we can easily fix by just stipulating that in our example S *does* justifiedly believe the conclusion.

However, for a refutation of JTBm, we need it also to be the case that in the example S does not know his conclusion to be true. What gives us the right to conclude that this is so, simply from how the example is described? This is at best unclear. Nothing in the example rules out S's knowing the conclusion *in some way or other, possibly in a way unrelated to the argument specified in the example*. Williamson's objection against traditionally conceived Gettier counterexamples to JTBm does then apparently succeed.¹²

5. Compare, on the other hand, the following question: Does S *thereby* know his conclusion to be true? Gettier reflection on *this* question has a quite different outcome. Our intuitive reaction now *can* be decisive. It is quite plausible that if S in the example does know his conclusion to be true, he must know it *in some other way*. He does not know it *simply by believing it correctly and justifiedly*.¹³ Traditional armchair Gettier reasoning against JTBe is thus protected from Williamson-style refutation.

¹² Arguably, but not conclusively so, given a proposal by Thomas Grundmann and Joachim Horvath in their "Thought Experiments and the Problem of Deviant Realizations," *Philosophical Studies* 170 (2014): 525–33.

¹³ Objection: *This assessment of the original Gettier cases relies on counter-closure (that is, on the principle that one knows p inferentially only if one knows all the premises on which one's conclusion depends essentially). However, some have argued against this. If they are right, then, at least in some cases, S's knowledge is constituted by her inference from a justified but false belief. Here is one such case, from Ted A. Warfield, "Knowledge from Falsehood," Philosophical Perspectives 19 (2005): 405–16.*

Ted printed 100 copies of the handout for his talk, and, after carefully counting 53 people in attendance, he concludes that 100 handouts are enough. However, Ted double-counted one person who moved during his counting of heads, making his premise that there are 53 people in attendance false.

Prima facie, counter-closure gives the wrong result here, since it says that Ted does not know 100 handouts are enough.

Reply: I am not convinced that this refutes counter-closure. It seems to me that Ted has at least one other good basis for his belief that 100 handouts are enough: namely, that there are approximately 53 people in attendance. I think this also functions (implicitly) as an epistemic basis for his belief, and this is something

The same sort of contrast can be drawn generally as we consider proposed philosophical modal claims. The ones of main interest to the philosopher are the explanatory ones. And not all modal biconditionals are explanatory. Consider the gallery of philosophically interesting phenomena such as justice, virtue, happiness, rightness of action, action itself, and so on, as well as knowledge. Not all modal biconditionals concerning such phenomena will be of interest to the philosopher. We seek rather biconditionals (or, on occasion, at least conditionals) that will *explain* why the phenomenon of interest obtains when it does.¹⁴ This is why we focus on JTBe rather than JTbm.

6. Let us now return to the threefold distinction between language, psychology, and objective phenomena. In addition to the relevant expressions and concepts, there is also presumably the phenomenon of people knowing things. How is this different from the words and concepts, singly or in combination?

Closely analogous to the metaphysics of knowledge is the metaphysics of persons. Here again is the familiar threefold divide: the

he plausibly *does* know to be true. Moreover, it is not clear that this must itself be justified on the basis of the false belief that there are *exactly* 53 people there. Rather, the approximation content might itself be arrived at on the basis of counting, not just by inference from the exactitude content. One might even hesitate and *never* accept the exactitude content but go directly from the counting to the approximation content.

Besides, our reasoning need not commit to counter-closure, since all we need is a strong enough intuition about the *particular* case described, that *in that case* the subject would not know the conclusion that is deduced based on a false premise. Still, the counterexample seems more conclusively persuasive if we find counter-closure very plausible intuitively with full generality.

¹⁴An example may help. You come across a cubical box and ask, "Why is this a cube, specifically, rather than some other sort of box?" It is true that this box is a cube iff it has the form of a Platonic solid other than a tetrahedron, an octahedron, a dodecahedron, or an icosahedron; and necessarily so. Yet its satisfying that Platonic condition will not explain why the box is a cube in the way we might explain it by saying that it is made up of six congruent square sides while the angle between any two adjacent sides is a right angle. This is what constitutes that box as a cubical box and explains in virtue of what it is a cube. Unlike this, the Platonic condition is a necessary and sufficient condition, and necessarily so, but it is not a constitutive condition.

words; the concepts; and the extra-linguistic, extra-conceptual entities, the living persons. If we were to decide that persons are just animals of a certain sort, then of course they would be subject to biological analysis. On another level of analysis, however, we would find metaphysical options such as substance dualism, animalism, and so on. According to an Aristotelian view, the relation between a person and that person's body would be constitution, not identity. The body would need to be alive, and would need to underwrite or ground certain powers and abilities, in order to (thereby) constitute a person.

This style of view involves a concept of metaphysical dependence, of one thing existing or being actual dependently on certain other things and their properties. The dependent thing would then exist or be actual *by virtue of* how the other things were propertied or related.

7. We can now discern three Gettier problems that are quite distinct, however closely they may be interrelated.¹⁵ First is a problem of *semantic* analysis: What is the semantic analysis of the linguistic expression "S knows that p"? Second is a problem of a certain sort of *conceptual* analysis: With concepts understood either as meanings or as psychological entities, what is involved in someone's possession and/or deployment of a given concept? This problem thus concerns people's minds, their psychology, their mental states or the contents of those states. Third is a problem of *metaphysical* analysis. Here our focus is on an objective phenomenon that need be neither expression nor concept. Our focus is rather on a state that people host or an act that they perform. This is the phenomenon whose ontology we now wish to understand. What is the nature of such a state or act, and how is it grounded? In virtue of what is it actual when it is actual?¹⁶

¹⁵Here I understand "the Gettier problem" broadly, as the problem of what knowledge is, in light of the examples showing it *not* to be justified, true belief.

¹⁶Here I have lumped together questions of grounding, questions concerning the *in virtue of* relation, and questions of *nature, essence, or constitution*. I am leaving open whether these various ontological issues should be distinguished.

8. On what epistemic basis might we assess a proposed analysis of that sort?

A metaphysical analysis might have a direct intuitive basis, as it attracts our assent directly upon consideration.

Constitutional analysis may also gain plausibility through its explanatory payoff, even for those who lack naturalist intuitions. We intuitively believe, for example, that if someone knows that *p*, then it is true that *p*. And we might wonder why this is so. At this point we could explain that because it is *necessarily* so, therefore it is generally, universally so. About this necessity we might in turn ask: Is it brute, or can we hope to explain why it is so? In answer we could then appeal to the ontological nature of knowledge, observing that because knowledge amounts *constitutively* to justified true belief, therefore, necessarily, whenever someone has knowledge, he must also have relevant truth.

Thus might we explain how it is that wherever there is knowledge there is (necessarily) belief, truth, and justification. The explanation for each of these facts would derive from the constitution of knowledge as justified true belief.

9. So we have arrived at three distinct theses concerning propositional knowledge:

| | |
|---------------------------------|---|
| JTB _m | <i>Necessarily</i> : S knows that <i>p</i> if and only if S has a justified true belief that <i>p</i> . |
| JTB _e | <i>Necessarily</i> : Whenever S knows that <i>p</i> , S knows that <i>p</i> by (<i>in virtue of</i>) having a justified true belief that <i>p</i> , and whenever S has a justified true belief that <i>p</i> , S <i>thereby</i> knows that <i>p</i> . |
| JTB _c (ontitutional) | Knowledge is (ontologically) constituted by justified true belief. |

If JTB_c can explain JTB_e and JTB_m, it enjoys additional explanatory payoff.

10. We had committed to considering two Platonic problems in epistemology: What is the nature of our knowledge, and how is it constituted? And what is the distinctive value of human knowl-

edge? In particular, why is it better than mere true belief? We now turn to these questions in the following three chapters.

APPENDIX (MORE ADVANCED MATERIAL)

The ontological theory of knowledge as constituted by belief that is both true and justified would gain standing through the metaphysical explanations that it provides. On its basis: first, one could explain how knowledge supervenes at least in part on belief, truth, and justification; second, one could explain also why knowledge entails each of belief, truth, and justification; and third, one could have explained why JTb_m and JTb_e were true had they indeed been true (contra Gettier). The ontological theory might gain *further* support through its direct intuitive attraction when one inquires with sufficient clarity into the constitution of knowledge.

Moreover, the distinction among the three JTb theses, along with their explanatory relations, would seem to be matched by corresponding theses concerning other traditional targets of philosophical analysis and theorizing, such as justice, virtue, rightness, happiness, and the rest. We are thus encouraged to think that a similar philosophical structure might be found more generally (though *how generally* remains to be explored).

Going beyond the case of knowledge to metaphysical explanations more generally, the template would be this:

ME *Necessarily: Whenever X is the case it is so by (in virtue of) Y's being so, and whenever Y is the case, X is thereby so.*

For X we could have *justice (as it applies to a polity), personal identity, personhood, free action*, and so on for all the various phenomena that have been of interest to philosophers.

The Williamson objection is obviated if we think of a counterexample as targeting a view of the form of ME, such as JTb_e above. A counterexample to JTb_e could then take the following form: it is a case in which someone has a justified false belief that *f* and competently forms a true belief that *t* by deduction from his believed premise that *f*, such that he is justified in believing that conclusion

only through that competent deduction. All we need now for a counterexample to JTBe is the intuition that our subject would *not thereby* know that t. This would show it to be false that whenever one has a justified true belief that p, one *thereby* knows that p. Our subject would have a justified true belief that t, but in the example he would not *thereby* know that t.

FURTHER READING

- Brown, Jessica. "Doubt, Circularity, and the Moorean Response to the Sceptic." *Philosophical Perspectives* (2005): 1–14.
- Gardiner, Georgi. "Normalcy and the Contents of Philosophical Judgments." *Inquiry* 58 (2015): 700–740.
- Moore, G. E. "Proof of an External World," "Certainty," and "Four Forms of Skepticism." In *Philosophical Papers*, chapters 7, 9, and 10. London: Allen and Unwin, 1959. Key excerpts appear in *Epistemology*, ed. Ernest Sosa, Jaegwon Kim, Jeremy Fantl, and Matthew McGrath. Oxford, UK: Blackwell, 2008.
- Pritchard, D. H. *Epistemic Luck*. Oxford, UK: Oxford University Press, 2005.
- Pryor, James. "The Sceptic and the Dogmatist." *Noûs* 34 (2000): 517–49.
- Roush, S., *Tracking Truth*. Oxford, UK: Oxford University Press, 2005, especially chapter four.
- Williamson, Timothy. *The Philosophy of Philosophy*. Oxford, UK: Blackwell, 2007.

Knowledge as Action

A. WHAT CONSTITUTES A FULLY APT PERFORMANCE

Judgment and knowledge itself are forms of intentional action—that is the thesis to be argued in this chapter. Such action falls under a certain normative structure of success, competence, and aptness, or success that manifests competence. Judgment is a special case falling under that structure.

A concept of the fully apt performance figures importantly in my account. In order to introduce this concept, we turn first to practical performances, to how they are constituted and to a special normativity that pertains to them. An example yields initial insight, whereupon we turn to performances that are epistemic rather than practical, on which a further example will also shine its light. Those examples show the way to a better view of human knowledge, which is our main objective.

1. *Actions and Performances*

a. Intentional actions come in two sorts. An attempt is an intentional action, an endeavor to attain a certain objective. An attempt can fail and remain a mere attempt, whereas an achievement is a certain sort of *successful* attempt. Thus, if someone intentionally flips a switch in the endeavor to turn on a light, that is an *attempt*. If the attempt succeeds, then the agent does turn on the light by flipping that switch, which, if all goes well, is an *achievement*.

b. Intentional actions (both mere attempts and also achievements) are one sort of performance. Some performances are also aimings, however, without being intentional. The heart, for example, aims by pumping to circulate an organism's blood, which it does through functional teleology without doing so intentionally.

Such functional aimings also come in two sorts. These, too, can fail and remain mere aimings, to be distinguished from those that are successful.¹

2. *The AAA Structure*

a. In what follows we focus mainly on (intentional) attempts. These fall under the sort of normative structure that pertains, for example, to an archer's aim to hit a target. The shot might hit its target, in which case it is accurate. Second, it might be competent: as the arrow leaves the bow, its speed and orientation derived from the archer's skill would, in normal conditions, take it to the bullseye. A shot can be accurate without being competent and competent without being accurate. But even a shot that is both accurate and competent might still fall short in a further respect. Thus, a gust of wind might take the arrow off course, whereupon a second gust might come along and put it back on course. That shot is then accurate and competent, but the accuracy is attributable not so much to the archer's competence as to the luck of the second gust. Such a shot then falls short in a further respect. A shot is *apt* iff its success manifests the competence then exercised by the archer. Its *success* must manifest that competence seated in the agent and exercised on that occasion. Real achievement would seem to require not just success but apt success.

b. Here is the AAA structure, which applies as follows to that shot, and then more generally to attempts. In first approximation:

A shot is *accurate* iff it hits the target.

It is *adroit* iff it is an exercise of competence.

It is *apt* iff it is accurate because adroit.

¹Thus, the expression "intentional attempt" is redundant. Not so "intentional aiming," since an aiming could be just functional or teleological without being intentional. The heart beats, for example, so as to circulate the blood; that is its (teleological) end in so doing.

Generalizing to attempts generally, because these all have constitutive aims, we can also distinguish as follows:

An attempt is *successful* iff it attains its aim.

It is *competent* iff it is an exercise of competence.

It is *apt* iff it is successful because competent.

3. Epistemology as a Special Case

Epistemology concerns a domain of *epistemic* performance, such as belief and judgment (and also inference, learning, teaching, inquiry, and so on).

We will focus here on the performance of *judgment*, to be explained in due course. Although the focus here will be on that particular performance, our account should extend to epistemic performances generally, including functional beliefs that are not judgmental.²

We begin, however, with another example from athletic performance: the basketball shot on goal.

4. An Example: A Basketball Shooter's Animal, Reflective, and Full Aptness

a. Even a player who overconfidently takes low-percentage shots too frequently may retain an excellent ability to sink her shots when close enough to the basket. Success even in her low-percentage attempts may still be creditable to her competence, moreover, and properly so. This need not be affected by her lowered reliability when she tries too often beyond her safe zone, *not* if she is *well aware of her limits* while taking a deliberate risk. What is more, her success *within her proper range* will still seem creditable even if, *while unaware of those limits*, she continues to shoot confidently beyond them.

²For example, any belief whose content is nonconceptual will not be *fully* capturable through linguistic formulation. Thus, you know what a friend looks like. There is a certain physiognomy such that you know your friend to have *that* specific look. Yet it is not something you can capture in English, nor need you even be able to draw it. Your knowledge is still manifest when you recognize your friend.

Before us then are two interestingly different cases:

In one case, the player is unaware of the limits of her competence and shoots indiscriminately at distances too near her threshold of sufficient reliability.

In a second case, the player still shoots at distances beyond her limits of competence, below her threshold of reliability, but now well aware that she is so doing, as she willingly runs the risk involved.

Here is a notable difference between those two cases. In the latter, the player can still perform with full aptness because she knows her limits and can knowingly perform within them when she knowingly ascends above her threshold of reliability. In the former, the player no longer performs with full aptness, not when so near her threshold of reliable enough performance. In that zone—*barely above the threshold*—she is still likely enough to succeed, even without knowing that she is. So even in her ignorance she can still perform with *animal* aptness, but full aptness is now beyond her because, unlike the earlier player, this one fails to know her limits.

The term “animal” here is a technical label. Relative to a certain order, a performance attains “animal” status as adroit (and as apt) so long as it is an exercise of competence on that order (or its success is attributable to competence on that order). No claim is made as to the sophistication required in the animal kingdom in order to be able to act reflectively, *with* endorsement from a higher order. Animals that are well enough endowed, such as dolphins and monkeys, seem capable of endorsement from a higher order, even if this takes the form of implicit awareness that enables functional guidance and not that of conscious, explicit linguistic performance.

b. Does a basketball player normally aim merely to get the ball in the hoop? *That aim can* be attained creditably (somewhat), even with the shooter too far from the basket, below the relevant threshold of reliability. This is so especially if the success attained at that distance is due to a level of competence well above

the average.³ However, basketball players aim not *just* to succeed, no matter how aptly. Normally they aim *to succeed aptly* (aptly enough, through competence) *while avoiding too much risk of failure*. Their shots are assessed negatively when they take too much risk.

A shot that makes a goal in the closing minutes of a game may be welcome in *one* respect: it does score a needed goal! But it is extremely unwise, poorly selected, if the shot crosses the whole court when there was plenty of time to dribble safely to a much better range. So it is a successful shot provided that the player's aim was just to score a goal. However, no good player will normally have only that aim in the thick of a game. Relative to a more ambitious aim, to shoot aptly, that shot is deplorable. In that situation, moreover, the player is negligent even if she does not flout, but only disregards, this fuller aim. The coach may well deplore the shot and scold the player for ignoring the importance of shot selection. Due to the player's negligence, the shot is poorly selected, and thus inferior.

What *more specifically* is the required, more ambitious aim? What is reliability *enough*? This obviously varies from domain to domain. In basketball we know at least roughly where it lies, with due allowance for the position of the player and her teammates, the time remaining on the clock, whether the shot to be taken is a three-point shot, and so on. Many factors thus bear in diverse ways, and good players will take them into account, aiming not *just* to sink a basket but to manifest in so doing the full competence required.⁴

³Even when we put it aside for simplicity, how she does as a *team player* is relevant to this assessment. For example, should she pass rather than shoot? This, too, can affect the quality of her action, as she intentionally shoots by choice.

⁴Granted, it seems initially implausible that any significant difference in credit could derive simply from surpassing a sharp threshold of reliability. Would not a shot *barely* below that line be about as creditable as a shot *barely* above it? But this supposed problem is at least in part a figment of our simplifying assumption that what separates competence from incompetence is a thin line. Even given the vagueness in the concept of competence, it is still plausible that a player could fail to recognize that she would shoot reliably enough from where she stands (when she is barely so competent).

c. Consider now a shooter as she approaches a distance to the basket near her relevant threshold of reliability. And suppose that she is above the threshold, but *indiscernibly so to her*. A statistician-coach-observer might know perfectly well that the player is now barely above the threshold. Suppose that he has studied her success rate extensively, aided by a device that measures with exactitude her distance from the basket. That way he can tell that she *is* reliable enough at that distance (given the circumstances noted earlier). But she herself is very far from knowing any such thing.

The player may still attain her basic aim: namely, to sink that shot in the basket. In that respect, her shot may also be apt. Its success may manifest the competence that the statistician knows her to possess even at that distance.

What, then, is she missing? Anything? Well, although her first-order animal aim *is* attained aptly, not so the *reflective* aim of succeeding thus aptly, an aim that she should also have, whether or not she has it. She does aptly score her goal, but she fails to attain aptly the aim of *aptly* scoring it. Unlike the statistician, she is unable to tell that her shot is still reliable enough at that distance. If she shoots anyhow, and her shot turns out to be reliable enough, she may *aptly* reach her aim of scoring that goal. What she does *not* aptly reach, however, is the aim of *aptly* scoring. And the success of her shot is hence not *fully* creditable to her, given this important element of luck. Even if her first-order success *is* apt, it is not guided to aptness through apt meta-awareness that the shot will be apt; hence, it is not a *fully* apt shot.

d. Our basketball example suggests a distinction between first-order safety and second-order safety. The player's shot will be safe when she is (even barely) above her threshold of sufficient reliability. So situated, not too easily will she then fail in her attempt to make that goal. Unlike that first-order performance, however, her second-order performance may still be unsafe. Unaware of her threshold's location, she might too easily shoot inaptly, below that threshold. Properly situated as she is *in fact* (though barely), she *is* thereby disposed to shoot *successfully* and *aptly*. Because she is unaware of her threshold, however, she might too easily shoot inaptly, because she might so easily be *improperly* situated and still shoot just the same.

Suppose, moreover, that the lights might easily dim just as she is taking her shot. Because of this, she might easily shoot inaptly, in a way that does not manifest complete competence. Even if the lights *might* then easily dim, however, so long as, in fact, they *do not*, her shot can still be apt. Two things are here plausibly compatible: first, she might too easily shoot inaptly (because the lights might so easily dim); but second, she in fact shoots aptly, with a shot whose success does manifest her complete shooting competence, which *in fact* is present in its entirety.

When our basketball player is disposed to succeed reliably enough, this may be because she satisfies the SSS (skill/shape/situation) conditions of a relevant first-order competence. The *skill* involved is still retained by her even when asleep in the dark, although when asleep one is not in proper *shape* for basketball shooting, nor is one in a proper *situation* when in the dark. These are conditions that determine whether one is disposed to shoot *accurately*, and reliably enough. But the SSS conditions of *the competence to shoot aptly* are not the same. The competences are clearly different: the SSS conditions are different. A shot aimed at hitting a target is apt if its *success* manifests *first-order competence*. For a shot to be *fully apt* it must succeed aptly, and, in addition, it must *aply* succeed aptly. It must be aptly apt, with its *aptness* manifesting *second-order competence*.

e. Take any domain of human intentional action, whether that of athletic performance, of the performance arts, or of service-professional performance, as in medicine or the law, and so on. In any such domain, achievement is creditable to the extent that it is attained through competence rather than luck. Whenever an aim is attained, there will be a dimension with pure luck at one end and pure competence at the other. And there will be a threshold below which the aim is attained too much by luck, too little by competence. This is the threshold below which the agent's attempt, given his or her relevant skill/shape/situation, would be *too risky*. This means "too risky relative to the goal, internal and proper to the domain, at which the agent's performance should aim on that occasion." This is a notion familiar to spectators and assessors of athletic performance. Thus, a swing by a batter is too risky if the pitch is way out of the strike zone, a basketball shot too

risky if taken beyond a safe range, a tennis serve too risky if hit too hard and too flat for the occasion, and a bow hunter's shot too risky when conditions are too unfavorable, with respect to lighting, wind, distance, and so on.

What sets such a threshold? This will vary from domain to domain. It may be conventional and formalized, as in some professional contexts, or it may be less formal, more intuitive, as in the domain of a hunt. In each case, the threshold will be set by considerations distinctive of the domain and the proper basic aims of performances in it, and not by external pragmatic aims that the performer might *also* have for his or her performance.

Aims external to the domain might, of course, properly motivate a performer to take an outrageous risk. Even so, from a domain-internal perspective the performance can still be *too* risky and the performer negligent in deliberately taking such a risk, or even in being too insensitive to the risk taken. Thus, a basketball player might be offered a vast sum for taking a shot from across the full length of the court and *might* thus act quite rationally and appropriately in taking that shot, all things considered, especially if the offer is innocent and not a bribe. But the shot is then still bad as a basketball shot because of how poorly selected or negligent it is *as a basketball shot taken in a game* if there is plenty of time to dribble safely to within a safer range.

f. Performing with *full* aptness would normally require knowing (at least at the animal level) that one *would* then perform aptly. This is the knowledge that must *guide one's performance* if it is to be fully apt.⁵ This will be seen to play a role in epistemology.⁶

⁵Here I do not take up the question of how one attains one's knowledge of subjunctive conditionals. One interesting proposal is in Timothy Williamson's *Philosophy of Philosophy* (London: Blackwell, 2007), chapter 5. Without necessarily endorsing that proposal, I do share the assumption (p. 141) that we regularly know the truth of such conditionals, knowledge that we put to much important action-guiding use. (But I do not attribute to Williamson any agreement that subjunctive conditionals help guide our actions in the specific way that I suggest: that is, by enabling their full aptness.)

⁶We should take note of a qualification. Many performances cannot aspire to success that is fully apt in the demanding way specified. Athletic performances, for example, nearly always aspire not to antecedently *assured* success but only to

B. FULL APTNESS IN EPISTEMOLOGY

1. Consider next an example in which a *guess*, surprisingly enough, might still qualify as a case of “knowledge.” Please recall your yearly eye exam as I describe mine:

When I go for my exam, I am asked to read the lines of a chart with letters that shrink line by line, from a huge single letter at the top to those barely visible at the bottom. At some point I start to lose confidence that I am getting the letters right, but I keep going until the technician tells me to stop and then records some result. At that point there are many cases in which I am quite unsure as to whether a letter is an E or an F, say, or a P rather than an F, and so on. Suppose, however, that it turns out that (unbeknownst to me) I am in fact unfailingly right year after year about a line at which I am thus unsure. At that point I am in effect “guessing.” I do affirm what I believe each letter to be, to myself in private and to the technician in public, and I do so in the endeavor to get the letters right. That is, after all, what the test requires: that I *try* to answer correctly. And we can surely stipulate that I thereby manifest a competence, one I do not recognize as reliable enough. *This latter* is why I resort to guessing when I continue to affirm what I believe the letters are as I undergo the test. Unbeknownst to me, however, my affirmations turn out to be surprisingly reliable.

How, then, do we assess my performances? We are here conflicted. *Somehow* I *do* know what letters I see, as shown by my impressive

success that is likely enough. But what needs to be assured is that the success *would* be likely enough. This is not quite like assurance, based on a fair sample, about the percentage breakdown of an urn's contents.

In what follows we will mostly focus on cases of assured success and the corresponding full aptness. This is for simplicity, though a complete account needs to cover the more general phenomenon of performance that is apt “fully” *to the extent possible without flaw*. Thus, the complete account will need to cover also performance guided by prior knowledge that it would *likely enough* be apt.

reliability. But there is also a pull to say that I do not *really* know. What accounts for this?⁷

2. First we need a distinction. I still affirm what I believe the letters to be on the lower rows while aiming *to get them right*. I give it my best shot, because only thus will I undergo the vision test properly. But do I aim to get it right *aptly, reliably enough*? No, by that point it matters little whether I do get it right *reliably at all*. I just make my best guess, aiming to undergo the test successfully so as to get the right eyeglass prescription. Whether on the lower rows I am still reliable does not matter much, because the letters there are tiny and my vision will be fine either way. Accordingly, I make my guesses without *endeavoring* to get the answers right *aptly*.⁸ Still, it turns out that I am nearly infallible on those rows; that is how it turns out *by hypothesis*.

Many will insist that *somehow* a vision-test subject does know, even without knowing that he does. All who attribute knowledge to the blindsighters and chicken sexers of Gettier lore will agree. Moreover, the vision-test example is easily conceivable, with no need to indulge in science fiction.⁹

⁷ And it would be easy to construct a similar example concerning one's memory. With aging comes the need for assurance that ostensible memories are still reliable enough.

⁸ However, even if we *do not normally* aim to get it right *aptly*, we still *might* so aim, and we might *still* be guessing. In order to go beyond guessing, we must affirm confidently enough *both* on the first *and* on the second order. If we confidently enough affirm on the first order but remain too unsure on the second order, then we are still *in that way* guessing.

An eye-exam subject will normally guess both on the first and on the second order when viewing a row of the eye chart far enough down the chart. Take subjects who feel confident on the first order. They may still feel unsure on the second order as to how or even whether they have a competence that reliably delivers that first-order assurance. Presumably their stance is possible even if not perfectly coherent.

⁹ By contrast, some of the earlier cases are hard to imagine for an actual human, since the attribution of knowledge clashes so radically with our background knowledge. That applies both to Norman the clairvoyant and to Truetemp, who can tell ambient temperature directly, and does so (unbeknownst to him) through a thermometer implanted in his brain. For all I really know, however, I myself in fact qualify as the vision-test subject (without the confidence or confirmation

3. Again, as the letters get smaller, even as we start guessing, we might still know in some very basic way, with a subcredal “animal knowledge” *below* even the animal knowledge that requires belief. What are we missing as we descend to that lower level of knowledge? What distinguishes the higher knowledge that we enjoy on the rows with bigger letters?

Perhaps what makes the difference is just more confidence. In order to know when the letters get smaller, do we just need to be more assertive? Given how reliable we are by hypothesis, is that all we need at that point: just more confidence? Is that the only relevant difference that gives us knowledge of the bigger letters?

Some of us are constitutionally assertive risk takers; others can acquire confidence through therapy. Suppose we gain our confidence only through therapy, with no other change. That would not give us the sort of knowledge we enjoy with the larger letters. Indeed such artificial increase of confidence can *worsen* a subject’s epistemic position.

Compare someone who gains not just confidence but also *confirmation* that he remains reliably right even when the letters shrink. *This* gain, beyond mere confidence, might raise him to a higher epistemic level. Now he might attain the knowledge requiring judgment, not just a guess. *His* knowledge will then comprise not only more confidence but *also* the proper *meta*-assurance that, even for those very small letters, his level of competence limits epistemic risk within proper bounds.

Absent such additional confirmation, the vision-test subject lacks well-founded confidence *on the second order* that his first-order affirmations are more than sheer guesses. Even if by hypothesis his guesses are *not* right just by luck, he cannot be sure of that, not competently.

By contrast, when the letters near the top of the chart are clearly and distinctly discernible, we do not just guess, which comports with our knowledge *that we can aptly identify those letters for what they are*, that our affirmations on those rows will be apt.

but with extremely reliable guesses). The vision-test subject fits better with our background beliefs, as do blindsighters and chicken sexers, by contrast with clairvoyants or truetemps.

Perhaps what is missing on the lower rows is our taking our “guesses” to be reliable enough. As our confidence wanes, we still affirm, even once we start to guess. What, then, is missing? Reflective competence is missing, as is reflective aptness, and also aptness full well.¹⁰

4. Again, epistemic agents do not aim just for correctness of affirmation. They also *judge*, aiming for *aptness* of affirmation. So even a properly confident subject who affirms aptly might fail with his judgment. Why so? Because even while affirming aptly in the endeavor *to affirm correctly*, he might fail to *aptnly* affirm aptly: that is, he might fail to *judge* aptly, because to judge is to aim to affirm not only correctly but also aptly. While affirming in the endeavor *to affirm aptly*, he might *fail* to affirm *aptnly* in that endeavor. In other words, his *alethic* affirmation, aimed at truth, might be apt without being fully apt, in which case his *judgment* will not be apt.

We rely here on two facts: (a) that one may do something as a means to more than one end and (b) that aptness pertains not just to the means but to the whole structure of the form: *taking means M to end E*. Thus, one might flip a switch aptly in an endeavor to illuminate a room without doing so aptly in the endeavor to alert someone, even if by flipping the switch one aims concurrently to attain each of those two aims. That is to say, the first aim might be attained aptly thereby, without the second aim being attained aptly, or at all.

Similarly, one might affirm aptly in the endeavor to affirm correctly (an aim required if the affirmation is to be alethic and not just pragmatic). Compatibly, one still might *not* affirm aptly in the endeavor to (alethically) affirm aptly, not even if one then endeavors to affirm not just correctly, with truth, but also (thereby) aptly. Only with *apt* success in this second endeavor does epistemic affirmation count as fully apt. And only thus is *judgment* also apt, beyond the aptness of one's affirmation.

¹⁰ A fuller treatment would need to consider also how confidence itself can bear on knowledge.

Equipped with our distinctions, we next, in chapter 6, consider some problematic issues of epistemology.

APPENDIX (MORE ADVANCED MATERIAL)

1. Affirmations are intentional acts. These can be conscious or subconscious. If we define choice and decision so that these must be conscious and deliberate, then much of what we do freely and responsibly is not done by *choice* or *decision*. Even what is not done by choice or decision might, then, be done through a free act of intention or design, however, whether conscious or subconscious. By means of such an act one might come to believe something.

Would such belief be *voluntary*? If the “voluntary” is only what is done at will, so that one is then free to refrain from doing so, or even to do the opposite, at will, then belief does not qualify as voluntary. We do not believe at will if this requires that we be free at will to not so believe, and even to believe the opposite. I am not now free at will *not to believe* that I am awake, and much less am I free at will to *believe the opposite*. Nor am I free to *believe* for any reason whatsoever, including that it will be personally advantageous to so believe. I can, of course, freely put myself in a position wherein I *would* believe something that for whatever reason I would like to believe. Thus, if in the dark I would like to believe that the lights are on, I could freely turn them on, whereby I would freely put myself in a position wherein I *would* believe that the lights are on, thus satisfying my desire to so believe. What I am unable to do at will is to right now directly form the belief that the lights are on even while my basic relevant evidence remains the same as I stay in the dark. Nor can I even stop myself from then believing that the lights are *not* on. At least none of this can happen with rational or epistemic propriety.

2. We can agree with those considerations, however, while rejecting the view that restricts our free exercise of will to a freedom of indifference, whereby in order to be free the will must be able to disregard and even to oppose the obvious force of reason. More

properly understood, what voluntary freedom requires is only the ability to override *improper* influences so as to align our action with the requirements of reason.

We may still be urged to explain why we are not free to believe freely and at will based on what reason determines to be our best course, all things considered at the time. Fair enough. Our explanation does not so far explain why pragmatic reasons cannot directly guide belief, at least not properly.

The fuller explanation must focus on how belief is constituted. Let us focus in particular on judgment and on judgmental belief. Here the constitutive act is affirmation and the constitutive aim is truth. So the constitutive act is *alethic* affirmation, or affirmation aimed at truth.

In order to supplement our explanation, we must distinguish extrinsic from intrinsic reasons. *Extrinsic* reasons pertaining to a means/end action are reasons that pertain to whether we should engage in such actions at all, including reasons for adopting that end at the time. *Intrinsic* reasons pertaining to a means/end action are reasons for adopting certain means, *given the end*, where the end is not at issue.

3. And now it may be seen that we remain free concerning whether to form a belief on a certain question based on any reasons whatsoever. What we cannot properly settle at will, based on any reasons whatsoever, is what to do as a means to a certain end once the end is set and beyond consideration. We *can* properly determine what to do in such a context based on our belief as to what *is* a means to that end. In order to judge that *p*, let us suppose, we must affirm that *p* in an endeavor to answer correctly and aptly the question whether *p*. In this case, the reasons that bear properly on what we are to do are reasons that bear properly on whether *by affirming that p* we will be likely to affirm correctly and aptly. Consider here the *cogito* as an example. In order to bear properly, however, a reason cannot be blatantly circular or question-begging.

4. Take reasons that might bear on whether one should affirm that *p* in an endeavor to answer correctly the question whether *p*.

They may bear in the way Descartes's rationale bears on whether he should affirm "cogito" propositional contents such as <I think> or <I exist>. When we consider whether we face a hand or a fire, however, we determine what to affirm based on a different sort of rationale. It is not a second-order rationale concerning what we would be likely to affirm correctly given our relevant skill set, shape, and situation at the time. When we determine what to believe perceptually based on our relevant sensory experience at the time, our experience has a content that may attract our assent in various degrees of strength, and other factors may then do so as well on that same question. Some of these factors may be consilient, some conflicting, until a seeming results whose degree of strength, in turn, helps us determine whether assent is or is not then proper. If rational, the assent is then based on the resultant seeming *directly*. A resultant seeming is automatically a possessed rationale. It is in the subject's possession and can perform its job of rationalizing affirmation (or assent) on its own through the proper operation of a psychological mechanism. This is all properly so if when things then seem that way to the subject his cognitive psychology is functioning properly.

A *proper* rationale for *alethic* affirmation *in the endeavor to attain truth* must be of a sort that bears on the truth of that affirmation. It would be rationally improper to affirm *in that endeavor* based on a rationale that did not bear on the objective at which that affirmation aims. Consider the constitutive aim of an affirmation that qualifies as a judgment. It must be, at a minimum, an affirmation that is alethic, that aims at truth, and, as we have seen, also at aptness. Reasons that bear on other ends that could be attained by so affirming are hence not relevant to judgment, and indeed one cannot coherently affirm *for such reasons* when one is aiming at aptly attaining truth.

Consider one's attempt to turn on a light by flipping a switch. Coherent reasons for flipping that switch *in that endeavor* would then be reasons that at least ostensibly would bear on how flipping it would aid one's attaining the aim constitutive of the endeavor. Compare one's judgment that *p*, which is constitutively an attempt by affirming that *p* to answer aptly the question whether *p*. Proper, coherent reasons for affirming that *p in that endeavor*

would then be reasons that bear on how so affirming would aid one's attaining the aim constitutive of the endeavor. This is what is so incoherent about judging that *p* for the reason that it would please someone or would yield some other pragmatic objective. Attaining that objective is evidently irrelevant to attaining the aim constitutive of judgment if judgment on a question (whether *p*) is constitutively an attempt to get it right on that question.

FURTHER READING

- Greco, John. "Safety in Sosa." *Synthese*, forthcoming. This will be a special issue on Ernest Sosa, edited by Christoph Kelp.
- Hieronimi, Pamela. "Believing at Will." *Canadian Journal of Philosophy* 35 (2009, supp1): 149–87.
- Shah, Nishi. "How Truth Governs Belief." *Philosophical Review* (2003): 447–82.

Varieties and Levels of Knowledge

A. GETTIER CASES

1. Consider the examples that have proved important in the Gettier tradition:^{*}

- (a) *Inferential cases.* A subject S believes on excellent evidence that Nogot here has a Ford. But Nogot has no Ford, unlike Havit, also here and, unbeknownst to S, a Ford owner. S concludes that someone here has a Ford from his premise that Nogot does. His conclusion belief is thus justified, and also true. Yet S does not thereby *know*, since that belief is true only because of Havit, not Nogot. Here S fails to know through his justified true belief.¹
- (b) *The barns example.* Barney sees a barn and believes accordingly, although he might as easily be viewing a mere barn façade, of the many nearby. Most epistemologists deny that he knows, but a substantial minority is not so sure.²

^{*}Here I understand the Gettier tradition *broadly*, to include not just the original Gettier examples but the whole tradition meant to understand what *propositional knowledge* is in the light of examples that bear on whether it is *justified true belief*.

¹Edmund Gettier's own examples are inferential ("Is Justified True Belief Knowledge?" *Analysis* 23 [1963] 23: 121–23). Our Havit/Nogot example is from Keith Lehrer ("Knowledge, Truth and Evidence," *Analysis* 25 [1965] 25: 168–75).

²Alvin Goldman (Carl Ginet), *Epistemology and Cognition* (Cambridge, MA: Harvard University Press, 1986). However, some X-phi survey results imply that, on intuitions about Barney, the folk are in sharp disagreement with what has been the philosophers' consensus. See David Colaco, Wesley Buckwalter, Stephen Stich, and Edouard Machery, "Epistemic Intuitions in Fake-Barn Thought Experiments," *Episteme* 11 (2014): 199–212. This paper tests versions of the original fake barn case. The number of "knowledge" attributions is high. Josh Knobe and John Turri have both independently replicated those findings (the "knowledge" attribution

Compare Simone, who nears the end of her pilot training and each morning might, to all *appearances*, from her viewpoint, as easily be piloting in a simulator as in a real plane. When she happens to be really in flight, as she believes, can she know that she is? How could she, given that she might just as easily be in a simulator, with no way to tell the difference?

- (c) *BonJour's clairvoyant, Lehrer's Truetemp, the chicken sexer, the blindsighter.* These subjects find themselves believing quite reliably but with no idea how they are doing so or even *that* they are doing so. These cases divide respondents more evenly. Some take such subjects to know; others think the opposite. And some can feel the pull of the opposing view but decide to "bite the bullet."

2. Subjects of sort 1b, such as Barney and Simone, seem to fall short because of their poor situation. What deprives them of knowledge seems to be the danger posed by nearby possibilities because of the nearby mere façades or the simulation cockpit. But just where is the *unsafety* located? This question arises once two orders of safety emerge, as in our basketball example. Where and how do Barney and Simone perform *unsafely*? Is it mainly on the first or rather on the second order?

3. Plausibly, Barney falls short of knowledge because he affirms *somehow* unsafely. Too easily might Barney have affirmed that he faced a barn when it was only a barn façade. However, that would not explain why it is so intuitively attractive that Claire (the clair-

findings, though not the age effects) with very large sample sizes. Overall, on a 0–6 scale, the mean response is about 4.75. See also J. Turri, W. Buckwalter, and P. Blouw, "Knowledge and Luck," *Psychonomic Bulletin and Review* 22 (2015): 378–90. This article tests a more general category of cases, "failed threat" cases, to which fake barn cases belong. The percent of "knowledge" attributions is high (sometimes topping 80 percent) and does not differ from ordinary examples of perceptual knowledge. Finally, there is also a paper by Turri "Knowledge and Assertion in 'Gettier' Cases," *Philosophical Psychology* (under review), which tests a version of the original fake barn case. The result here is that the percent of attributions of "knowledge" and "assertability" is high (>80 percent) and does not differ from ordinary examples of perceptual knowledge.

voyant) and Truetemp likewise fail to know. The alethic *affirmations* of Claire and Truetemp might be confidently safe!

By contrast, if knowledge requires aptness of *judgment*, this explains why Claire and Truetemp fall short and *also* why Barney's affirmation needs to be safe.³ Aptness of judgment *entails* safety of affirmation.⁴ In *judging* as he does, Barney must on some level be aware of the following, by taking it for granted, perhaps, or presupposing it: *that not too easily might he go wrong on the matter of a facing barn if he affirmed at all by then exercising his perceptual competence*. In judging as he does, after all, Barney aims to affirm aptly that he does face a barn. So his *judgment* will be apt only if he aptly attains *this* aim: that of affirming aptly. And this will happen only if he is guided to the aptness of his affirmation (that he does face a barn) by his second-order awareness that if he then affirmed (that he faced a barn) he *would* (likely enough) be right.⁵ From this it follows that if his judgment is apt, then his affirmation is safe. We thus co-opt *this* safety requirement and conclude that Barney and his like fail to judge aptly, which explains why they fail to know full well.

Again, in order to know full well, Barney must know that if in his conditions he affirms that he faces a barn, not easily will he thereby go wrong. He needs to know this conditional so as to guide himself to apt affirmation, as required for *full* aptness of affirmation, which is what apt judgment requires. So his judgment can then be apt only if safe. Accordingly, Barney knows full well only if his judgment is safe. But the nearby fakes make his judgment unsafe, which precludes his knowing full well. Thus we can understand the attraction to think that Barney does not know if the knowledge of interest to us is judgmental knowledge full well.

We are focused on apt judgment (and judgmental belief) and on corresponding knowledge above mere subcredal animal knowledge (as in the eye-exam case). Judgment is *affirmation* in the endeavor to *affirm aptly*. In judgment one aims to *alethically* affirm aptly. Judgmental affirmation that p must hence manifest competence

³One might of course find *different* explanations for these different phenomena, but it is always a point in favor of an explanation that it explains a broader range of cases.

⁴This will be developed later.

⁵Again, emphatically, this awareness need be neither conscious nor temporally prior.

not only to get it right on the question whether *p* but also to do so aptly. In order for a *judgment* to be apt, the subject must *aptly* attain aptness of affirmation.

Accordingly, Barney's *judgment* succeeds if and only if it attains not just correctness but also aptness. And *this* aim is aptly attained, finally, iff the *aptness* of Barney's affirmation manifests his competence to succeed in *affirming aptly*.

That being so, in order to be apt, Barney's *judgment* must be safe. Since it is constitutively an affirmation aimed at getting it right *aptly*, a judgment that *p* will succeed only if it meets two requirements:

- (a) The embedded affirmation that *p* must of course be apt.
- (b) The embedded affirmation must be one that *would* be apt, in the following way: The subject must possess a complete competence in virtue of which he *would* then reliably enough affirm correctly *and aptly* on that question if he affirmed through exercise of that competence.

In the best-case scenario, the agent who *judges* aptly knows that he would likely enough affirm correctly if he affirmed as he intended. The agent *affirms alethically* fully aptly only if guided to a correct *and apt* affirmation by second-order awareness of his competence to so affirm. It follows that if a *judgment* is apt, the embedded *affirmation* is then *fully apt* and must hence be safe, so that the judgment must be one that *would* be true, and must itself in that way be safe. The affirmation must be safe because the agent must know that he *would* succeed aptly if he tried, so that *if he affirmed he would do so correctly*, which is tantamount to safety of affirmation and, in turn, to safety of judgment.

By contrast, an affirmation can be apt without being safe, without being one that *would* be true if made.⁶ This is because an apt affirmation can be *unsafely* apt, unlike a *fully* apt affirmation.

What is more, if knowledge full well does require *full* aptness of affirmation, that *also* explains the pull to think that Claire and Truetemp do not really "know." This is because we are pulled to think that they do not know full well. Supposing that they do not

⁶Note the specific sort of safety involved here: a performance constitutively aimed at *X* is "safe" iff it *would* then (probably enough) attain *X* if executed.

guess, but *judge*, this judgment of theirs will *not* be apt. Even if the embedded affirmation *is* apt, as stipulated, these subjects still do *not* succeed *aptly* in thus *affirming aptly*. They affirm aptly through luck, not competence. Although Claire and Truetemp have animal knowledge, as do our eye-exam subjects, they still lack knowledge full well. And the same goes for blindsighters, chicken sexers, et cetera.⁷ What pulls us to deny that these subjects really know is the evident fact that they fall short of knowledge full well.

Consider the intuition that our eye-exam guessers do not *really* know (even if somehow they do “know”). That intuition plausibly flows from their failure to know *full well*. They lack the required apt awareness that if they affirmed on the first order they would (likely enough) do so correctly and aptly. Of course, what they are missing is not just full meditative conscious awareness. *Rather* they lack even a presupposition or implicit awareness that their first-order affirmations are and would be apt.⁸

We can thus explain our divided intuitions in each of the 1b and 1c cases.⁹ The subjects in those cases do enjoy the sort of animal

⁷We are helped to see more clearly into this nest of issues if we distinguish (a) S's aptly affirming that he would affirm aptly if he answered the question whether p from (b) S's aptly attaining his objective of aptly answering the question whether p. Only the latter requires that S's affirming, in answer to the question whether p, be *guided* by his apt second-order awareness that his answer to the question whether p would be an apt answer. And for *this* to be so, it does not suffice that one enjoy second-order awareness that one's first-order affirmation would be apt. This second-order awareness must also *guide one* to the relevant aptness on the first order. Our attainment of such aptness on the first order must *manifest one's* pertinent second-order competence.

⁸Some features of a situation might rationally sustain our implicit assumption that the light is good for color judgments. We might note that the light is natural, for example, which we know based on the open window, the sunny scene, the light bulb that is off, and so on. When we judge that the surface seen is red, we may thus rely not only on its red look but also on those other factors. That is indeed so if specific reasons are needed for taking the lighting conditions to be appropriate for color judgments. And this is not to suppose that one *always* needs such specific positive evidence for the appropriateness of one's first-order SSS conditions. On the contrary, such appropriateness might be presupposed by default absent specific signs to the contrary. (Of course, we could rather say that the very absence of such signs is then an implicit factor to which one responds rationally in making one's first-order judgment. But this would just rephrase the essentials of our view.)

⁹People do seem divided. Some are pulled to say that these subjects do “know,” others that they do not “know,” while others yet (myself included) feel a pull in

knowledge that comes with apt affirmation. But they do not know full well. As for the inferential 1a cases, these subjects clearly have no knowledge of any sort: not of the animal sort, and even less of the reflective sort. Nor, of course, do they know full well.

The requirement of full aptness accounts for all three sorts of examples. The fact that knowledge has varieties accommodates the apparent conflict. Some intuitions are sensitive to the lower and some to the higher varieties of knowledge, even if none need be conscious or explicit in order to have its effect.

B. ON THE EPISTEMOLOGY OF THE REFLECTIVE

1. Independently of the importance of knowing full well, human knowledge intimately involves higher-order phenomena. Here are some reasons why that is plausibly so—a, b, c, and d below—even without a full display of the supporting arguments.

- (a) Judgment is affirmation with the intention to *thereby* affirm competently enough, and indeed aptly. That distinguishes judgments from mere guesses. The quiz-show contestant affirms in the endeavor to thereby affirm correctly (and thus win the prize) while taking his affirmation to be a sheer guess, far from apt epistemic performance.¹⁰

Judgment thus involves a second-order stance regarding one's own affirming. When one judges, one affirms with the aim to get it right aptly by so affirming (where aiming is more than just wishing or hoping).

each direction. Of course, once we focus on metaphysical rather than conceptual or semantic analysis (a distinction spelled out in a later chapter), we need not quarrel, especially if we can agree that across the board there are significant phenomena having to do with truth and getting it right. It is helpful and interesting to consider how those phenomena are constituted and interrelated. Compatibly with all this, we might also conclude that they can vary in importance and value.

¹⁰Objection: *But a guesser might try not just to get it right but to get it right aptly, while recognizing that the chance of his getting it right aptly is very low.* Reply: Still, the two actions are quite distinct: aiming to get it right is still different from aiming to get it right aptly. Thus, we can distinguish the guess from a judgment made through the same act of affirming. (But a fuller reply would also invoke degree of confidence.)

- (b) Suspension of judgment is an intentional double-omission whereby one omits affirmation, whether positive or negative. Inherent to rational suspension is the assessment that affirmation is then too risky, which implies that, whether positive or negative, it would not then be apt, or at least implies the *absence* of assessment that it *would* be apt.

Suspension thus understood involves a second-order intention to double-omit (an intention that need not be conscious or temporally prior). Judgment, whether positive or negative, is on a level with suspension as part of a threefold choice: affirming, denying, suspending judgment. Judgment, too, thus involves a second-order intention to affirm in the endeavor to affirm aptly, whereas intentional suspension involves intentional double-omission (of affirmation both positive and negative) in the endeavor to affirm *only* aptly. So here is an aim shared by epistemic affirmation and suspension: the aim to *affirm aptly and only aptly*.

- (c) For competence of judgment on a first-order question, epistemic negligence must be avoided through responsiveness to reasons that a fully proper judgment must weigh.

Reasons of what sort? Reasons against judging that *p* can be *counter-weighting* reasons, reasons for the *opposite* judgment, the judgment that *not-p*.

By contrast, *undermining* reasons work against the judgment that *p* without favoring the opposite judgment. These might be reasons, first, for thinking that no good reasons are available. Second, they might be reasons for thinking that one is not then in good *shape* to render a judgment. Third, they might be reasons for thinking that one's *skill* for rendering a judgment is degraded or altogether absent. (Skill is an innermost competence that one might keep even when in bad shape or poorly situated.)

These three, at least, are reasons to which one must be sensitive so as to avoid epistemic negligence in first-order judgment. And note what is true about them all: one's sensitivity is on the second order, involving second-order competence.

Because first-order judgment must be responsive to such factors, which must not be neglected, first-order judgment must be aided by second-order competence. Consider a belief that stays

in place because no second-order defeaters emerge. But suppose that, *being unresponsive* to the presence or absence of defeaters, the agent would have believed the same *even if defeaters had emerged*. Here again is a kind of luck that reduces, and may even remove, credit that might otherwise be due.

- (d) Finally, consider that historical paradigm of the highest form of certain knowledge: the Cartesian *cogito*, the thought *that I now think*. What accounts for its particularly high status? Our approach offers a distinctive way to illuminate that status. Consider my judgment *that I now think*, and suppose this to be my affirmation of the indexical content <I now think>, an affirmation made in an endeavor to affirm aptly (indeed with utmost, infallible aptness). If so, my judgment will attain its objective iff my affirmation attains aptness *and does so aptly*. This means that my aptness in so affirming must itself be secured aptly. So I must then enjoy a second-order competence to secure aptness in my first-order affirmation. It is not enough that my first-order affirmation attain correctness, or even that it do so aptly (with enough reliability). It is required *in addition* that *this* aptness of one's affirmation be *itself* attained aptly. And this is precisely what the *cogito* affirmation does, paradigmatically so, after the thinker's tour through the gallery of skeptical scenarios in the earlier Cartesian meditations.

Once we arrive at the *cogito* passage at the beginning of Meditation Three, Descartes has satisfied himself that in affirming the indexical content <I now think> and in the correlative judgment, he is right, *and he is bound to be right*. Suppose this insight guides his apt affirmation. In that case, the *aptness* of that affirmation (of <I now think>) is now attained under the guidance of his apt belief that this *would* be an apt affirmation, which makes his knowledge that he then thinks a case of knowing full well.

2. The foregoing considerations—in a, b, c, and d of the preceding section—suggest why *reflective* competence is epistemically significant.

It might be replied that such second-order competence is *not* necessarily competence to form a *judgment* that one's first-order judgment would be apt. And of course no such conscious judgment is *required*, even if in certain cases it is importantly present, as in the case of the *cogito*, and even if it *is* required for a level of meditative ascent of special interest to the thoughtful. But the second-order stance required need not *in general* take the form of a conscious judgment. It might just be a presupposition, an implicit awareness that all is well enough for first-order judgment. And in certain basic cases this might even be a default stance properly sustained absent defeaters. No special rational basing is required; one need only be sensitively ready to detect defeaters.

So much for the distinction of animal versus reflective knowledge and for reasons in favor of its deployment. What speaks in favor of going beyond that, to a more developed virtue epistemology?

C. THE REFLECTIVE: IMPORTANT BUT SUBSIDIARY

1. A further theory will be more explanatory than the simpler account. Why does the second-order dimension matter for our first-order knowledge? Why does it matter whether one has merely animal knowledge or "ascends" rather to the more reflective levels? Why doesn't reflective knowledge involve *just more knowledge*, animal knowledge on top of animal knowledge? Why should the second-order knowledge *improve* the knowledge on the first order, raising *it* to a *better* level of knowledge?

Reflective quality is important for human knowledge *largely* for the reasons already suggested: (a) because of the nature of judgment and how it differs from guessing, (b) because of the nature of suspension, and (c) because of how competence must avoid negligence and insensitivity to defeaters. All of these—a, b, and c—involve ascent to a second order, in the ways noted.¹¹ And, as a bonus,

¹¹Objection: *It does not follow from the fact that judgment must be responsive to reasons that judgment must be aided by second-order competence. I think judgmental beliefs are just beliefs that are subject to rational control, where rational control is just control by deliberation on the reasons. None of that must involve anything higher-order (even implicitly), as far as I can see.* Reply: Consider, first, the

(d) we also gain insight into the special status of the Cartesian cogito.

But let us now consider whether and how we should go further.

The importance of the reflective is not explained *fully* until we see what really matters: namely, that the aptness on the first order be attained *under the guidance* of the second-order awareness. The performance on the first level must be guided to aptness through the apt second-order awareness (explicit or implicit) that the subject is *in that instance* competent to avoid excessive risk of failure. This would comport with the subject's apt awareness that if he performed on the first level, he *would* (likely enough) do so aptly.

Requiring *full* aptness provides a more satisfactory treatment of the complete array of data. The account of human knowledge as requiring knowledge full well attains that further success. We have seen how much better it fares in guiding us through the Gettier problematic. But it is not just its explaining of Gettier data that matters. The account of a desirable level of human knowledge as knowing full well is in fact a special case of something much more general.

The fully desirable status *for performances in general* is full aptness: it is aptness on the first order guided by apt awareness on the second order that the first-order performance would be apt (likely enough). A first-order performance will aim to attain a certain basic objective: hitting a target, as it might be. This will in-

second-order ascent that I have argued to be involved in the nature of judgment. This is not ascent beyond judgment. It is rather ascent in the affirmation that is constitutive of judgment. The relevant affirmation must be in the endeavor to attain truth and aptness. But in order to do something X as a means to an end Y, you must endeavor to bring it about that X contributes to Y. Perhaps one could conceptualize this so that it does not involve one's taking an attitude toward X (which would hence be a second-order attitude). But that would need to be shown. Second, even if not every double omission of affirmation and denial constitutes suspension, it still seems right that every instance of suspension is a case of such double omission, though it may need to be a double omission aimed at taking proper and only proper epistemic risk. Third, to avoid negligent insensitivity to defeaters, one must be poised to respond properly to them. But take an undercutting defeater. It is a reason to forbear affirming specifically, without being a reason to deny. It is thus a reason not to perform a certain way, that is, not to affirm. And this seems a second-order attitude.

duce the correlated aim to succeed *aptly* in that basic aim. Attaining the basic success aptly is better than attaining it inaptly.

Moreover, attaining aptly the aptness of one's success is also better than attaining *this* inaptly. That is a lesson drawn from our basketball player near her threshold of competence, and especially from the amazing shot that spans the full length of the court to score a goal. This success suffers by comparison with that of a shot aimed not just at scoring but at doing so competently enough, even aptly. A shot suffers if it falls short of *full* aptness: that is, if the player fails to guide herself to aptness through knowledge that her shot *would* then be apt. What is needed for these further levels of success is what is still lacked by our player when she stands too near her threshold of reliable-enough competence. She fails to know that she is above the threshold (even if she now is barely so).

The amazing shot that suffers because it is not fully apt is of course still admirable in many ways. It scores the winning points for the player's team, for one thing, and it is quite creditable if it manifests competence far above the ordinary. But the coach is right anyhow in chastising the player for having taken such a risky shot when she need not have done so. So the shot still falls short in being very poorly selected. Moreover, it is not just the temporally prior selection that is deplorable. The more importantly deplorable aspect of the shot is that the decision to shoot is *carried through intentionally*. After all, the player could have reconsidered right up to the point at which she released the ball. What matters most is the player's intentionally shooting in disregard of whether the shot would likely enough be apt.¹² That is, moreover, an important respect of evaluation for performances generally. A performance is poorly selected if it falls below the threshold of reliable-enough competence operative in its domain. Hence, even while succeeding in its basic aim, a performance falls short if it neglects

¹²Objection: *The choice of the shot falls short, and the player falls short for shooting. But the shot was—as you said—amazing. I don't see why it couldn't have been flawless.* Reply: The reason the shot itself falls short, the flaw in it, resides precisely in the choice and in the ensuing sustained and guiding intention (design). This induces a flaw in the shot itself so long as we are thinking of the shot as an intentional action constituted in part by that (admittedly flawed) intention.

attaining that aim *apty*. This latter is an aim also required for full credit in the domain of that performance. What is more, full credit requires that *this* aim, too, be attained *apty*. The performance still suffers if it attains aptness only by luck.

That is so for all first-order performances, in whatever domain, whether cognitive or not.¹³ The case of cognitive performance is a special case in which the aptness of *epistemic affirmation* on the first order is attained *through the guidance* of apt second-order awareness that such affirmation *would* then be apt. *This* case of fully apt performance is just that of *knowing full well*.¹⁴

2. Our advance over simpler versions of virtue epistemology takes two steps. First we highlight that full aptness of performance re-

¹³ Although the emphasis on full aptness seems to institute a potential regress, I can't see that it is vicious. True, as we ascend to the second order we get a boost of epistemic standing (as does a basketball player through enhanced risk assessment based on knowledge of her own competence in the situation in which she now performs). Arguably, one might then get a further boost if one's competence assessment is *itself* not just apt but fully apt. But this need not keep going forever. Returns may in fact diminish quickly to the effect of an asymptotic approach to a limit near to what one reaches already with ascent to the second order. After all, we soon hit a limit at which human competence gives out as we ascend through the higher orders. Beyond that limit, creatures better endowed might attain incremental enhancement not attainable by limited humans. Because *ought implies can*, however, failure to surpass that limit is no human flaw.

The order of a performance is determined by its propositional content and, more specifically, by the number of nested references made in that content to the performer's own possible or actual performances (or competences to perform). Mainly for simplicity, I have restricted myself to two orders, to the first and the second. We must of course allow ascent to higher orders, as I do repeatedly. And we can make terminological room for such ascent by indexing the reference to fullness; so we can have fullness-*i*, or fullness to the *i*th order. And I can see no problem with allowing such possible indefinite ascent.

¹⁴ Granted, animal-level attainments are importantly different from a lucky grasp of gold in the dark (as in our opening analogy from Sextus). The latter is more like a gambler's lucky guess. But this just reveals how epistemic luck can be found at different levels and in different ways. Suppose our seeker of gold enters a dark room that contains only gold objects, and *enters such a gold-containing room unknowingly*. In one way, then, it is no accident that what he pockets is gold, but in another way it surely is. Thus tweaked, the Pyrrhonian gold-in-the-dark example still helps to show why and how one needs a reflective perspective aiding the apt correctness of one's belief while avoiding an important element of blind luck.

quires guidance through the agent's knowledge that his performance would be apt. Through that fact we then, secondly, explain why knowing full well is not just animal knowledge on top of animal knowledge, which explains also our spread of intuitions in the full array of Gettier cases. It is crucial to this further explanatory benefit that in order to know full well the subject must be aptly aware that his affirmation would be apt.¹⁵ Barney fails to know—*fails to know full well*—because his judgment, his affirmation *in the endeavor to affirm aptly*, might so easily have failed.

In falling thus short of knowing full well, does Barney fall short of something epistemically significant? In assenting here, I am guided by a more general, intuitively plausible principle: namely, that in *any* domain of human interest, first-order performance suffers if it is not fully apt.¹⁶ That is a lesson of our basketball shot across the full length of the court, and even of the mid-court shot taken in ignorance that it lies just barely above the threshold of reliability.

Fully apt performance goes beyond the merely successful, the competent, and even the reflectively apt. And it is the human, rational animal that can most deeply and extensively guide his performances based on the risk involved, in the light of the competence at his disposal. That is why reason must lord it over the passions, both the appetitive and the emotional.¹⁷ Opting for such rational guidance does involve judgment and risk in any given case, and a faith in reason as our best guide.

The dimension of cognition is, moreover, just the special case in which our modularly competent seemings must be subject to a rational competence that assigns them proper weights. It is such

¹⁵ And, more generally, that their *representation* would be apt, if we include cases of functional belief, including belief owed to psychological or biological teleology.

¹⁶ So does any potential performance beyond the limit of human competence, where *oughts* lapse for lack of *cons*.

¹⁷ When research reveals the hidden influences that move us, moreover, that may just help us to enhance the competence of our management. Those revelations may lead us to avoid certain situations, competently, *virtuously* so, and to undergo relevant therapy over time, thus counteracting inappropriate influences. The evidential basis for situationism, such as it is, may hence aid virtue without refuting virtue theory.

rational competence that properly determines the balance required for the wise, fully apt choice, which for cognition is the choice of whether to affirm, positively or negatively, or whether to suspend judgment.¹⁸

D. THE EPISTEMIC NORMATIVE HIERARCHY

In conclusion, the epistemic normative hierarchy (on the left below) is a special case of the action hierarchy (on the right), as follows:

| | |
|--|----------------------------------|
| (1) (alethic) affirmation | attempt |
| (2) successful affirmation | successful attempt |
| competent affirmation | competent attempt |
| (3) competent and successful affirmation | competent and successful attempt |
| (4) apt affirmation | apt attempt |
| (5) reflectively apt affirmation | reflectively apt attempt |
| (6) fully apt affirmation | fully apt attempt |

(*Affirmation* here means, again, *alethic affirmation*: that is, *saying that is aimed at truth*.)¹⁹

This highlights how the distinctions and value judgments that we have laid out in epistemology are really just a special case of something with far broader import.

In the next chapter we broaden the scope of our performance-theoretic approach to include mind-world relations other than knowledge, such as perception and action.

Appendix (More Advanced Material): A. Agency and Reflection

1. The epistemic agency involved in our judgment and judgmental belief can be either free or unfree. This aligns with a distinction

¹⁸The case of functional belief would require an analogical extension of this framework based on the shared focus on performances that are *aimed*.

¹⁹But here I have in mind a very broad usage according to which you can say things to yourself “silently,” and according to which, when someone nods, he can mean to say something (and is not just exercising), something specific, such as saying yes to an invitation. (Of course, he does not say it *out loud*.)

between two sorts of knowledge, the fully reflective and the animal. We next take up the nature and significance of these two sorts of knowledge, starting with the reflective.

Reflection has two aspects: first, reflection as careful, conscious thought, as *meditation*, and, second, reflection as thought that turns back on itself, as higher-order thought. The two aspects come together in Descartes's *Meditations* and also in British Empiricism, where reflection is the *operation of the mind by which it is consciously aware of its own conscious contents*.

The two aspects are separable, since higher-order reference to one's own mind can be *subconscious*. So we can distinguish two degrees of reflection. A first degree involves the mind's turning back on itself, whether consciously or subconsciously. A second degree also requires the higher-order thought to be conscious.

2. Both aspects of reflection are important in epistemology—both its conscious and its higher-order character—each in its own way.

The importance of the *higher order* emerges, first, with epistemic suspension of judgment. Suspension is constitutively a second-order mental phenomenon, as I argued earlier and will do so later (in section B). The higher order matters, second, because of our aim to keep epistemic risk within proper bounds in our search for truth. This risk assessment is inherently second-order. It is an assessment of how risky it would be to judge that *p*. Does the risk permit judgment, or does it require suspension instead?

So much for the *higher order*. The importance of *consciousness* emerges through contexts of critical assessment, whether in private thought or in public dialectic, where one must weigh all pertinent reasons while judging in light of the total evidence. Among the relevant reasons as one ponders a question is the fact that one already holds a certain belief on that question, if one does, even one just stored in memory. Take any context of critical assessment: that of legislation, for example, or of the courtroom, the doctor's office, the lab, the criminal investigation, or the philosophy seminar. Questions in such contexts must be addressed by weighing all the reasons in view. Yes, the fact that one already believes a given answer is among the reasons in view as one considers whether to (continue to) believe going forward. But it cannot be allowed

to trump automatically any contrary reasons that may also come into view. Nor will the stored answer necessarily outweigh the newly available reasons *simply because of how very reliable is the perception-plus-memory that diachronically delivers that stored answer*. A believer who has forgotten just how reliably he acquired and retained his belief cannot now draw the belief from storage *with a weight determined simply by its diachronic reliability*. The believer now needs reason to self-attribute such reliability. And this self-attribution will be on the second order and also conscious.²⁰

3. Reflective knowledge of the highest degree involves reflective thought that is both conscious and higher-order. This is the *scientia* that Descartes takes as his epistemic aim. It is knowledge consciously endorsed on the second order as reliable enough, as belief whose correctness manifests superlative competence. Here we have both components required for appropriate reflection of the higher sort. A first-order belief is endorsed *consciously* on the *second order* as one whose correctness manifests superlative competence.

Should we also allow a lesser degree of reflective knowledge requiring no *conscious* awareness? Nearly all one's knowledge remains implicit at any given time, and not all implicit beliefs are epistemically on a par. Only some derive from proper risk assessment, for one thing, and from a proper grasp of one's relevant competence. Some are on a merely animal level, unaccompanied by sufficient assessment of risk or grasp of competence, whether conscious or subconscious. For example, the beliefs of a blind-sighter ignorant of his competence fall short epistemically, despite their animal epistemic standing. We should thus recognize a distinction between knowledge that is merely animal and knowledge that is reflective even subconsciously. This distinction is epistemi-

²⁰True, the degree of confidence of a stored belief may well bear on how much weight it should be given at a later juncture, just in virtue of being so stored with that degree of confidence. But this enhancement of the case will not affect our treatment of it in the main text.

cally worth drawing not just because the knowledge that is thus modestly reflective already seems superior *in that respect* to the knowledge that falls short. Our distinction is epistemologically significant also for a further reason: because main traditional arguments for skepticism threaten reflective knowledge irrespective of whether the second-order endorsement is conscious or subconscious. The threat of vicious regress or circularity applies either way.

4. What is that threat? We attain the reflective level of knowledge only when we self-attribute the competence manifest in the correctness of a first-order belief. It is the ostensibly vicious circularity involved in seeking knowledge of such competence that poses the threat. The targets include first-order sources like perception, testimony, and introspection but also the faculties of arm-chair thought, such as rational intuition and deduction. Reflective knowledge would allegedly involve a vicious circle or regress, because it requires second-order endorsement of the reliability of one's sources, a requirement that supposedly cannot be met without relying on those very sources. This gives us reason to distinguish between animal knowledge, on one side, and reflective knowledge, on the other, whether this latter derives from subconscious or from conscious endorsement. Either sort of reflective ascent would be blocked by the skeptic's argument.

B. Epistemic Negligence

A "shot" is a shooting performance attributable to the agent, perhaps a deed, whether intentional or not; or, alternatively, it is a full-fledged intentional action, an attempt. If the latter, then it constitutively contains an intentional aim: indeed, an aim at a specific target.

What should we say about an action that is negligent? Suppose someone drives drunk and swerves over a sidewalk. There happens to be no one there, so he has good moral luck. Is there, then, nothing wrong with his act of driving drunk on that occasion, over that spatiotemporal stretch? He is just aiming to drive home, and he does drive home. He attains that objective. He also aims

to do so without harming anyone. And he attains this objective as well. So do we say his drive home drunk is a fine act since it harms no one? Is there something wrong only with the agent, while the act itself is just fine?

What, then, makes the agent blameworthy? That he neglects to take into account the risk of harm derived from driving drunk? On this account, what he does is fine. What is wrong is what he fails to do: namely, to take his inebriation properly into account.

In my view, the action itself is wrong, with a wrongness resident in the attendant neglect. It is a negligent action and *for that reason* wrong. The neglect that attends the choice and the resultant intention (design) infects at least one action that our driver performs. It infects his action of driving as he does on that occasion. He does this intentionally, by design. But his choice and ensuing intention are wrong, and the ensuing intention is a constitutive part of that action. Yet this intention is morally wrong and bad, and the action of which it is an essential constituent is thereby wrong and bad. The wrongness derives from the neglect.

What is more, the neglect is on a second order and bespeaks second-order moral incompetence. The agent makes his choice and continues to host the corresponding intention, *in doing which* he fails to take properly into account the risk he runs by so doing. This risk assessment is incompetently done. And it could have been even worse. It could have been that, well aware of the risk, he assessed it competently but ignored it callously and recklessly. So the neglect might have resided not in the assessment but in the ignoring of its result. Either way, the choice is bad, and so is the intention going forward, and so is the act itself, since it constitutively contains that intention.

We shall return in chapter 10 to the similar importance of negligence for epistemology, in connection with Saul Kripke's dogmatism paradox.

FURTHER READING

Greco, John. "A (Different) Virtue Epistemology." *Philosophy and Phenomenological Research* 85 (2012): 1–26.

- Ichikawa, Jonathan Jenkins, and Matthias Steup. "The Analysis of Knowledge." In *The Stanford Encyclopedia of Philosophy* (Winter 2014 edition), ed. Edward N. Zalta. Stanford, CA: Metaphysics Research Lab, Center for the Study of Language and Information. Available at <http://plato.stanford.edu/entries/knowledge-analysis/>.
- Kornblith, Hilary. *On Reflection*. Oxford: Oxford University Press, 2012.

The Value of Human Knowledge

TWO OF PLATO'S best-known dialogues are inquiries about knowledge. The *Theaetetus* inquires into its nature, the *Meno* also into its value. Each dialogue invites the same more basic question: What sort of normativity is constitutive of our knowledge? A belief that falls short of knowledge is thereby inferior. It is better to know than to get it wrong, of course, and also better than to get it right just by luck. What is involved in such evaluation?

The two Platonic problems are closely related. A satisfactory answer to the value problem likely requires a correct view of the definitional problem. We shall now revisit a solution to this latter problem, at least in first approximation, enabling, in turn, an approach to the value problem.

We assume that knowledge requires, at a minimum, a belief that is true. Our inquiry into the nature of knowledge thus takes a more specific form. Our question is this: What condition must a belief satisfy, in addition to being true, in order to constitute knowledge? This question about the nature of knowledge has been central to epistemology in recent decades, as it was for Plato.

The celebrated Gettier problem derives from the fact that the further condition that a belief must satisfy cannot be just its being competently held, competently acquired or sustained. More strictly, this is *a form* of the Gettier problem, its main form as we saw already two chapters back. And here we understand "competence" as short for "epistemic competence." A belief or judgment is thus *epistemically* competent iff it is formed well enough in epistemic respects. This is how we shall think of what so often goes under the label "epistemic justification." What makes such justification *epistemic* is how it relates to belief's aim at truth, which is what we

assume, while taking the label “competent” to fit this conception better than the more customary “justified.”

This Gettier problem arises because a belief can be false despite being competent. If the believer then competently deduces something true from his false belief, this true conclusion cannot thereby amount to knowledge.¹ Yet, if we competently deduce a conclusion belief from a premise that we competently believe (even after drawing the conclusion), we thereby *competently* believe that conclusion as well.

Post-Gettier, our first Platonic problem takes the following form: *What further condition, added to or in place of being competently held, must a true belief satisfy in order to constitute knowledge?*

On the contemporary scene, the second Platonic problem, that of the value of knowledge, has more recently moved to center stage. We can interpret Plato’s *Meno* so that it presents the problem of how knowledge can be quite generally more valuable than its corresponding true belief if a merely true belief would be no less useful. Thus, a true belief as to the location of a certain city, Larissa, will guide you there no less efficiently than would the corresponding knowledge. In line with this, we ask: *How, if at all, can knowledge be as such always better than the corresponding merely true belief?* Harking back to the Platonic dialogue, this problem has been much discussed in contemporary epistemology.

In connection with both problems—the definition and the value problems—we have assumed that beyond being a belief and being true, there is some further condition (however simple or complex) that a belief must satisfy in order to constitute knowledge. We should thus be able to explain how it is that knowledge, which must satisfy this further condition, is as such always better than would be the corresponding merely true belief. When one ponders a question, for example, there is some respect in which it would always be better to answer knowledgeably than to answer correctly but just by luck.

¹ Again, some have argued against such counter-closure, unsuccessfully so in my opinion, as laid out earlier (in a footnote to chapter 4, section 5).

We turn in this chapter to that *Meno* problem based on our earlier approach to the *Theaetetus* problem. Succeeding chapters will further develop and defend our virtue-theoretic account of propositional knowledge, of its nature and its place in human life.

A. THE VALUE OF KNOWLEDGE

The aim of belief is said to be truth. When you sincerely pose a question to yourself, for example, you want a correct answer. When you reach an answer to your question by adopting a certain belief, the aim of your belief is to answer correctly. If the aim of a belief is merely truth, then, if true, that belief would seem to have what really matters epistemically, irrespective of its etiology.

How, then, can a truth-reliably produced true belief be better than one that is no less true, regardless of how reliably it may have been produced? Conclusion: Knowledge is really no better as such than merely true belief.

“Any argument leading to that conclusion,” it might be replied, “must have its premises examined. Perhaps the aim of belief, and of inquiry, is not just truth but also knowledge. This would explain how and why it is that knowledge (with its required etiology) is, after all, better than merely true belief.”

What follows will defend this reply by placing it in context, by explaining its content, and by drawing some implications.

B. A STEP BACK

How, indeed, is truth our aim? How should we understand the value we place on it? More explicitly, our aim is to have the truth. So it is the attaining of truth that has corresponding value. How, then, should we more fully describe our true objective? Is it just the accumulation of true beliefs? Compare how we assess accurate shots, those that hit their targets. What is it that people value under this rubric? Is it the accumulation of accurate shots?

Someone casually draws a large circle on the beach right by his feet, aims his gun, and hits the target. Does he thereby attain, at

least in some small part, a previously standing objective: namely, that of securing accurate shots? Is that an objective we all share, given how we all share the concept of a good shot? Don't we all want good things (other things equal)?

That, I trust we agree, is absurd.

Yet the shot at the beach could be an accurate, good shot nonetheless, as the marksman hits his target in the sand. Although, from one point of view, given the low or even negative value of the target, this accurate shot has little value of its own, yet from another, performance-internal perspective, it is graded as quite accurate, a good shot, maybe even an excellent shot if the marksman steps back far enough from the target. Even when the shot is difficult, however, its status does not derive from any standing preference of people for an accumulation of accurate, difficult shots. There is no normative pressure on us to bring about good shots, not even if we grasp perfectly well what it takes to be a good shot and have this criterion uppermost in our consciousness at the time. There is no normative pressure to bring about even excellent shots, none whatever that I can discern. (What we are *not* normatively pressured to accumulate is *shots*, regardless of how excellent they may be.)

Compare now our intellectual shots, our beliefs. A belief may answer a question correctly but may have very little value nonetheless if the question is not worth asking. The status of a target will bear on the worth of any shot aimed at that target. Arbitrary selection of an area by your feet at the beach yields a silly target. Similarly, suppose you scoop up some sand and proceed laboriously to count the grains. You then take up the question of how many grains are contained in that quantity of sand. If you reach your objective of answering that question correctly, what is your performance worth? Do you thereby fulfill, at least in some small part, a previously standing objective, that of securing more and more true beliefs? This seems about as implausible as is the corresponding view about the shot at the beach. In what way, then, does the truth of our beliefs have value?

One thing that does plausibly have *prima facie* value is the satisfaction of our curiosity. So even if the question as to the number of grains of sand is of little worth, if someone gets interested

in that question anyhow, then the satisfaction of his curiosity will in an obvious way have value to him (and perhaps even, to some small extent, for him). This is, of course, a way for the truth to have value to someone and for someone. After all, if one is curious as to whether *p*, this is just to be curious as to whether it is true that *p*. There are not two instances of curiosity here: (a) as to whether *p* and (b) as to whether it is true that *p*. So what we want when we value the truth in that way is to have our questions answered, and of course answered correctly. Sheer curiosity, whatever its basis, thus invests the right answer to a question with some value, though the value might be small and easy to outweigh, as with the question about the grains of sand. Having the answer to that particular question may add so little to the life of the believer, while cluttering his mind, that it is in fact a detriment, all things considered, if only through the opportunity cost of misdirected attention.

Similarly for the shot aimed from a foot away at the sand on the beach: the sheer desire to hit that target, whatever its basis, gives value to the agent's hitting the mark. But it might well be that hitting that mark imports little value for anyone. Spending his time that way may even be a detriment to the agent's life (even if the sand is soft and yielding enough that the shot poses no physical danger). Nor is it plausible that we humans generally have a standing desire for accurate shots or that we place antecedent value on securing such shots. Accuracy will give value to that shot at the sand only through the gunman's whim to hit that target.

Even if that shot at the beach fulfils no human interest other than the gunman's whim, it may still be a better shot, better *as* a shot, than many with higher overall value. Take a shot at close quarters in self-defense that misses the targeted head of the attacker but hits him in the shoulder and stops the attack. A bad, inaccurate shot, this one, but more valuable than the accurate shot at the beach. (Had it been better as a shot, moreover, a more accurate shot, it might have constituted a terrible murder, since the attack did not justify shooting to kill.)

Are beliefs like shots in that respect? Is a belief a performance that can attain its internal aim while leaving it open whether it has any intrinsic value and whether it serves or deserves any external aim? Let us explore this view of belief.

C. KNOWLEDGE AS A SPECIAL CASE

All sorts of things can “perform” well or ill when put to the test. Rational agents can do so, but so can biological organs, designed instruments, and even structures with a function, such as bridges. A bridge can perform well its function as part of a traffic artery. When a thermostat activates a furnace, it may perform well in keeping the ambient temperature comfortable. When a heart beats, it may perform well in moving some blood. And so on.

A wooden puppet performs well under the control of a puppeteer if its hinges are smooth, not rusty, and well oiled, so that its limbs are smoothly responsive. A bridge might perform well by withstanding a storm. We credit the puppet, as we do the bridge, if its good performance flows appropriately from its state and constitution.

The puppet “performs” (well or ill), as does the bridge, and thus produces “performances.” But it would be quite a stretch to consider it an “agent.” Human beings are different, in any case, if only because we are rational agents. Not only are there reasons *why* we perform as we do. There are also reasons *for* performing as we do, and reasons *that we have* for so performing, and *for which*—motivated by which—we perform as we do. This is not just a matter of having aims in so performing. After all, the thermostat and the heart do have their aims. But they are motivated by no such aim; no such aim gives them reasons *for which*—motivated by which—they perform as they do.²

Human motivation is on another level, even when the performance is physical, as in athletic or artistic performance.

The archer’s shot is a good example. The shot aims to hit the target, and its success can be judged by whether it does so or not,

²True, we could perhaps, just barely, make sense of an extended sort of “motivation” even in those cases, as when a nearby torch fools the thermostat into activating the air conditioner even when the room is already cool. The thermostat still “has a reason” for performing as it does, in some broad sense, a “motivating reason.” Despite the nontrivial resemblance, nonetheless, this is clearly a metaphorical extension, if only because of the vastly greater complexity involved in human motivation.

by its accuracy. However accurate it may be, there is a further dimension of evaluation: namely, how skillful a shot it is, how much skill it manifests, how adroit it is. A shot might hit the bullseye, however, and might even manifest great skill, while failing utterly, as a shot, in a further dimension. Recall the example of a shot diverted by a gust of wind initially so that it would miss the target altogether but for a second gust that compensates and puts it back on track to hit the bullseye. This shot is both accurate and adroit, yet it is not accurate *because* adroit, so as to manifest the archer's skill and competence. It thus fails in a third dimension of evaluation, besides those of accuracy and adroitness: it fails to be apt.

Performances generally admit this threefold distinction: accuracy, adroitness, and aptness. At least so do performances with an aim (even assuming a performance could be wholly aimless).

A performance is better than otherwise for not having failed, that is, for not having fallen short of its objective. In line with that, it is good if it succeeds, if it reaches its objective. A performance is at least good as such for succeeding, even if it is a murderer's shot. The shot itself may still be an excellent shot, despite how deplorable is the broader performance in which it is embedded.

A performance that attains its first-order aim without thereby manifesting any competence of the performer's is a lesser performance. The wind-aided shot scores by luck, without thereby manifesting appropriate competence. It is hence a lesser shot by comparison with one that in hitting the mark manifests the archer's competence.

A shot might manifest an archer's competence without its accuracy doing so. The shot with the two intervening gusts is a case in point. How does that shot manifest the archer's competence? By having at the moment of release an angle, direction, and speed that would take it to the bullseye in relevantly normal conditions.

A blazing tennis ace is a lesser shot if it is a wild exception from the racket of a novice, by comparison with one that manifests the superb competence of a champion in control. And so on. Take any performance with a first-order aim, such as the archery shot and the tennis serve. That performance, then, has the induced aim of *intentionally (by design) attaining* its first-order

aim.³ A performance X attains its aim <p> not just through the fact that p but through its bringing it about that p.⁴

The case of knowledge is just the special case in which the performance is cognitive or doxastic. Belief aims at truth and is accurate or correct if true. A guess can also aim at truth, however, if that is how a contestant will win a prize. Belief aims not just at accuracy (truth) but also at aptness *period* (knowledge of a sort). A belief that attains both aims, that of truth and that of knowledge (of that sort), is for that reason better than one that attains merely the first. That, then, is a way in which knowledge seems as such better than merely true belief.⁵ But let us take a closer look.

D. THE VALUE OF HUMAN KNOWLEDGE

1. In what ways is knowledge valuable, and how well does virtue epistemology accommodate its value? How might our virtue epistemology accommodate generally accepted platitudes about the normative standing of knowledge and about its superiority to mere true belief?

³As I argued in my *Judgment and Agency* (Oxford, UK: Oxford University Press, 2015), the intentional (by design) attainment of an aim requires some degree of aptness, hence some correlated degree of competence, but not necessarily a degree sufficient to meet the standard of reliability operative in the relevant domain (the domain of one's performance) for competence *period*, nor correlative for aptness *period*. Even a neophyte archer might attain his hitting the target, despite lacking the competence *period* that would make his hit *apt period*. But he cannot attain any such thing if he has *no* degree of competence at all, so that he lacks even any idea of the direction in which to aim in order to hit a target in that direction, or of which way to move the string of the bow so as to speed the arrow on its way, and so on.

⁴Just as its being true that p entails its being true that it is true that p, so one's bringing it about that p may entail that one brings it about that one brings it about that p, assuming such iteration always makes sense.

⁵And here we have an account of why knowledge is better than merely true belief, since apt performances, in general, are as such better than those that attain success only by luck. So beliefs are just a special case of that general truth. This account still depends, of course, on our view of knowledge as apt belief, belief that manifests the relevant competence of the believer in reaching its aim of truth.

2. Here are some thoughts about knowledge, belief, and truth with respect to yes/no questions taken up “alethically,” *with the aim of answering them correctly*:

- (a) Any question taken up alethically is better answered knowledgeably than ignorantly.
- (b) It is also better answered correctly than incorrectly. Correct belief is better than error.
- (c) And it is better answered competently than incompetently.
- (d) Finally, it is better answered *both* correctly and competently than while falling short *either* way.

These are platitudes that our epistemology would do well to help explain. And here, now, is a proposal for doing so within virtue epistemology.

Each of those four platitudes is a special case of a more general platitude, as follows.

3. Take any attempt to attain a given aim.

- (a) Any attempt is good as such if it succeeds aptly and better than a corresponding failure.⁶
- (b) An attempt is better as such if it succeeds, however aptly, than if it fails.
- (c) An attempt is better as such if competent than if incompetent.
- (d) It is better in an attempt *both* to attain success and to do so competently than to fall short *either* way, either by not succeeding at all or by succeeding *incompetently*.

4. That shows a way in which the performance normative of virtue epistemology can help throw light on our platitudes 2a–2d. But note how implausible 3a–3d seem when interpreted as principles of *objective axiology*. Obviously, the success of a murder attempt is not necessarily better than its failure. It is not even *prima facie* or *pro tanto* objectively better, in virtue of its being a success. Indeed, failure in a murder attempt would most often or

⁶One might fail to do so by trying without success but also, for example, by negligently not trying.

always be better than success, and *incompetence* in the execution of such an attempt would most likely be better than competence, even *prima facie* or *pro tanto*.

That is analogous to the axiological ideas that pleasure in the good is incrementally good, pleasure in the bad incrementally bad, while displeasure in the bad is good and displeasure in the good bad. The plausibility of these principles remains when we switch to intentional attempts in place of the attitudes of favoring or being pleased that such and such.

Thus, intentionally bringing about the bad seems incrementally bad. The pleasure involved when one (factively) takes pleasure in the bad does not add positively to the overall value of the factive state. It may even add negatively. Similarly, then, the intentional bringing about of the bad seems to reduce rather than enhance the overall value of that factive state.

Apt judgment or belief would thus have incremental objective value only if true belief as such has positive objective value. But I myself doubt the objective value of having truth, just as such, *any* truth. There is no objective, intrinsic value in trivial enough truth.

5. That is how we speak with plausible truth in saying that knowledge is always *necessarily* better than would be the corresponding merely true belief. Knowledge is not *objectively* better by necessity, as we have just seen. The plausible truth expressed must be more subtle than that. What is plausibly true is, for one thing, rather this: that, *necessarily*, the truth seeker would always proceed properly if he preferred knowing to merely believing correctly. After all, for *any* endeavor one undertakes, it is always *necessarily* proper to prefer that one succeed, and indeed succeed aptly, not just by luck. Our beliefs, our endeavors after truth, are just a special case. It would always be proper to prefer one's attaining one's aim, and attaining it aptly, not just by luck. (That is, one would *properly* have that preference *at least in one respect*, since it is the preference required for rational coherence.)

6. Might knowledge be more valuable than true belief in any more objective way? Well, knowledge of *certain matters* does add importantly to the flourishing of one's life individually and of life

in community. Mere true belief on those important matters does not contribute as much as does the corresponding knowledge. However, not *every* instance of knowledge, *regardless of its subject matter*, adds in those important ways. Not *every* instance of knowledge adds to the flourishing of that knower or community. Knowledge might *still* be better, more valuable, than true belief simply because of how knowledge *of certain important matters* contributes to human flourishing, more so than would the corresponding merely true belief.⁷

7. Thus, the account of epistemic normativity as a sort of performance normativity has two virtues. First, it provides an explanation of the nature of knowledge, which amounts to belief that is apt, belief that is an apt epistemic performance, one that manifests the relevant competence of the believer in attaining the truth. And second, it also helps explain how knowledge can be plausibly thought to be “better” than merely true belief.

E. HOW KNOWLEDGE ALSO MATTERS: A FURTHER DIMENSION

1. Finally, we shall explore a further way in which knowledge can be crucial for the proper value of much that matters most in human life. We turn to how knowledge can matter because of what it constitutes rather than because of how it is itself constituted.

Distinguish between the following two situations:

Situation 1: Paul’s favoring Mary’s enjoying life, when, *as it happens, unbeknownst to Paul*, she is indeed doing so.

Situation 2: Paul’s taking pleasure in Mary’s enjoyment of life.⁸

⁷Here the important subject matter could be important in general or important to that subject in particular. Compare how plausible it is to think that health, wealth, love, and friendship are important goods, and yet how obvious it is that not every instance of these is a good to the one who has it. These can, of course, contribute enormously to the flourishing of a life or of a group, but the specifics have to be right. It all depends. And the same goes for knowledge.

⁸Such taking of pleasure seems clearly not to be a mere hosting of sensory pleasure (whatever the source). But in any case, “satisfaction” could replace the word “pleasure” with little to no effect on the argument.

Clearly the second is better, significantly better. This is of a piece with Aristotle's emphasis on full-fledged action and how it constitutes flourishing. Full-fledged *passion*, beyond the false passions of a Matrix scenario, is similarly important.

Wherein resides the further value? We need a further attitude beyond merely abstractly favoring M (the possible state of affairs of Mary's enjoying life), with no idea whether M obtains (is actual) or not. *Taking pleasure in* goes beyond that, surely. How so? Try this:

S takes pleasure in M iff S favors M and M obtains.

No, this seems obviously inadequate. S doesn't take pleasure in M if M is simply true without S's having any opinion on the matter. So perhaps we should expand as follows:

S favors M, M obtains, and S believes that M obtains.

No, that still falls short. S may mistake Nancy for Mary, and only thus may he then believe that M obtains. It seems we need to go at least as far as the following:

S takes pleasure in M iff S favors M and S *knows* that M obtains.

But this *still* arguably falls short. S may know "deep down" that M obtains, although what really drives his present favoring is *not* that bit of knowledge but something else, maybe even his temporary mistaking of Nancy (obviously enjoying herself) for Mary. So we need something further. Perhaps the following:

S takes pleasure in M iff S favors M *because* S knows that M obtains.

No, this still seems wrong. S might favor M just the same even if he did not know that M obtains. So it need not be true that he favors M *because* he knows that it obtains.

2. That line of thought has led nowhere; let us try a different tack.

Suppose we recognize the following adverbial modifications of "believes": believes approvingly, believes disapprovingly, believes indifferently. Whatever *believes approvingly* may be, it is not just

the conjunction of *believes* and *favors*. One might believe that *p* and favor the possibility that *p* without putting two and two together. One believes it implicitly while forgetting at the moment, and while then favoring it, and in this case one might *not* believe approvingly.

And now, given that attitude, we can try this:

(PKA) *S* takes pleasure in *M* iff *S* knows approvingly that *M* obtains.

Plausibly, the value of *S*'s taking pleasure in *M* is more than the sum of the values of its components. Suppose that *M* has great positive value (in part because Mary is deserving). In that case, *S*'s taking pleasure in *M* will have value that derives from the value of its constituent *M*. But it does seem plausible that *S*'s taking pleasure in *M* has *greater* value than just the value of *M*. So it has "organic value," value over and above the measures of value that derive from the value of *M* and the value of *S*'s degree of pleasure. Here the value of the whole is organic in being greater than the sum of those parts. It is a further value that derives from the appropriateness of *S*'s pleasure taking *M* as its object.

That shows a way in which knowledge is of pervasive value in human life. Moreover, the additional value involved in PKA involves more than the mere valuing of the valuable, and more along several dimensions. It seems more in line with an Aristotelian approach whereby (in opposition to the Stoics) there is more to what makes life worth living, more to flourishing, than the inner value of getting our attitudes in proper internal order. Nor is the additional content of the human good just a matter of efficient causation. While involving causation in some broad sense, it does so in the way of competence. Thus, as soon as knowledge appears, as in PKA, competence comes along, and so does the relevant sort of causation.

What is more, one might well reject that *S*'s taking pleasure in *M* *does* necessarily have value greater than its component *M*, or even any value at all, despite the value of *M* (given how deserving Mary is). Isn't that going to depend on how appropriate is one's *approvingly knowing* that *M* obtains? Maybe the reason why one *approvingly knows* it is only that one thinks it will cause untold un-

deserved suffering, through some god's wrath. In the end, we will likely need to consider whether the approving is apt or not. But the knowing seems required in any case.⁹

Next we broaden the scope of our performance-theoretic approach to include mind–world relations other than knowledge, such as perception and action.

FURTHER READING

Hyman, John. *Action, Knowledge and Will*. Oxford, UK: Oxford University Press, 2015.

Kvanvig, Jonathan. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge, UK: Cambridge University Press, 2003.

Pritchard, Duncan, and John Turri. "The Value of Knowledge." In *The Stanford Encyclopedia of Philosophy* (Spring 2014 edition), ed. Edward N. Zalta. Stanford, CA: Metaphysics Research Lab, Center for the Study of Language and Information. Available at <http://plato.stanford.edu/entries/knowledge-value/>.

Zagzebski, Linda. "Intellectual Motivation and the Good of Truth." In *Intellectual Virtue: Perspectives from Ethics and Epistemology*, ed. M. DePaul and L. Zagzebski, 135–55. Oxford, UK: Oxford University Press, 2003.

⁹Someone might question whether S's taking pleasure in M has more value than one of the earlier versions, such as "S favors M, M obtains, and S believes M obtains." But here I would guess that most of us would prefer real interactive experienced enjoyment to the false pleasures of a Matrix victim or a BIV (a subject with a brain in a vat). I at least would not hesitate, and this seems to reveal my view of a difference that really matters; it does not seem just an idle, idiosyncratic preference on my part, like my preference for chocolate over vanilla.

Mind-World Relations

ACTION, PERCEPTION, KNOWLEDGE

1. Metaphysical analysis goes beyond conceptual or semantic inquiry and also beyond necessary biconditionals, which can fail to provide the metaphysical explanation of special interest in philosophy.¹

Recall our earlier example of the metaphysics of persons. In the broad domain of persons, we find a threefold divide among (a) words, such as the word “person”; (b) concepts, such as the concept of a person; and (c) extralinguistic, extraconceptual entities, the living persons. Concerning the latter we find various metaphysical options. According to one, an Aristotelian view, a person is never identical with, but only constituted by, a body. The body needs to be alive and in possession of certain powers and abilities in order to (thereby) constitute a person.

This Aristotelian view on the metaphysics of persons involves metaphysical dependence. One thing exists or is actual dependently on certain other things and their properties. The dependent thing then exists or is actual dependently on the other things and on how they are propertied or related.

2. Moreover, an analysis might still be helpful even if it is only partial. Take the familiar example of a cat’s lying on a mat. Consider first the pragmatics and semantics of the corresponding sentence. For example, when exactly must the cat be on the mat in order for an utterance of that sentence, which takes time, to be true?

¹We shall soon consider what is distinctive of such philosophical explanation in at least some of its varieties. In section B7 of chapter 4, we considered an example of a biconditional concerning the Platonic solids: a necessary biconditional that nevertheless has no explanatory potential for the seeker after metaphysical understanding of what constitutes a cube or a cubical box.

Consider next the analysis of the sentence's conceptual content. One sort of conceptual analysis would involve a biconditional claim of necessity, with the target content on one side and the explaining content on the other. On pain of vicious circularity, the explanation could not contain the content to be explained. Grasp (understanding) of the target content would be explained through (prior) grasp of the explaining content.

The metaphysical analysis of *the cat's lying on the mat* is distinct from such linguistic or conceptual analysis. A metaphysical analysis, with respect to that cat and that mat, of the former's lying on the latter, would involve *the cat*, *the mat*, and a certain binary relation, that of *lying on*. But what is this relation? We might try this: it is a relation of being adjacent to and above. But that is incorrect. Suppose the cat lies on a mat that is glued onto the inside and bottom of a large wheel. And suppose the wheel starts to rotate rapidly so that the cat stays on the mat throughout the rotation due to the forces involved. The cat continues to lie on the mat all through the rotation, but at the top is *below* the mat.

One reaction to the example would be to try to explain how the relation of *lying on* is a more complex relation than it might have appeared. In doing this one might appeal to resultant forces, so that the actual relation involves previously unsuspected factors. Without the concept of force, one would not be able to entertain that proposal. Yet one could still gain a partial account of the relevant relation as follows: it is a relation of being adjacent *in the right way*. While unsure of what that way is, one might at least know that it is some way of being adjacent.

Such partial metaphysical analysis might or might not correlate with similar success involving linguistic or conceptual analysis. In any case, the subject matter is different. Our target concerns a state in the world of *the cat's lying on the mat* or, more generally, a *sort* of state, instances of which we find in the world, of one thing's lying on another. And this is different from any words or any concepts.

Similar distinctions apply generally in philosophy. Thus, consider the linguistic analysis of the terminology of "persons" and the conceptual analysis of a concept of a person, and how these differ from metaphysical inquiry into the nature of persons, such as the living human beings among us.

3. Turning to epistemology, we can discern three problems that are distinct, however closely they may be interrelated. First is a problem of semantic analysis: What is the semantic analysis of the linguistic expression “S knows that p”? Second is a problem of a certain sort of conceptual analysis: With concepts understood either as meanings or as psychological entities, what is involved in someone’s possession and/or deployment of a given concept? This problem thus concerns people’s minds, their psychology, their mental states or the contents of those states. Third is a problem of metaphysical analysis.² Here our focus is on an objective phenomenon that need be neither expression nor concept. Our focus is rather on a state that people host or an act that they perform. This is the phenomenon whose ontology we now wish to understand. What is the nature of human knowledge, and how is it grounded? In virtue of what is it actual when it is actual?³

4. Performance whose success manifests the relevant competence of the performer avoids thereby a kind of luck. According to competence virtue epistemology, knowledge is a special case of that. Knowledge of a sort is belief whose correctness is attained through epistemic competence, belief that is thus “apt.” This account of a sort of knowledge has a serious problem similar to one that afflicts the metaphysics and ideology of perception and action. This is the problem of deviant causation.

We shall consider a solution for the problem in its three varieties. But first we examine accounts of action, perception, and knowledge. And we explore how our take on the point and substance of metaphysical analysis bears on the problem and on competing reactions to it.

²Thus, how persons themselves are constituted is a question other than that of the conceptual content of the *concept* of a person.

³Here I have lumped together questions of grounding, questions concerning the *in virtue of* relation, and questions of *nature, essence, or constitution*. I am leaving open whether these various ontological issues should be distinguished. These are the issues at the focus of my own interest in the ontology of persons, as in “Subjects among Other Things,” *Philosophical Perspectives* 1 (1987): 155–87. Questions of grounding have recently attracted intense attention among metaphysicians such as Kit Fine, Gideon Rosen, Jonathan Schaffer, and others.

What follows is divided into three parts. In the first part, the main lines of the view are laid out and it is shown how it applies with the same basic structure in all three domains. The second part then considers how our account improves on earlier main alternatives. The final, third part reflects on a methodology that will fit our approach.

A. THREE MIND-WORLD RELATIONS

1. Action

What is it to act intentionally? As a first approximation, you might think, to act intentionally is to succeed in a certain intentional aim (by *design*), where the success is owed to the agent's intention.

But that has counterexamples, such as the following:

A waiter intends to startle his boss by knocking over a stack of dishes right now, which makes him so nervous that he involuntarily staggers into the stack and knocks it over, thus startling the boss. But this is not something he does intentionally.

So we should require that the agent's intentions must bring about the success in the right way, with "the right kind of causation." This is the view of Donald Davidson.⁴ The waiter's knocking over of the dishes is not caused in the right way by any such intention. But no account of "the right way" has won consensus.

⁴Davidson advises accordingly in his long struggle with the problem and in his parting thoughts on the matter. Here is how he puts it (with minor variations): "What is it for an agent to F intentionally on a particular occasion? There must be some G such that the agent's intending to G must cause . . . in the right way, the agent's particular act of F-ing." See p. 221 of his "Reply to Vermazen," in *Essays on Davidson: Actions and Events*, ed. Bruce Vermazen and Merrill B. Hintikka (Cambridge, MA: MIT Press, 1985). Davidson's thought evolved from "Actions, Reasons, and Causes," *Journal of Philosophy* 60 (1963): 685–700, to "Intending," in his *Essays on Actions and Events* (Oxford, UK: Oxford University Press, 1980), and then to his replies in the Vermazen and Hintikka collection.

2. *Perception*

a. What is it to perceive an entity? Here is an early answer:

X perceives M iff X has a sense experience that is causally dependent on some state of affairs involving M.⁵

This, however, is subject to counterexamples. Our visual sense experiences while in the sunlight, for example, are causally dependent on the sun even when we look away from it. Nor do we perceive our eyes even when our visual experiences are highly dependent on the state of our eyes.⁶

b. The account is then revised to say that an object is perceived iff some condition involving it is a differential condition that affects some but not all of the perceiver's relevant sense experience at the time of perception. The sun is not seen when we look away from it; on the revised account, this is because no condition of it affects only some and not all of one's visual sense experience.

However, the revised account, too, has counterexamples. Spotlights can shine respectively on statues viewed concurrently, each spotlight thus affecting the perceiver's visual impressions differentially, though only the statues are seen, with the spotlights blocked from view.

c. Thus we eventually arrive at approximately the following view:

X perceives M iff X hosts a sensory experience for which M is causally responsible in the right way.

This is what Paul Grice's view comes to, given how we are said to grasp the "right" way, to be considered later.

3. *Knowledge*

In first approximation, propositional knowledge can be understood as belief that attains its aim (truth) and does so not just by luck but through competence. Such knowledge, then, is a special

⁵This is drawn from H. H. Price's *Perception* (London: Methuen, 1932).

⁶See Grice, H. P., "The Causal Theory of Perception," *Proceedings of the Aristotelian Society* 72 (1961 suppl.), 121–53. The counterexamples in our main text come from this work of Grice, as does the alternative view developed.

case of performance that is not just lucky but apt: that is, performance whose success is owed to the performer's relevant competence. The aptness of a performance is thus supposed to block an important sort of luck, the sort that precludes Gettiered subjects from knowing what they believe both correctly and competently. A belief falls short of knowledge when its truth is owed to luck and not to the believer's competence.⁷

That view of the apt (not just lucky) performance has ostensible counterexamples. Take an archer's competent shot that (a) would hit the target absent intervening wind and (b) does hit the target because, although a first gust diverts it, a second gust puts it back on track. Here the agent's competence yields the early orientation and speed of the arrow, and this combined orientation and speed, together with the two compensating gusts, results in the bullseye. So why is this shot not apt after all? A performance is apt when it succeeds because of the agent's competence. But our archer's wind-aided shot does seem to succeed because of his competence! If the agent's competence had not resulted in the right orientation and speed upon release from the bow, then the arrow would not have hit the target.

Taking a leaf from Davidson and Grice, we might judge success to be apt only if it derives causally from competence in the right way. Success essentially aided by lucky gusts of wind would not derive in the right way from the archer's competence.

4. *Assessing the Three Accounts*

a. All three accounts may be rejected as unsatisfactory until we are told what it is for success to derive "in the right way" from the relevant causal sources.

b. We are considering accounts of phenomena that are broadly "factive," such as perceiving *x*, killing *x*, perceiving that *p*, intentionally ϕ ing, and knowing that *p*. These involve relations spanning mind and world, relations between the subject/agent's mind

⁷We have reached the view that knowledge is true belief out of intellectual virtue, belief that turns out right by reason of the virtue and not just by coincidence" (Ernest Sosa, *Knowledge in Perspective* [Cambridge, UK: Cambridge University Press, 1991], 277).

and her environing world.⁸ Philosophical analyses of these various relations are then proposed. And the analyses repeatedly appeal to some essential causal relation.

Thus we reach the nub of the problem. The problem is often posed by deviant causation, wayward causation that gives rise to counterexamples whether the analysis targets action, perception, or knowledge. Time after time, a kind of “luck” derives from the deviant character of the causation, incompatibly with appropriate success and relevantly creditable perception, action, or knowledge.

c. For all such “factive” phenomena, there is a good case and a bad case. In the good case, the agent fully succeeds. In the bad case, she fails in some way or other.

All three accounts—that of intentional action owed to Davidson, that of perception from Grice, and the account of knowledge as apt belief—are analyses of the good case into factors, and in all three there would be a highest common factor shared by the good case and the bad case: respectively, an intention, a sensory experience, and a belief. But in none of them would this highest common factor figure as a conjunct in a conjunctive analysis of the good case into independent conjuncts, as in the analysis of knowledge as justified, true belief.

Why is there no such analysis? The reason is uniform across the three cases, as they all involve a causal connection said to be present in the good case and absent in the bad case. In each case, the causal connection is just not relevantly detachable from other factors conjoined with it, and with each other, in the analysis.

No state of affairs that comprises relata related by causation will have a metaphysical analysis such that no two factors constitutive of the whole are necessarily related by logical or metaphysical necessity. X's causing Y thus comprises X and Y as relata related by causation, but the whole causal state of affairs is not fully resolvable by analysis into logically and metaphysically independent factors. Even if factors X and Y are logically and meta-

⁸Why “factive”? In the case of propositional attitudes, because the attitude must necessarily have a fact as its object and in the case of objectual attitudes because it must necessarily have an existing individual as its object.

physically independent, there is no way of adding a further factor independent of these two that will secure the required causal connection crucially involved in X's causing Y.

d. Some disjunctivist opponents of traditional analyses argue that the relevant entity (say, knowledge) that is traditionally thought to have an analysis is in fact prior, so that any proper analysis would follow the reverse order. The upshot of our discussion is, however, that if causal accounts of perception, action, and knowledge are to be rejected in favor of disjunctivist views, the objection will need to go beyond any argument against analysis based on the assumption that proper analyses must be conjunctive analyses of mutually independent factors. In order to clinch their case, opponents of traditional analyses must argue more fully than has been done to date. They must show not only that there is no factorizing analysis of the relevant phenomena into independent conjunctive factors but also that there is no acceptable causal analysis.⁹

Suppose even that essential appeal to "the right way" spoils semantic and conceptual analysis. Consider the view of "firsters," who believe that there is no such linguistic and conceptual analysis in any of our three domains: that of perception, that of action, and that of knowledge. Whether we are in a position to give it

⁹A more specific objection might be lodged as follows: *A central case for disjunctivism about objectual seeing is the argument that disjunctivism gets right the phenomenology of perception, which is a phenomenology that strongly recommends naïve realism: it is the transparent phenomenology of acquaintance with objects in one's environment, where these objects seem to be constituents of the experience. Isn't this a further argument pressed by disjunctivists, at least in the case of perception?* Reply: This argument, too, is inconclusive, however, because there is no such sensory phenomenology. There is only a "phenomenology" of seemings, inclinations to believe or the like. This is the sort of seeming in play, surely, when it "seems" to anyone that objects seen can be constituents of the corresponding visual experience. But such a prejudgment "seeming" has scant probative force to countervail how plausible it also "seems" that the same sensory phenomenology can be shared by a hallucination and a veridical experience. In addition, the believer in purely subjective sensory experience need not deny that there is a sort of sensory experience that is constituted by a seen physical object. For he can also insist that this sort of experience has a metaphysical analysis in terms of the subjective experience involved in such objectual experience.

in full detail or not, however, there might still be a metaphysical analysis, even if the formulation of the analysis must make use of “in the right way.” Any formulation we could give might then have to be partial, not complete. Recall in this connection our cat-on-the-mat example.

Firsters thus owe further argument that there is no metaphysical analysis of perception, action, or knowledge into phenomena metaphysically more fundamental. Such metaphysical analysis is not precluded even if there is no interesting, noncircular, informative semantic or conceptual analysis of the words or concepts in the relevant domains (that of action, that of perception, and that of knowledge).

5. An Approach through Performance Theory

a. What follows aims to turn the tables on objections to traditionalist causal analyses. The use of a concept of “manifestation” will enable causal analyses in all three cases. Appeal to manifestation helps us to develop a better solution to those problems. The notion of aptness (success that manifests competence) promises to be helpful not only in the theory of knowledge but also in the theory of action and in the philosophy of perception.

Both Davidson and Grice make a crucial move in defending their respective accounts. Even though their formulations are different, the move is essentially the same. They both, in effect, require a particular sort of causation while ostensibly assuming that no verbal formula can nontrivially define it. Davidson then says that no such formula is needed, and Grice adds that a grasp of the right sort of causation can be attained through examples. Let us have a closer look.

b. Recall the waiter who intends to knock over a stack of dishes right now but does so only through an attack of nerves caused by the nervy intention. Why is this not a way in which a doing can relate to an intention so as to constitute intentional action? What is the required causal relation? Can it be defined so as to reveal why the waiter’s doing does not qualify? Davidson claims that we need no armchair analysis of this matter. In his view, intentional actions are analyzable as doings caused by intentions in the right

way, and no further explication of the right way is possible or required. We might ask: "No further explication is required for what?" And here is one plausible response: We need not provide a further explication (of what that "right way" is) in order to make any progress. We can at least partially formulate an analysis of intentional action through appeal to appropriate causation "in the right way."

Still, it would be nice to be able to make further progress by going beyond such "appropriateness," beyond invoking "the right way."

Let us try an account in terms of competences and their manifestations. Consider:

Knowledge is apt belief.

Perception (propositional perception, perception that such and such) is apt perceptual experience, experience whose success manifests competence. A perceptual experience succeeds when it is veridical or accurate. An apt experience is one whose accuracy manifests the relevant competence of the subject's perceptual systems.

Action is apt intention.

In all three cases, the following factors come to the fore:

Success, the attainment of the aim.

The competence of the performance.

The aptness of that performance: whether the success manifests competence.

And it is not just an accident that competence is the key to "the right way." Again, all three human phenomena involve aimings, performances with an aim. Perception involves functional, teleological aimings, through the teleology of our perceptual systems. Intentional action involves aimings that are full-fledged intentions. Knowledge divides into two sides: a functional perception-like side, and a judgmental action-like side.

The sort of causation essentially involved in all three phenomena is hence the causation of aptness. It is not enough that the success be just caused by competence, for it may be caused

deviantly, by luck. Rather, the success must be apt. It must *manifest* the performer's competence.¹⁰

6. Preliminary Conclusion

We find unity across action, perception, and knowledge. All three are constituted by aimings, by performances with a constitutive aim. In perception the aim is functional through the teleology of our perceptual systems.¹¹ The aim of an intentional action is obvious in its constitutive intention. Knowledge comes in two sorts. One is functional, so that its aim can be teleological, like that of perception. The other is rather like that of intentional action, and that is because judgment is a kind of action, with judgmental belief the corresponding intention. When causation figures in the right way in all three of these phenomena, it is hence the causation of aptness.

B. GOING BEYOND GRICE AND DAVIDSON

1. Understanding and Ineffability

a. Our account goes beyond Grice or Davidson by specifying, in a performance-theoretic way, what the "right way" is in which causation must join together the relevant items: intentions with doings in intentional action, sense experiences with objects in perception, and beliefs with truth in knowledge.¹²

¹⁰Where this competence might be subagential, such as in perceptual competence.

¹¹But is the teleology of our perceptual systems really veridicality of experience? Is it not rather representation of the world to us in such a way that it allows for successful and cognitively efficient interfacing with the world or some such more vital end? Although I lack the space required for a full discussion of this fascinating issue, let me just suggest a plausible compromise. Perhaps the right teleological aim is to represent properly so as to further such success and efficiency of interfacing.

¹²(Here "object" is used broadly to cover any "object of perception.") We also go beyond Grice by placing perception in the domain of biological and psychological proper functioning (functioning that is in that extended sense "competent")

We would like to understand the metaphysics and epistemology of action, perception, and knowledge, which we must attain through certain concepts, even when these are not helpfully expressible through verbal formulas. Nor need they be thus expressible even when widely shared among us.

Really? How do we understand those oracular claims?

b. Just compare how we manage to grasp what politeness is, what it requires. No verbal formula can fully convey or determine (by explicit convention) what is or is not polite conduct. Polite face-to-face conversation sets limits to the proper distance between the partners and limits the volume of voice and the tone. How is any of this to be captured nontrivially through verbal formulae? It seems quite hopeless. Yet somehow, antecedent community convention sets those limits. Such convention requires antecedent agreement, at least implicit agreement, which in turn requires shared content, which must apparently be shared even without the sort of formulation that would be required for explicit conventional agreement.

Compare the “manifestations” of a competence. A community might similarly agree (however in the end we understand such implicit “agreement” and its content) on what are cases of “manifestation” of a given competence, even with no helpful verbal formula to cover all such cases. This is like “politeness,” in general and in specific respects. Now a complete competence can be broken down into three components: the relevant Skill, Shape, and Situation. Consider such SSS competences, our concepts of these, and the induced SS and S correlates. Take, for example, our complete driving competence on a certain occasion, including (a) our basic driving skill (retained even when we sleep), along with (b) the shape we are in at the time (awake, sober, and so on), and (c) our situation (seated at the wheel, on a dry road, and so on). Drop the situation and you still have an inner SS competence. Drop both shape and situation and you still have an innermost S competence:

and in the domain of biological and psychological performances that satisfy the AAA structure of accuracy, competence, and aptness.

the basic driving skill that one retains even asleep (in unfortunate shape) in bed (in an inappropriate situation).

Such concepts are broadly shared with no benefit of linguistic formulation.

What counts as manifestation seems also graspable only in implicit ways, as in the case of etiquette, and not through explicit (and nontrivial) verbal formulation.

2. *Competences, Dispositions, and Their Manifestations*

a. It is plausible that the competence to drive on a certain road comes in three varieties: skill (basic driving competence), skill + shape (skill plus being awake, sober, and so on), and skill + shape + situation (skill plus shape plus being at the wheel of an operative car while the road is dry, and so on). Only with the SSS competence are we fully competent to drive on that road. What determines whether we have the innermost S competence? It is presumably a modal matter: very plausibly what is implicated is that if we tried to drive safely we would reliably enough succeed. But in any shape or condition? Surely not. It is not at all likely that we would drive safely, even if we tried, when dead drunk or on an oily road. But this does not bear on our competence to drive safely on a given road. There is an array of SSS conditions that would likely enough ensure the success of our attempts to drive safely on that road. This would involve certain ranges of the shape we need to be in and of how we must be related to the road, including the condition of the road. Communities that use cars and roads are interested in certain particular combinations of shapes and situations, and we are pretty well implicitly agreed on what those are. Innermost driving skill is then determined as the basis for our likely enough succeeding if we tried in *those* shape + situation combinations.

b. We have a large and varied array of common-sense dispositional concepts: fragility, flammability, malleability, and so on. These all can then perhaps be understood in terms of our SSS structures, along with relevant triggers and outcomes. For an outcome behavior of an object to manifest a given disposition, then, is for it to flow causally from its triggering event involving the

object, when the object has the relevant skill/seat and is in the relevant shape and situation. What are the relevant shape, situation, trigger, and outcome associated with a certain dispositional concept? This may simply not be formulable in full explicit detail by humans who nonetheless agree sufficiently in their grasp and deployment of the concept. A particular disposition, then, will have a distinctive SSS profile, with restricted shape and situation. Not every disposition to shatter amounts to fragility! Thus, suppose an object's disposition to shatter upon its impact with the hard floor is dependent on the presence of a powerful zapper ready to zap any nearby object when it hits the floor. Such a zapper-dependent disposition would not count as fragility. An iron dumbbell is not made fragile by the presence of such a zapper, even though the truth of the relevant conditional is now dependent on the zapper. It is because of the zapper that objects dropped nearby would shatter upon their impact with the floor. Only the zapper ensures the truth of that conditional even as concerns nearby iron dumbbells.

c. But why should we have all this implicit agreement on how to categorize dispositions and their special cases, such as abilities, and in turn competences? Why do we agree so extensively on whether an entity's output is to be attributed to a disposition hosted by that entity as its manifestation and, by extension, attributed to the entity itself?

When the output is good, it is then generally to the entity's credit, when bad to its discredit. The entity might be an agent who manifests a competence, or it might be a lifeless patient manifesting a mere disposition. Why do we agree so extensively on these dispositions, abilities, and competences; on the credit and discredit that they determine (whether this be to the credit of a moral agent or to that of a sharp knife); and on the sortals that they help constitute?

Is that not all just part of the instrumentally determined common sense that humans live by?

Such common sense helps us keep track of potential benefits and dangers and how the bearers of these are to be handled. As a special case of how to properly handle things and agents that

manifest dispositions and competences, we have the propriety of encouraging praise or approval or discouraging blame or disapproval, which in turn helps to determine the relevant dispositions, abilities, and competences in ourselves and in our fellows.

Such an instrumentally determined common sense must, of course, be structured against background implicit assumptions about what is normal or standard, either in general or with respect to the specific domain of performance that may be contextually relevant. Many are the domains of human performance that allow and often require degrees of expertise beyond the ordinary: athletic, artistic, medical, academic, legal, and so on. Expert perception, agency, and knowledge would be determined proportionally to the respective levels of competence set for the specific domain. This is often set largely by convention, or, for more basic competences and dispositions, by the requirements of success in our evolutionary niche. After all, how we credit, discredit, trust, and distrust has a large bearing on human flourishing, individually and collectively.

d. Manifestation determines credit and discredit and also is attributable causally to the host of the manifest disposition in a way that is projectable, though this is no more amenable to formulation than seems the projectability of greenness (or “green”) by contrast with grueness (or “grue”). When something shows its true colors through manifestation, we can take notice and revise our view of what to expect from the host of the manifest disposition. This is in contrast to when the disposition is only mimicked, so that the correlated trigger prompts the correlated ostensible manifestation but only through the trumping action of the mimic. Such fake manifestation is not to the relevant credit, causal or otherwise, of the disposition. And the host, in turn, acquires no credit or discredit thereby.¹³

¹³The agent does not get full credit for the *success* of his performance, but he may well get full credit for qualities of the performance that make it a skillful performance. Thus an arrow may leave the bow well oriented, in such a way that, absent interfering conditions, it would succeed. And this property of the *shot* may be fully creditable to the archer, even if the *success* of the shot is not.

e. Consider the mimicking of fragility when a fine wine glass shatters because it is zapped upon hitting the hard floor. By hypothesis, the causal action of our zapper trumps the inner structure of the glass, whereby it normally shatters on impact. Still, that inner structure can be causally operative, as it is through the agency of the zapper (who hates the impact on the hard floor of the fragility that he spots in the fragile glass). Despite being causally operative in that way, through the knowledge of the zapper, that inner structure is not causally operative in the right way. And this is why the fragility that we normally attribute to the glass is not really manifest on that occasion.¹⁴

That is why, although the competence-theoretic account goes beyond the proposals of Grice and Davidson, we may need to rely on our shared grasp of how competences are constituted and manifested: for example, of when the relevant causation is not deviant or wayward.

That requires an SSS-joining of seated skill, shape, and situation so as to cause the manifestation upon the onset of the trigger. And this must occur in an appropriate way. Consider, for example, what is required for a true manifestation of fragility as a fine wine glass shatters upon hitting the hard floor. By common consent, such true manifestation excludes the zapping of the glass by someone who hates fragility meeting hardness, even if, through

¹⁴We might, of course, understand a broader, more determinable sort of “fragility” that generalizes from the situations we require for our ordinary fragility. This more determinable fragility would allow for an object to acquire a temporary fragility in the presence of the hateful zapper. *This* sort of fragility the glass might even share temporarily with an iron dumbbell (so long as the zapper hovers and extends his hatred beyond fine-glass-hitting-hardness to iron-meeting-hardness). However, that would be an extension of proper English, and also of proper cross-linguistic ideology, because the same surely goes linguistically for other natural languages. Our discussion in the main text suggests reasons that it might or might not be advisable to so extend our language and ideology. This would likely depend on how likely it is for the relevant community to encounter such zappers. Thus, recall the suggestion in the text: “Such an instrumentally determined common sense must of course be structured against background implicit assumptions about what is normal or standard, either in general or with respect to the specific domain of performance that may be contextually relevant.”

his knowledge and action, he does manage to link causally the fragile structure of the glass with the shattering upon impact.¹⁵

3. *How to Go Beyond Appeal to "The Right Way"*

a. That is to present our account with a certain modesty. More boldly we might claim that when the zapper shatters the glass because he knows it to be fragile, the shattering does not intuitively manifest the glass's fragility. Anyone who joins me in finding that plausible enough can make the following bolder claim:

Manifestation enables us to go beyond the need to appeal to "the right way." The manifestation of competences and other dispositions, then, provides a solution to the problem of specifying "the right way" as it pertains to action, perception, and knowledge.

The problem of specifying the right way is solved through a primitive relation of manifestation that has outcome manifestations (successful performances) on one side and competences (perceptual, agential, epistemic) on the other.

b. Some may well remain skeptical of the powers alleged for our primitive concept of "manifestation." To such skeptics we can offer,

¹⁵ As we focus on examples such as that of fragility zapped, it emerges that we may still need reference to a showing, a demonstration, that goes beyond any ability to tell. We may require that a disposition be manifest in a certain outcome only if it accounts for the outcome in a certain way that we can distinguish only in practice, through demonstration. This requires a joining of seated skill, shape, and situation so as to cause the manifestation upon the onset of the trigger. And this must take place in a normal way identifiable only demonstratively, not through any verbal formula, a way that by common consent excludes the action of a zapper, even if he does manage to link the fragile structure of the glass with the shattering upon impact.

However surprising it may be to the verbally accomplished, it seems increasingly plausible that much of our conceptual repertoire is not given substance, nor is even so much as adequately described, through linguistic formulation. Our shared conceptual scheme of dispositions, abilities, competences, and their manifestations is plausibly a case in which our agreement lacks explicitly formulable content. Accordingly, we also lack any nontrivial way of *securing* it through explicit convention. All of this is in line with how etiquette is constituted, learned, and invoked.

as a fallback, a more modest option whereby, perhaps through examples, we can still explain what is required for proper manifestation. And we could even disown any ambition to rely exclusively on explicit verbal formulation (as by invoking “manifestation”).

Even on this more modest option, we will have made progress. We will have specified more fully the sort of causation involved. And we will have seen that it is the same sort of causation in all three cases: that of perception, that of action, and that of knowledge.

C. METHODOLOGICAL CONTEXT FOR OUR INQUIRY

We often appeal to what we would ordinarily say, and even to what we would ordinarily think, in the exercise of generally shared concepts. But our main interest in philosophy is not restricted to semantic or conceptual analysis. When we wonder about personal identity, freedom and responsibility, the mind and its states and contents, justice, rightness of action, happiness, and so on, our main focus is not, or not just, the words or the concepts. There are things beyond words and concepts whose nature we wish to understand. The metaphysics of persons goes beyond the semantics of the word “person” and its cognates, and even beyond the correlated conceptual analysis.

The same goes for epistemic concerns such as the nature of knowledge and other epistemic phenomena. Consider the semantics of epistemic vocabulary, and even the conceptual structure of epistemology and its normativity. It seems an open possibility that our words and concepts are not in the best shape, just as they stand, for grasping and understanding the relevant domain of objective phenomena. Why not leave open the possibility of terminological and conceptual improvement in epistemology in a way analogous to what is familiar in science when we reconfigure the terminology and ideology of fish, for example, or of vegetables, fruits, and much else? If so, semantic and conceptual analysis might still remain an excellent start in epistemology. Such analysis would remain important in various ways to the epistemologist, and to the philosopher generally, but we might also be able to

delineate phenomena whose importance is obscured by ordinary speech and thought.

If so, that might also, as a bonus, help shine light on pervasive and persistent disagreements so common in philosophy. Some of us may just be trying too hard for the exact, fully general analysis that will apply smoothly and directly to all thought experiments. For the sake of a simple and illuminating take on some range of phenomena, we may just be right to “bite certain bullets” if by so doing we can distinguish a type of phenomenon that seems plainly important in the domain of interest, such that it should also prove illuminating to consider its relations to other such phenomena. We may then reject an ostensible counterexample while allowing that the example points to some further interesting phenomena interestingly related to the phenomena of more direct and central interest to us in our specific inquiry.

Philosophical progress might then take a form similar to the kind of scientific progress that involves conceptual innovation. We may find in the phenomena themselves differences that seem important even if there are no proprietary terms or concepts that correspond to them neatly and without exception. If so, it may behoove us to stretch close terms or concepts so that they will help us to mark the relevant phenomena and to cut the domain more closely at the joints.

Finally, once our objective is analysis that is metaphysical rather than linguistic or conceptual, the bullet-biting recommended does not amount to giving up on intuitions. Our metaphysical project may then be driven by intuitions concerning the phenomena themselves, and not the (proper use of) the language used to describe them or the content of the related concepts. After all, to bite the bullet, on the proposal floated, is precisely not just to describe or understand but to *change* our given language or concepts, at least by addition but perhaps also by subtraction, or by modification. True, the relevant metaphysical intuitions will need conceptual content, but our focus on the phenomena may lead to concepts that are modified, or even quite new. We need not be restricted to concepts used when we begin our inquiry, as our inquiry may properly lead to revision.

FURTHER READING

Choi, Sungho, and Michael Fara. "Dispositions." In *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition), ed. Edward N. Zalta. Stanford, CA: Metaphysics Research Lab, Center for the Study of Language and Information. Available at <http://plato.stanford.edu/entries/dispositions/>.

Two Forms of Virtue Epistemology

TWO QUITE DISTINCT forms of virtue epistemology are generally recognized. One of these finds in epistemology important correlates of Aristotle's moral virtues. Such responsibilist character epistemology builds its account of epistemic normativity on the subject's responsible manifestation of epistemic character. The other form of virtue epistemology cleaves closer to Aristotelian intellectual virtues while recognizing a broader set of competences still restricted to basic faculties of perception, introspection, and the like. How more specifically are these two forms constituted and related? This chapter goes into that question.

A. EPISTEMIC VIRTUES, CONSTITUTIVE VERSUS AUXILIARY

1. Here is a view of intellectual virtue that is thought to accord with *competence virtue epistemology*:

[Intellectual virtues are] personal qualities that, under certain conditions and with respect to certain propositions, are a reliable means to reaching the truth and avoiding error.¹

Moreover, agential virtues are thought to share certain distinctive features:

- (a) They are virtues exercised in intentional agency.
- (b) They are developed through repeated agency.
- (c) They bear on the personal worth of the possessor.
- (d) They aid agential success.

¹Jason Baehr, *The Inquiring Mind* (Oxford, UK: Oxford University Press, 2011), chapter 4, section 4.2.

- (e) In epistemology, they concern intentionally conducted inquiry.²

Because of its focus on traditional faculties such as perception, memory, and inference, such virtue reliabilism is said to *overlook character traits* such as open-mindedness and intellectual courage. These traits are said to possess the five features of agential virtues listed and to satisfy the formal conditions accepted by competence virtue epistemology (spelled out earlier). Such overlooked character traits are indeed, under certain conditions and with respect to certain propositions, a reliable means to reaching the truth and avoiding error, and their exercise can most saliently explain why the subject gets it right in believing as he does. Does our competence virtue epistemology (virtue reliabilism) plead guilty?

2. Are intellectual virtues to be understood *simply*, in line with the assumptions listed earlier, as traits (a) whose manifestations reliably yield true belief and (b) that *play a salient role* in explaining why one reaches the truth in cases where one does so? That is indeed a plausible account of epistemically relevant belief-yielding sources. But there is a better account, one that differs in a way that matters, or so I will argue in what follows.

3. Competence virtue epistemology (CVE) aims to solve two Platonic problems: the *Theaetetus* problem as to the nature of knowledge and the *Meno* problem as to its distinctive value. In connection with the definitional problem, CVE proposes that knowledge is belief whose correctness manifests the believer's pertinent competence. So the pertinent competence (the pertinent reliabilist intellectual virtue) must be one whose exercise can help *constitute* knowledge. That is what knowledge is said to *be*: belief that is correct, that thus succeeds, through the exercise of competence. That is at least what an important, basic sort of knowledge—credal animal knowledge—amounts to.³ However, it is

² *Ibid.*, *passim*; e.g., section 2.2.1, pp. 22–25.

³ Recall the distinction in sections 1–4 of chapter 4, part B, between credal and subcredal knowledge.

crucial to restrict the “through” appropriately. When the correctness of a belief is due to competence in a way that *constitutes* knowledge, it is not enough that the competence reliably puts one *in a position to know*, in a position where one can exercise one’s knowledge-constitutive competences, those whose exercise *does* constitute knowledge.

It may be thought that a virtue such as open-mindedness or intellectual courage could be such a directly knowledge-constitutive virtue. And it may be concluded that the reliabilist competence-based view neglects responsibilist virtues that it should welcome within its fold since they, too, can be important in explaining how a subject gets it right.⁴ Granted, in certain instances, a responsibilist virtue can provide the salient explanation, especially where the truth must be won through complex and competent effort. Courageous and open-minded pursuit of truth—by a scientist, journalist, or detective—might well enable someone to uncover a truth that escapes all others. Consider a reliabilist virtue epistemology that requires for knowledge only that the correctness of a belief derive somehow, perhaps at a great remove, from the exercise of a certain intellectual virtue that is normally a reliable aid to reaching the truth. Such a form of virtue epistemology *would* be negligent if it ignored, or declared irrelevant, any responsibilist virtues that did help one attain truth, including open-mindedness and intellectual courage.

However, that objection need not be relevant to our form of virtue reliabilism, since the intellectual virtues or competences that matter for us are not simply those whose exercise through inquiry can reliably help one reach the truth. Rather, they are competences whose exercise can help *constitute* knowledge. And a competence whose exercise reliably aids our search for truth—even so as to be the salient explanation of why truth is then attained—might easily be one whose exercise would *not* constitute knowledge. It may just fail to be of the right sort to be thus constitutive.

For example, a scientist may follow a healthy regimen with great discipline, and her good health may help explain why she makes her discoveries, by contrast with her wan, depressed rivals; her regimen may even be the salient explanation.

⁴Baehr does so conclude in *The Inquiring Mind*.

Or it might work the other way around. It might be that someone's obsessive pursuit of truth, even at the cost of malnourishment and depression, puts her in a position to attain truths that are denied to her healthy rivals. Even if such obsession to the point of ill health does reliably lead to truth on certain matters inaccessible otherwise (even *if*, I say), the exercise of such personal qualities (obsessiveness) would hardly *constitute* knowledge. The long hours, the intense concentration, the single-minded avoidance of distractions, might put the inquirer in a situation, or enable her to attain a frame of mind or certain skills, through all of which she can have and exercise the competences more directly relevant to the attainment of knowledge. She might acquire important data through a perilous voyage to distant lands, or through extensive observation of the night sky, none of which she could have done without persistent dedication over many years with enormous care.

4. But the point does not require reference to the heroics of a Darwin or a Brahe. A simple example from everyday life should suffice. Suppose a mysterious box lies closed before us, and we wonder what it contains. How can we find out? We might, of course, just open the lid. In pursuit of this objective we will then exercise certain competences, perhaps even character traits (if the box is locked or the lid stuck), such as persistence and resourcefulness. And perhaps these qualities (in certain contexts and in certain combinations) do lead us reliably to the truth. Nevertheless, the exercise of *such* intellectual virtues need not and normally would not *constitute* knowledge, not even when it does indirectly lead us to the truth.

Contrast what happens when we manage to open the lid and look inside. Now we may immediately know the answer to our question, with a perceptual belief—say, *that there is a necklace in the box*—which manifests certain cognitive competences for gaining visual experience and belief. Perhaps this complex, knowledge-constitutive competence first leads to things seeming perceptually a certain way, and eventually to the belief that things are indeed that way, absent contrary indications. A belief manifesting *such* a competence and, crucially, one whose *correctness* manifests such a competence, *does* constitute knowledge, at a minimum animal knowledge, perhaps even full-fledged knowledge (including a reflective component).

It is such knowledge-constitutive competences that are of main interest to a competence virtue epistemology (CVE) aiming to explain human knowledge. Other epistemically important traits—such as open-mindedness, intellectual courage, persistence, and even single-minded obsessiveness—are certainly of interest to a broader epistemology. They are, of course, worthy of serious study. But they may fall outside the charmed inner circle of traditional epistemology. They may be only “auxiliary” intellectual virtues, by contrast with the “constitutive” intellectual virtues of central interest to virtue reliabilism.

My distinction has on one side intellectual virtues whose manifestation helps to *put you in a position to know* and on the other side intellectual virtues whose manifestation in the *correctness* of a belief thereby *constitutes a bit of knowledge*. In my view, competences that *constitute* (credal) knowledge are generally dispositions to believe correctly, ones that can then be manifest in the correctness of a belief. A competence in general is a disposition to succeed when one aims to attain certain objectives, and a competence to believe correctly is a special case of that.

The crucial point is that a competence whose exercise can aid one's attaining a correct and even an apt belief is not necessarily one that *manifests* in any such attainments. For it *need not* be a competence *to attain* any such things as correctness or aptness, despite being a competence whose exercise *further*s such attainments. The competences of focal interest to CVE are those whose manifestations in such attainments *constitute* knowledge. These are the competences whose manifestations constitute apt belief.

5. We have seen how it is that important virtues of inquiry are auxiliary to the attainment and exercise of knowledge-constitutive competences. As we have also seen, however, an auxiliary epistemic virtue might be an overall personal vice. The example of obsessiveness to the point of ill health is already suggestive.

Irrespective of all that, aesthetics and epistemology still enjoy important autonomy. Artistic and epistemic performances are properly assessable within domains unto themselves, the artistic or epistemic domains relevant to such performances. And such assessment is autonomous from the values pertinent to other do-

mains, *as well as* from overall value or moral standards. When we say that knowledge is better than mere true belief, the alleged superiority is epistemic. This superiority must be understood autonomously so that it does not derive from any moral or personal worth that may attach to the believer's motivation. And this is so *even if* some particular epistemic and artistic accomplishments can withstand conflicting extraneous values.

B. THE PLACE OF CONSCIENTIOUSNESS

1. Conscientiousness is the contrary of the carelessness of negligence or recklessness. Might it be epistemically constitutive and not just auxiliary? Let us consider how.

Take our physical ability to pry open a box so as to determine what it contains. The exercise of that ability seems inessential to our coming to know the contents of the box by seeing what it contains. The lid might have even sprung open on its own compatibly with our enjoying the same perceptual knowledge of the necklace. Our physical prying skill would thus be merely auxiliary to the competence that constitutes our knowledge in that case. The knowledge would still have been constituted by our perceptual competence as we perceptually spotted and identified the necklace and aptly believed it to be there.

So are character traits relevantly similar to our prying skill? Is the exercise of such conscientiousness just as external and auxiliary to knowledge as is our physical prying skill in opening the box?

2. Negligence and recklessness preclude full aptness, whatever may be the domain of performance. Distinctive varieties of that phenomenon are found across domains, whether the domain be one of sport, such as chess or archery, or one of artistic accomplishment or professional performance, such as medicine or the law. The epistemic domain is no exception.

Consider open-mindedness and intellectual courage. Each of these forms of conscientiousness seems an ability and a competence to avoid pertinent negligence and recklessness. One's mind

should be open to an *appropriate* extent, not too little and not too much. One should be willing to take an *appropriate* risk of failure by judging in a certain way on a given question: not too little risk, and not too much. Risk-determining factors are set by the specific domain of performance. The relevant failure to be avoided is the failure of one's pertinent attempt in the given domain.

Thus, one is negligent or reckless *epistemically* in making a certain judgment iff one fails to take properly into account the risk of failure *in the relevant attempt to affirm aptly*, the one that constitutes one's judgment.

However, one can have several and sundry aims in making the affirmation that constitutes one's judgment. Among these may be the aim of doing the morally right thing. And this is an entirely different endeavor subject to its own assessment. It, too, will fall under the AAA + SSS framework and will be subject to the question whether it attains not only aptness but full aptness. Moral questions will then arise concerning the distinctively moral forms of negligence and recklessness and more specific varieties of these, such as that involving the dimension of open-versus close-mindedness and that involving the dimension of courage versus rashness or cowardice. Opposed to such moral negligence or recklessness, then, is a corresponding moral conscientiousness.

3. A performance can thus fall short morally even when it succeeds aptly and even fully aptly in various other domains wherein it also falls, such as sports, the arts, or inquiry of various sorts. Recall, for example, how Gauguin may have succeeded artistically with his Tahitian project while failing miserably in the moral domain.

C. THE PROPER CONTENT AND SCOPE OF RESPONSIBILIST VIRTUE EPISTEMOLOGY

1. Prominent among responsibilists is a high-minded conception of intellectual virtues, which are then held to be character traits that bear on the personal worth of the *person*. Such virtues are then viewed as inherently motivational, as character traits manifest in actions motivated by a virtuous pursuit of the truth. In this

view, a belief that derives (at least in important part) from such a virtue must derive from actions that express the subject's love of truth.

2. This form of character-based responsibilist epistemology is then addressed to the traditional problematic of epistemology, at the core of which is the project of defining knowledge. Indeed, it is this *motivational* component that is then held to explain the distinctive value of knowledge above whatever value might be found in the corresponding merely true belief. According to such responsibilism, knowledge is best understood as belief that gets it right through such responsibilist intellectual virtue. Thus argues Linda Zagzebski.⁵

3. That approach seems subject to simple counterexamples, such as a pang of pain or a strike of lightning out of the blue, which one

⁵*Virtues of the Mind* (Cambridge, UK: Cambridge University Press, 1996). Occasionally, and more recently, Zagzebski has taken the somewhat different view that it is only the knowledge that *does* manifest, such high-ranking virtues, that has relevant distinctive value, even if there is a lower order of knowledge that lacks it (e.g., *On Epistemology* [Belmont, CA: Wadsworth Cengage Learning, 2008]). But this will not help with the *Meno* problem, which cannot be solved through appeal to such worthy belief motivated by the love of truth. What makes knowledge of the right way to Larissa better than mere true belief need not depend on such knowledge's being an achievement that deserves admiration or on its having pragmatic value. This is increasingly clear if we switch the example to one of knowing which is the *shortest* road to Larissa. Of the two obvious roads, the shortest may be just infinitesimally shorter, so that its increment of pragmatic value is negligible. Moreover, one's knowledge may have been attained through the most ordinary testimony, such as by asking a passerby, which would merit little personal credit or admiration. And yet knowing what one believes is in that case still better than merely getting it right by luck. The sense in which it is still better comports with the fact that epistemology is not a department of ethics. Epistemic attainments, like good shots, are not quite generally and inherently valuable in any objective sense. In spite of that, the good ones are still "better" than alternatives. Knowledge is in that way a better attainment than belief that does not succeed or does so just by luck. But this *general* superiority is not a quasi-ethical matter of motivation. It is rather a matter of competence, which is often and importantly enough a matter of intentional agency but can also be just a matter of functional, biological, or psychological teleology.

knowingly discerns with no delay. These one *can't help knowing*, sans deliberation and *unmotivated* by love of truth.⁶

Zagzebski responds:

[My definition] does not rule out easy knowledge by sense perception. A person who believes that she sees an easily identifiable object typically knows that she sees the object, provided that there are no indications in her environment that she should not trust her visual sense or understanding of the concept under which the object falls.⁷

She extends the point to testimony, and presumably would go further.⁸

4. However, we cannot explain the appropriateness of a belief that a room has suddenly gone dark as a matter of *non-negligent agency* if that belief is *not at all* a product of intentional agency, which is the sort of agency important to character epistemology. Surely *motivation* relates to agency, not to passive reactions that approximate or constitute mere reflexes.⁹

Suppose one's locomotive stays on track, despite one's having actively intervened *not at all*. One might still deserve credit even so *if* there have been junctures where as conductor one could have intervened, where one was free to intervene, and, without negligence, freely opted not to do so. Unfortunately, this will not do. The problem is that in the cases urged by the critics, there is no freedom to intervene in what seems clearly to be a belief, and even an instance of knowledge, as when one knows that the room has gone dark.

⁶Baehr, *The Inquiring Mind*.

⁷Zagzebski, *On Epistemology*, chapter 5, p. 128.

⁸Compare Sosa, *Knowledge in Perspective* (Cambridge, UK: Cambridge University Press, 1991), 240: "The beliefs of a *rational* animal hence would seem never to issue from *unaided* introspection, memory, or perception. For reason is always at least a silent partner on the watch for other relevant data, a silent partner whose very *silence* is a contributing cause of the belief outcome."

⁹See Baehr's *The Inquiring Mind*, p. 44: "[If as I work late at night there is a power outage] . . . I am, as it were, *overcome* by knowledge that the lighting in the room has changed. . . . Nor is it plausible to think that I am "trusting my senses" in the relevant, motivational sense. . . . Again, knowledge of this sort seems not to involve or implicate the knower's agency at all."

5. Here is the upshot. Consider responsibilist virtues that *both* are agential *and* bear on the personal worth of the agent (such as motivation by love of truth), in virtue of their motivational component. We cannot build a traditional epistemology on such virtues. Not even knowledge can be accounted for in those terms. However, in my view a traditional epistemology *can* be built on responsibilist virtues. Where both of the foregoing responsibilist approaches go wrong is in supposing that responsibilist virtues must involve the personal worth of the agent, must be virtues of *that sort*, involving motivation that passes muster.

Moreover, my point here should not be dismissed as merely terminological. Understood in a metaphysically interesting way, my claim is that the relevant kinds of virtues for building a responsibilist virtue epistemology are not just the following two: (a) nonagential faculties and (b) personal-worth-involving, motivationally appropriate agential virtuous competences. We may or may not consider the latter to be a category or kind worth emphasizing. We may or may not consider it worth emphasizing in a responsibilism that aspires to solve traditional epistemological problems. Regardless of all that, there is, at a minimum, *also or instead* the following epistemic kind: (c) agential virtues. *These obviously go beyond nonagential faculties. So they go beyond a reliabilism restricted to such faculties.* And so I submit that they can reasonably be considered “responsibilist” intellectual virtues in the sense that agents would be epistemically, agentially responsible in exercising them, irresponsible through their neglect, and even vicious through exercise of conflicting dispositions. In other words, they are traits or competences of agents as agents. And among these are the traits or competences of conscious, intentional agents as such.

D. VIRTUE EPISTEMOLOGY: RESPONSIBILISM AS A KIND OF RELIABILISM

1. In order to circumvent the impasse within responsibilism, we must first be clear that epistemology is not a department of ethics. An extremely high epistemic status, certain knowledge, can be attained with a deplorable state that represents a sad waste of time,

as when someone spends a morning determining with certitude how many beans are left in his coffee bag.

Moreover, that is quite compatible with there being special instances of knowledge that *are* outstanding accomplishments, which require an admirable love of truth (on a certain matter) and willingness to pursue it with persistent toil and sacrifice. And it is also compatible with the fact that *possessing knowledge of a certain sort*, for various sorts of knowledge, is an indispensable part of any flourishing life. Moreover, having sufficient knowledge of a certain sort may be indispensable without *any* particular bits of knowledge of that sort being indispensable, or even much desirable.

Independently of all that, it remains that there is a distinctive dimension of epistemic assessment isolated from all such broadly ethical (or prudential) concerns. Moreover, within this epistemic dimension, love of truth plays a negligible role *at most*, if any at all. Hedge fund managers, waste disposal engineers, dentists, and their receptionists can all attain much knowledge in the course of an ordinary workday despite the fact that they seek the truths relevant to their work only for their instrumental value. That is why they want them, not because they *love* truth. That seems, indeed, to be true of service professionals generally, including medical doctors and lawyers. It is not love of truth that routinely drives their professional activities, by contrast with desire for professional standing, wanting to help someone, or trying to make a living.

Disinterested, high-minded motivation must be distinguished from intentional, volitional agency. Dispositions to succeed when one tries need not be closely allied with, and much less do they need to be constituted by, a high-minded motivation, one that bears on the personal worth of the agent, on how fine a person she is. Professionals *are* indeed routinely engaged in intentional, volitional truth-seeking in their work lives, even when they do not disinterestedly, lovingly seek the truth. An assassin may even have *no desire whatever* for the truth about the location of his victim *except only* for the fact that it will make his crime possible. Indeed, if he thought a false belief would at that juncture get him more efficiently to his objective, he might heartily approve of his

so believing, and be glad he did so, with no regrets whatever. His search for truth, because agential, is subject to the full range of responsibilist assessment nonetheless. And his knowing the location of the victim, in believing as he does about that location, is still better *epistemically* than his *merely* believing correctly, and of course better epistemically than his believing incorrectly. Similarly, his shot may be an excellently apt shot, and thereby better than an inapt shot (whether successful or not), despite the murderous motivating intention. (That is to say, it is better *as a shot*; it is *a better shot*. It need not be a better event, or a better thing to happen, nonetheless.)

In conclusion, once we distinguish the *sort* of comparative evaluation (epistemic *performance* evaluation) that is involved in our taking knowledge to surpass merely true belief in (the relevant sort of) value, this removes any temptation to take personally laudable motivation to be the key, even if *in a broad sense* one's cognitive prowess may be a component of one's accomplishment worth, as might be the shooting prowess of our assassin. *Broad* "accomplishment worth" is not what responsibilist, character epistemologists have in mind.¹⁰ The assassin is not a better person for being such a good shot. A more accomplished person, yes, but not a better person, in any sense closely related to ethical assessment.

Let us turn now to a second distinction that will help accommodate responsibilism properly in epistemology.

2. At a certain level of abstraction, we can distinguish two sorts of "belief," one implicit and functional, the other explicitly, consciously intentional. It is the latter that needs to be explored in giving *responsibilism* its proper place in epistemology. This is because it is consciously reasoned choice and judgment that most fully manifest our rational nature. And it is with such choice and judgment that we reach the level of human functioning that most fully manifests our philosophical attainment. Accordingly, it is such consciously, rationally endorsed judgment that is at the focus

¹⁰ At least not Baehr, given his account of responsibilist virtues, which does seem representative of responsibilist views generally.

of the epistemological tradition from the Pyrrhonists through Descartes. It is not only the *act* of conscious, intentional judgment that is at the core, however, since by extension there is also the correlated *disposition* to judge upon consideration.

Still, although we do not here focus on functional, implicit belief, what we learn about conscious, intentional belief should carry over to belief generally, whether intentional or functional. The key to the carryover would be a conception of functional belief as still aimed at truth or at representing accurately and reliably enough. Functional belief might aim at truth functionally: for example, through psychological or biological teleology. This would enable thinking of functional belief also as a kind of action, even when it is only implicit, and neither *conscious* nor *intentional*. Anyhow, I distinguish such functional belief only to put it aside so as to focus on the sort of belief that does turn out to be a form of intentional action.

What is intentional belief? How is it structured? We focus on affirmation, and the corresponding disposition to affirm, in the endeavor to answer a given question correctly.¹¹ Consider the great importance of these for a collaborative social species. They seem essentially required for collaborative deliberation and for information sharing. Take *collaborative deliberation*, right up to the most complex, as in a nation's governance, and also *information sharing*, crucial as it is in a great many contexts, prominently in scientific inquiry.

Such affirmation is largely conscious and intentional. If you add a column of figures in your head, for example, you may seemingly obtain a certain result. But if the problem is complex enough, you may still hesitate to affirm accordingly. You may first take out pencil and paper or a calculator. Eventually, coincidence of results may provide strong enough evidence, which leads you to assent (perhaps properly so). You *decide* when to assent; you wait until the evidence is strong enough.

We focus on such intentional, judgmental belief. How is it structured? Judgmental belief is definable as a certain sort of dis-

¹¹In what follows, nearly always "affirmation" is short for "alethic affirmation" (where what makes it "alethic" is that it is aimed at truth).

position to affirm. What sort of disposition? For a start, let us take *judgment* that *p* to be a certain sort of affirmation *in the endeavor to get it right on whether p*. *judgmental belief* can then be understood as a certain sort of *disposition to judge* in the endeavor to get it right on whether *p*, if one so endeavors.

Compare pragmatic affirmation, whether as a means *to reduce cognitive or affective dissonance* or *to instill confidence that will enhance performance* or the like. In our conception, the latter is not proper belief. It is rather a sort of “make-belief” or mock belief.¹²

3. What distinguishes real belief from make-belief? The difference involves the subject’s intentions. In make-belief one affirms just in pursuit of some nonepistemic, practical aim. By contrast, in judgment and judgmental belief one constitutively aims at *getting it right* on the question addressed. Could that be all there is to the difference?

Before us so far is a partial account of judgment as *a certain sort of affirmation in the endeavor to get it right on whether p*.¹³ Judgmental belief could then be understood as a corresponding disposition: to judge in the endeavor to get it right *thus* on whether *p*, if one so endeavors.

4. *Suppose* that reflective knowledge is knowledge properly so-called and that the highest level of reflective knowledge, of distinctive interest to the philosopher, is the knowledge that gains

¹²The act of affirming can take either the form of public assertion or that of private affirmation to oneself. This is an act of crucial importance for a social species that depends as heavily as we do on collective deliberation and on the sharing of information. But it is worth distinguishing various importantly different objectives that one might have in performing that act. For epistemology, there is a particularly important intention that one might have in performing it, namely that of getting it right *thereby* on the relevant whether question: the act of *alethic* affirmation. And we do well to recognize as deserving special attention a further particular act: the act of affirming in the endeavor *thereby* to get it right *reliably enough and indeed aptly*—the act of *judgment*. Closely related to that is, of course, the corresponding disposition, which one might then label “judgmental belief”

¹³Judgment is, in our first approximation, “a certain sort” of (alethic) affirmation, of affirmation aimed at correctness. We have seen how one can thus affirm (alethically, in pursuit of truth) without judging, as does the quiz-show contestant.

conscious, agential, judgmental endorsement. This, then, is *more plausibly* the knowledge at the center of the epistemological tradition from the Pyrrhonists to Descartes and beyond. It is not just the knowledge that is acquired implicitly, with normal automatic processing, in the course of an ordinary day. Rather, it is the knowledge that does, or at least can, stand up to conscious reflective scrutiny, no holds barred. The knowledge that falls short includes not only the merely implicit belief acquired automatically. Also lacking and inferior is the explicit and conscious judgment that reflects what is absorbed uncritically through the culture's hidden persuaders. These judgments can be willingly rendered, explicitly and consciously, while still falling short because unendorsed *and not properly endorsable* by that subject, who lacks the rational wherewithal even dispositionally.

E. RESPONSIBILIST VIRTUE THEORY IN THE TRADITION

1. The Pyrrhonists stop short of *endorsed beliefs*. If we define *reflective knowledge* as animal knowledge properly endorsed, then the Pyrrhonists fall short of that and settle into a state of judgmental suspension, even while continuing inquiry. Moreover, they allow their *functional* seemings to have sway over their lives: they opt to live by appearances. Resultant functional seemings result from competing vectorial seemings (with either positive or negative direction, along with some degree of confidence representable on the unit interval). For example, one may thus resolve conflicting testimony from two friends in favor of the one trusted more. The Pyrrhonists seem to guide their quotidian conduct on the basis of such resultant seemings. That is the regimen they explicitly advocate. But they will not *judgmentally* endorse any such seemings. They are impressed by the fact that resultant vectors rarely or never warrant such endorsement. So they prefer to suspend conscious endorsement. This is why they remain skeptics (in a state of continuing *skepsis*, or inquiry, with little or no settled judgment or judgmental belief).

2. The endorsement of Pyrrhonian interest is agential. Crucial agential performance and competence thus attend even the most basic perceptual knowledge. Functional, perceptual seemings are passive states that we cannot help entering. But *endorsement of them* remains volitional, agential.¹⁴ And such endorsement is *required* for those functional states to ascend to the level of fully reflective, judgmental knowledge, the level to which the Pyrrhonists aspire, in which they are followed by Descartes.

3. Conscious epistemic agency can be found *not only* in the second-order endorsing judgments, moreover, but *also* in corresponding first-order judgments themselves: not in the *merely functional* introspective or perceptual beliefs but in correlated judgments and judgmental beliefs. This point can be developed through the following items and distinctions:

- (a) Affirmation, pure and simple, as a means to whatever objective, if any. (This is normally a free act, whether it is affirmation to oneself or to others.)
- (b) Affirmation in the endeavor to answer a whether question correctly (that is, *alethic* affirmation).
- (c) Affirmation in the endeavor to answer correctly *and also competently, reliably enough, even aptly*.

Consider *such* affirmation to oneself, or to others when this is called for. This is an act of crucial importance to a species as reliant as we are on collaboration and coordination, intellectual and deliberative. And it is thus an act that one would expect to be subject to social norms.

¹⁴I see no sufficient reason to accuse the Pyrrhonists or Descartes of lying or self-deception when they claim to withhold endorsement; not when we can understand such endorsement as a matter of the willful act of affirmation. Sceptics generally through the ages have often refused to flat-out affirm things that the rest of us do affirm flat-out. At least they refuse to do so publicly. But once we recognize a private correlate, that of affirming to oneself, I see no reason to deny that one can refuse to affirm even privately, as a skeptic, especially when our quotidian lives can go on with limited disruption, because they can proceed on the basis of credences understood more functionally in terms of degrees of confidence.

4. Consider, more specifically, item 3b. This particular form of affirmation is compatible with guessing: quiz-show contestants do affirm, even to themselves, and do so in the endeavor to answer correctly, because only thus will they win the prize. Only with item 3c do we have *judgment*. And judgmental belief is the disposition to judge when one faces a question *honestly*, with intellectual honesty, which does *not* mean *purely disinterestedly*, and this for more than one reason. Thus, one may *also* be looking for financial reward or professional recognition, or even to reduce cognitive dissonance.

5. For the Pyrrhonists, proper endorsement requires the ability to answer the skeptic satisfactorily. So there is a kind of knowledge of special value (reflective knowledge) that does require agency even when we most simply and passively take the given. Even when one (Müller-Lyer) line seems to us (passively) longer than the other, the question of endorsement is in place and can be pursued with inquiry (as it is pursued by the Pyrrhonian skeptics). And the endorsement is also paired with the questions consciously addressed even on the first order, where again intentional agency is required. Ironically, reliabilist competence epistemology is a more radical responsibilist, character epistemology.

FURTHER READING

- Baehr, Jason S. *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. New York: Oxford University Press, 2011.
- Battaly, Heather. "Virtue Epistemology." *Philosophy Compass* 3, no. 4 (2008): 639–63.
- Greco, John, and John Turri. "Virtue Epistemology." In *The Stanford Encyclopedia of Philosophy* (Fall 2015 edition), ed. Edward N. Zalta. Stanford, CA: Metaphysics Research Lab, Center for the Study of Language and Information. Available at <http://plato.stanford.edu/entries/epistemology-virtue/>.
- Zagzebski, Linda Trinkaus. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. New York: Cambridge University Press, 1996.

Knowledge, Time, and Negligence

HOW IS KNOWLEDGE affected by negligence and the passage of time? We shall consider some surprising ways.

A. SAUL KRIPKE'S DOGMATISM PARADOX

1. As night falls on your cabin in the woods, you do well to shut the door. Who knows what creature might come through that doorway to do you harm? Especially should you do so if you know that creatures *would* enter, bringing nothing that matters except possibly harm. On that basis, you should form the intention to shut that door, and you should act on your intention.

Consider now an analogy in epistemology. Once you know that *p*, you can deduce, with extreme competence and simplicity, that any evidence contrary to *p* would be misleading, whereas positive evidence would probably do you little good. After all, by hypothesis you already know that *p*! Given this, you should close your mind to any new potential evidence on the question whether *p*. If positive, the evidence will do little for you; if negative, it will harmfully pull you away from the truth and may even cost you the knowledge that you have.

This is a powerful argument for turning dogmatically close-minded once you know that *p*. And we do often act on similar reasoning. We may ignore the *National Enquirer* as we stand in line at a supermarket counter. Reading that tabloid will do us no good, we reason, so why spend the time? Here I assume what is of interest to be exclusively epistemological. One aims only to acquire interesting or important information. If one wants a good laugh, that's another matter.

In ignoring the tabloid we properly disregard a *source*. Alternatively, we may focus not on the source but on the question. Thus,

we take ourselves to know that astrology is ridiculously false and refuse to waste our time listening to any advocate. Any defense of it is not worth hearing out. Most of us would endorse that strategy to protect our time and energy.

For another example, you add up a restaurant bill and figure out what each diner owes, which gives you reason to turn a deaf ear to your fellow diners, refusing to double-check. Indeed, having attained knowledge, one should never try to make sure, since that would only risk being misled by counter-evidence.

The replication of experiments after one knows also comes under a cloud. Even when we barely know, further replication might just mislead. It should not be done, some will say; indeed, the scientific community should institute rules against such replication.

2. Still, we do not adopt any general policy of dogmatically closing our minds on the many questions to which we know the answers. We do not try to block incoming evidence that might have bearing. And we would seem deplorably close-minded if we did try.

Even though universal dogmatism is thus unacceptable, one still has *prima facie* reason to close one's mind, as explained earlier. If the reason is only *prima facie*, however, one may still prefer to continue inquiry. One may want to make *sure*, just for the fuller satisfaction of curiosity, or because it is epistemically better to know with certainty. One might thus have reason to double-check and to seek further evidence, reason *stronger* than the *prima facie* reason to close one's mind.

We earlier rejected dogmatic mind-closure because we understood it as final. Even if such peremptory closure is generally unacceptable, however, knowledge may still plausibly give us *prima facie* reason to close our minds.

What distinguishes cases in which such *prima facie* reason graduates into *ultima facie* reason? When are we right, *all things considered*, to close our minds?

3. Kripke's puzzle concerns a *prospective intention* to close one's mind.¹ A second puzzle is closely related. Instead of the prospec-

¹Saul A. Kripke, *Philosophical Troubles: Collected Papers*, volume I (Oxford, UK: Oxford University Press, 2011).

tive intention, it concerns an *antecedent belief*: namely, that any further evidence will be misleading if negative. Once you attain knowledge, you virtually know that any contrary evidence will be misleading and is best ignored. When negative evidence does come forth at some later time, *should you then* ignore it in accordance with what you already know? *This* problem is resolved by noting that knowledge can be lost with a change of evidence. As Gilbert Harman notes, later total evidence may deny you the knowledge that until then you enjoyed.²

The two problems of dogmatism are closely related, since the Harman belief might sustain the Kripke intention. If at a given time you know that *p*, you might form and implement an intention to close inquiry, regardless of how much is at stake. Even with the belief-related problem resolved as it is by Harman in the way just noted, however, the intention-related problem might still puzzle us. Why is it not appropriate to base on our knowledge an intention to close our minds and thus avoid epistemic harm?

4. Consider the choice whether to affirm, deny, or suspend judgment on a certain question. Even if we put aside which option will bring the greatest *overall* intellectual gain, one might still wonder how best to respond to the question with regard to *knowledgeably answering that question specifically*.

Our threefold option thus concerns a given question, whether *p*. One might affirm that *p*, deny it, or suspend judgment. Which is *epistemically* the best choice of epistemic attitude to *that* question, to the question whether *p*?

An example clarifies this distinction. One might act in a way that best promotes acquisition of knowledge by trusting a certain book. That book might still fail to deserve our trust, however; the best of reasons might strongly counsel that we distrust it. Believing what that book says is then no way to attain knowledge *in* so believing, despite how much truth and knowledge such trust might yield.

Let us linger over *this* status of justification or competence. I mean the normative status that a belief must *itself* have in order to constitute knowledge. Suppose we close our minds on a certain

²G. Harman, *Thought* (Princeton, NJ: Princeton University Press, 1973).

belief. What effect might this have on the epistemic standing of that belief? Can we close our minds in such a way that we *thereby* lose our knowledge? Might we lose it through a kind of diachronic epistemic negligence? Even if our belief was earlier competent, its status might change as time goes by and our minds remain closed. Sooner or later the belief might cease to qualify as competently held, because it now remains in place only through negligence.

Surely that can happen. Consider the overall epistemic competence that can holistically yield knowledge. If we are to give negligence its proper weight, we must include in such overall epistemic competence more than just how probably true our belief is if based on the total evidence. Overall epistemic competence must also include the proper conduct of inquiry. *Negligent* belief would not count as fully competent and would not constitute true knowledge.

5. Suppose we grant its proper status to conscientiousness (to the proper avoidance of negligence). What determines the epistemic justification (competence) of a belief seems, then, to surpass the reliability of the process that produces that belief. If so, the following questions will need attention: First, is epistemic negligence a causal factor in the formation of a belief? If not, how can it still bear on the relevant reliability of that process? Second, suppose that a believer is guilty of epistemic negligence. This would be for him to *fail* to do certain things. Can such a failure have positive causal bearing on the belief formed so as to affect the reliability of that process?³

B. ANIMAL COMPETENCE AND REFLECTIVE JUSTIFICATION: THE IMPORTANCE OF SYNCHRONISM

1. Plenty of beliefs are, of course, clearly based on conscious reasoning from given evidence and stored through retentive memory. Take initial uptake of something “given.” Suppose it to be per-

³Process reliabilism is discussed more fully in my “Process Reliabilism and Virtue Epistemology,” in *Alvin Goldman and His Critics*, ed. Hilary Kornblith and Brian McLaughlin (Chichester, UK: Wiley-Blackwell, forthcoming).

fectly competent and apt, through introspection or perception. Suppose further that the resulting belief is competent and aptly stored. This is all compatible with the believer's forgetting how he initially acquired that belief. One might later be able to say little more than "I just remember." What, then, is the later epistemic standing of that belief? About this we now face a puzzle.

We are focused on a time late in the life of the belief. No one can now detail how well that belief was first acquired or retained over time. Concerning its specific subject matter, moreover, the evidence now available is deplorably slight. If the belief must now rest on the evidence presently available, it may no longer count as epistemically justified. How competently can the believer now retain his belief? The belief's present epistemic standing must be assessed in the light of how good the subject's memory is for that sort of belief. Unfortunately, the competence involved in such second-order assessment might provide far less by way of reliability or justification than the competence that yields the belief on the first order. Such first-order competence will often combine excellent initial acquisition with excellent later retention, excellent perception with excellent memory.

What, then, is the believer to do as time passes? Should confidence dwindle in tandem with reduced qualification to endorse first-order belief? Consider the steady decay of the information required for endorsement. Compatibly with such second-order weakening, retentive memory can remain strong indeed, so as to ensure that the retained belief is very probably true, given the excellent perceptual competence that produced it and the excellent retentive memory that has kept it securely stored. That belief may thus constitute first-order animal knowledge of the highest quality. In many cases what decays over time is only the reflective second-order perspective.

Here is an example. At noon on a certain date, you are mistreated as a child. You know extremely well that it is noon on that date. You store that belief for months, even years. Maybe you retain it through excellent memory. In general, people would not remember so well. Maybe, in general, your own memory does not work so well. But it does in this case, on this sort of subject matter. That event stands out in your mind, and your memory of it is outstanding.

The perception-plus-memory manifest in your continuing belief has the highest quality. When perception and memory work as they do in this case, they are *extremely* reliable. Compatibly with that, your second-order competence can decay. Just based on common sense, you may come to doubt that your memory of that event is as good as in fact it is. You may even learn of accumulating evidence that memory of such mistreatment is much less reliable than common sense had supposed. Human beings, in general, do not recollect as reliably as had been thought, not even on such important matters, perhaps especially not on *such* subject matter. By hypothesis, however, *your* memory *is* in this case extremely reliable.

2. That is one example of the phenomenon I wish to highlight. But we need not suppose our protagonist to have abnormal powers. Another sort of example would involve just normal human perception and memory. In combination these might lead reliably to a correct present belief. Nevertheless, the believer may now be unable to detail how he acquired and retained his belief through the operation of his excellent faculties. Indeed, the time may come when the believer *still* retains his top-quality perception-cum-memory source for the present belief, although now he has the gravest *positive doubts* on that quality.

Moreover, the puzzle does not arise *merely* from a clash between externalist reliabilism and internalist evidentialism. What is crucially important is rather the clash between two distinguishable epistemic statuses that a belief can have:

First, there is the status that a present belief gains *diachronically* through the subject's thinking and memory over time, no matter how internal such thinking and memory may be over time and no matter how internal the subject's initial data may be.

Second, there is the status that a present belief gains at a given moment *synchronically* through the support of reasons present to the thinker's mind at that very moment.

So the problem transcends the divides between, first, externalism and internalism concerning epistemic justification and, second, between evidentialism and reliabilism.

3. A further question must be faced. Suppose we have an instance of such disparity between animal quality and reflective quality. This is a divergence between the following:

First, the high status a belief derives diachronically from a retention-involving first-order competence.

Second, the much lower status that the same belief might have synchronically at the given time due to the diminished epistemic quality of the believer's relevant second-order competence, either because this competence is much less reliable or because, in any case, it provides less by way of epistemic justification.

Later in life you may doubt your competence to remember having been mistreated as a child, despite remembering perfectly well. For another example, consider arithmetical calculation. You may doubt your ability to perform a complex addition without flaw despite performing it flawlessly.

Suppose that your reflective capacity to endorse a given first-order belief is much diminished. Your *endorsement from the second order* depends for its epistemic quality on the second-order competence exercised. What about your *judgment on the first order*? Here is the question on the first order: What should you affirm within the privacy of your own thought, and what should you assert to others?

4. Moreover, it is not just the *assessment* of a first-order judgment that depends on whether we invoke the first-order animal competence or, alternatively, the second-order reflective competence, especially its conscious variant. There is also this question: *Which perspective should have priority in your own continued believing and judging?* Should you trust your excellent first-order competence, or should you trump that competence once your belief is under present scrutiny, with the inevitable bearing of the second-order considerations? Should you now decide whether to trump based on all the reasons *presently* available to you for conscious consideration?

Suppose we give priority to the reasons presently available. This, in effect, recognizes the importance of a kind of reflective

knowledge involving a second-order judgment (or disposition to so judge). Such judgment itself depends for *its* standing on the quality of the competence that *it* manifests. Reflective knowledge will thus have a particularly important role. It will enable your conscious reasoning, practical and theoretical, and will sustain your place in the community as a bearer of information transmissible through testimony. Such reflective knowledge is constituted by an act of judgment (or a disposition to judge), indeed by two: one on the second order and one on the first order. And these are the judgments that figure in our conscious reasoning, and also in our sincere assertions when we inform others.

Although I have focused attention on second-order decay, the more general issue involves just *rational* decay, on whatever order, so that at the later time the subject has no available rational support for the truth of his belief (beyond the fact that he holds that belief) or may even have a balance of rational support against its truth.

We have been considering this question: *What should determine one's judgment and its quality at the present moment?* Is it *diachronic* competence, even if its initial inputs have long receded from view, or is it rather the reasons available and operative *in the present time slice*? If we opt for the present time slice, this upgrades the second-order perspective, simply because synchronic reasons for *stored* beliefs are now so frequently to be found within that perspective, often mainly so.

5. How properly one proceeds epistemically in endorsing (or not) a first-order belief is determined by the reasons consciously accessible at the time of reflection. It would seem deplorably stubborn to sustain and endorse a belief in the teeth of total available evidence strongly against it. And this is so *even if that belief is correct through diachronic first-order competence that is supremely reliable*.⁴

⁴In keeping with our irenic humor, we can make room for competences that are justification-reliable, as well as those that are truth-reliable. Our argument can thus be as relevant to the diachronic internalist evidentialist as to the diachronic externalist reliabilist.

A belief that is thus apt through diachronic competence can still fall short. It falls short if it turns out *not* to be endorsable synchronically through the balance of available reasons. And this failure must bear on how the subject proceeds on the first order. He cannot properly *judge* by relying blindly on his stored belief and on the diachronic competence that sustains it (blindly as of the present time slice, which is the time of decision). As our subject consciously considers his options, he must now prefer suspending to affirming. He may even need to contradict in synchronic judgment his diachronically sustained belief. Only rationally competent synchronic endorsement can now lift one's belief to the higher level of *reflective* knowledge. Once a belief is under scrutiny, only such reflective knowledge qualifies as a proper basis for conscious reasoning, practical or theoretical. And, finally, only such knowledge then sustains proper assertion, enabling the subject to fulfill his communal role as purveyor of information.

C. KNOWLEDGE, JUSTIFICATION, AND THE PROBLEMS OF DOGMATISM

1. If you already know the answer to a given question, how weighty is the *prima facie* reason that this gives you to close your mind?

Section A considered the importance of avoiding negligence for retaining your knowledge. In order to keep one's knowledge, often one must not ignore available evidence. Suppose someone you trust tells you that he will deposit what he owes you in your bank account by noon. Suppose, too, that his testimony and reliability are so trustworthy that in the afternoon you know the money is in your account. Might you now write an important check on the basis of that knowledge? Might you save some time by not verifying online? That depends. Suppose you next acquire evidence that your friend is not so trustworthy. Someone tells you he has early Alzheimer's, and someone else reports that he has once let her down. Suppose also that these reports do *not* remove your extremely well-based knowledge immediately upon receipt. What if you do write the important check while still not

bothering to verify? Would it be right to think that even now, later in the afternoon, when you write the check, you do so in the knowledge that the money is in your account?

2. I am *not* relying here on a contextualist point that our subject would speak incorrectly in *saying* to his protesting spouse, “But I knew the money was there,” that he would speak incorrectly because the spouse’s protest has raised the standards for the truth of “S knows” by highlighting how important it is that the check not bounce. Nor am I supposing that knowledge is pragmatically subject-sensitive. I am not supposing that, given how much turns on whether he is right, the subject no longer really knows if he fails to confirm. Nor am I supposing, finally, that the subject’s merely *knowing* that the money is in the account is not enough to justify his writing that check, that the high stakes require him to *make sure*, or to make *more* sure.

Of course, any or all of those factors might be operative, and advocates of the corresponding views of knowledge and of how knowledge bears on action might invoke their favorite views in accounting for the incorrectness of the subject’s words to his spouse. But someone *unpersuaded* by any of those views might *still* reject the subject’s claim to know that the money is in the account.

3. Let us put aside here all of those factors and the corresponding epistemological views. There is a further reason to think that the subject speaks incorrectly, namely that his negligence makes it *false* that he knows. One might accuse him of a kind of negligence in believing that the money is in the account and in voicing that belief sincerely to his spouse. In the circumstances, he *should* confirm online before sustaining his belief and acting on it. He should confirm, *of course*, given how easily he could do so and how much is at stake. But he “should” confirm *in any case*, even if he were not writing an important check. Arguably, he is not “epistemically justified” in continuing to believe absent confirmation. How would one argue for this?

The argument is *not* just that the subject is *obliged*, all things considered, to stop his normal life and go online. Nor need one

argue that he is *required* to confirm, for the proper conduct of his *intellectual* life. We might focus rather on the normative status that is constitutively involved in knowledge: that is to say, on the belief's "epistemic justification." In order for the belief to remain thus justified, we might contend, our subject needs to confirm. He should confirm in order just simply to keep on believing *knowledgeably*. He is hence *negligent* in not doing so, not just practically but also epistemically.

4. We thus recognize a distinctive reason why at that later time our subject "*should*" verify whether the money is in the account. This distinctive "*should*" is not a matter of what he should do, *absolutely all things considered*, nor is it even a matter of what he should do, *all epistemic things considered*. Rather it is a matter of what he should do if he is to retain his knowledge that the money is in the account, and even his epistemic justification for so believing.⁵ Compatibly with this, of course, he may have more important things to do than just to retain his knowledge, and he may even have more *epistemically* or *intellectually* important things to do. True, he *is* obliged to confirm if he is to remain sufficiently epistemically justified so as to continue to know. But from this it follows neither that he is obliged to verify, *all things considered*, nor even that he is obliged to verify, *all epistemic* or *intellectual* things considered.

5. A case about a pilot may help to clarify the intended point. The pilot may land his plane at the intended destination after difficult flying through very bad weather. Is the success of that flight creditable to the pilot's competence? His piloting competence was

⁵In "The Analysis of 'Knowledge That P,'" *Analysis* 25 (1964): 1–8, reprinted in part as chapter 1 of *Knowledge in Perspective* (Cambridge, UK: Cambridge University Press, 1991), I suggest that the analysis offered brings to light a "striking fact . . . besides justification for 'believing that . . . ' . . . justification for 'not believing that . . . ' is also involved in the concept of knowledge." This is then supported by examples of knowledge-precluding epistemic negligence. Intuitive confidence that negligence matters long preceded my advocacy of virtue in preference to process reliabilism. This intuition is not just theory driven. Theory in this case fits antecedent intuition.

certainly manifest in how he maneuvered so as to reach the destination and land safely. In some ways, then, his competence was manifest. He did have the proper skills seated in his brain and nervous system. He did have the constitutional competence required of a good pilot. He also was in good shape as he piloted that plane to safety. He was fully alert and sober, for example, and he was in good condition in every other relevant respect. Finally, he was also well situated; among other things, the plane was under his operational control, was in good shape, had enough fuel, and so on. And the storm was not unmanageable. It had to be managed, but that was not beyond the capability of a good pilot.

Given all of that, did the flight's success manifest the competence of the pilot? Well, it did manifest a relevant competence—of constitution, inner condition, and situation—that would yield success for the pilot's attempts. The things he did as he flew that plane and landed it safely derived from his complete first-order competence, which included his constitutional skill, his good condition, and his manageable situation. The success of the flight did manifest that threefold complete competence on the first order. Is the success of the flight not, therefore, fully creditable to the pilot? This turns out to be unclear.

Suppose the plane was small, and the pilot responsible for ensuring that it would be in good condition for the flight. In particular, he needed to make sure that the plane had enough fuel. But he neglected to do so. He took off without checking at all. Now is he still creditable? This is implausible. The flight's reaching its destination seems not after all fully creditable to that pilot. The pilot did reach his destination, but his doing so was not a fully apt performance on his part. It was not a performance that manifested all of the relevant competence required for it to be fully creditable. In particular, the pilot's negligence in failing to check the condition of the fuel tank puts a stain on his performance and makes its success substantially a matter of luck rather than competence.

6. Something similar would seem to apply in epistemic domains. Epistemic negligence detracts from epistemic performance, as when we

believe in pursuit of truth. Even if we do happen to succeed on the first order by attaining truth, the success of our performance might still fall short simply because of our negligence. If knowledge is a matter of apt intellectual performance in pursuit of truth, therefore, we get the result that negligence can deny us knowledge, or at least knowledge of a certain epistemically desirable level. We are denied fully apt attainment of truth when we attain truth despite intellectual negligence. We are negligent when we should be open to verifying evidence but close our minds instead. When this happens, our success is lucky. And the luck involved is no less luck when it is good luck, as when the tank happens to be full anyway, or when the further evidence would all be confirming evidence.

7. Finally, our results in section B provide additional reasons for retaining an open mind even when one already knows. These are again reasons independent of contextualism, subject-sensitive invariance, and pragmatic encroachment. Recall the emphasis of section B on your diminished awareness of how a belief is initially formed. Recall the gradual reduction of your ability to reliably endorse beliefs. Recall, finally, what you will need to rely on as your initial acquisition of the belief recedes into the past. Your later rational judgment will need to draw on the evidence available to you *at that later time*. And this will inevitably be new evidence that bears on the truth of the proposition believed. In blocking later evidence, *as soon* as you know that *p*, you will hence preclude or impair your later rational judgment that *p*.

That reveals a cost of closing your mind even when you have a *prima facie* reason for doing so: namely, the reason that you might otherwise be misled by contrary evidence. The cost is that you would now too often, too seriously preclude or impair your ability to judge rationally that *p*. You would negligently, willfully overlook evidence available synchronically that might bear importantly on your judgment. This would entail a massive epistemic loss in perpetuity.

FURTHER READING

- Ginet, Carl. "Knowing Less by Knowing More." *Midwest Studies in Philosophy* 5 (1980): 151–62.
- Harman, Gilbert. *Thought*. Princeton, NJ: Princeton University Press, 1973, especially chapter 9.
- Kripke, Saul. "Two Paradoxes of Knowledge." In his *Philosophical Troubles: Collected Papers*, volume I, 27–51. New York: Oxford University Press, 2011.
- Lasonen-Aarnio, Maria. "The Dogmatism Puzzle." *Australasian Journal of Philosophy* 3 (2014): 417–32.
- Sorensen, Roy. "Dogmatism, Junk Knowledge, and Conditionals." *Philosophical Quarterly* 38 (1988): 433–54.

Virtue Theory against Situationism

SUPPOSE SOMEONE YOUNG—let's call him Tom—gives up his seat on the bus to someone frail and elderly. Why does he? Perhaps only to impress his girlfriend but perhaps rather to be considerate, out of concern for the elder's welfare. And if the latter, maybe it is just a random act, entirely out of character, and due more to his being in a good mood. Alternatively, it might be quite in character for him to act kindly that way. Kindness may be one of his character traits, manifest in that act.¹

Or so one might think without a second thought. Any such scenario would seem a commonplace example of how we constantly try to understand people's conduct. Based on a body of troubling results in social psychology, however, an intriguing critique has been pressed in recent years against such a virtue ethics, raising doubts both about its moral psychology and about its normative content. In what follows we shall review some of the most striking, best-known results, and the arguments based on them.

Similar discoveries have been made by social psychologists about our belief management, moreover, so that a similar critique can be pressed against virtue epistemology. We here come to this

My thanks to Steven Lukes for illuminating formal comments at the Herdecke conference, where an earlier version of this chapter was delivered, and to Jason Kawall for helpful philosophical and bibliographical comments (reflected in notes 17 and 19).

¹Virtue ethics goes back to Aristotle and to ancient Greek philosophy more generally. Long neglected in the shadow of deontological and utilitarian approaches, despite its powerful advocacy by Hume and others, it has in recent decades regained much of its former luster and influence. The approach has both an empirical side and a normative side. It appeals to virtues, or to virtuous traits of character, in both the explanation and the assessment of human action. That an action manifests a virtue bears positively in its overall assessment; that it manifests a vice bears negatively.

topic through our virtue-theoretic approach to epistemology, which is broadly allied to similar approaches to ethics and even to parts of aesthetics. We focus in this chapter on virtue ethics and its correlated virtue psychology, then turn in the next chapter to virtue epistemology. The logical structure of our response to the critique of virtue ethics is closely replicated by a response available to the virtue epistemologist. Accordingly, rather than detail that obvious analogy, the space available in the next chapter will be devoted to an account of epistemic virtues or competences. That account is then defended briefly against its situationist critique by close analogy with how virtue theory prevails in ethics against a similar attack. So we begin with virtue ethics.

A. THE SITUATIONIST CRITIQUE OF VIRTUE ETHICS

1. In the early 1960s, experiments conducted by the psychologist Stanley Milgram at Yale University had disturbing results.² In multiple replications, moreover, the results have held up with impressive consistency. Milgram's subjects believed themselves to be playing the role of "teachers" in a study of the effects of punishment on learning. Here is the scenario. The participants were "teachers" expected to administer electric shocks to "learners" (who were, in fact, Milgram confederates). In one version, the "teacher" (the experimental subject playing that role) saw the learner/confederate strapped down to an electric chair in a separate room. The learner/confederate asked about the shocks and was told by the experimenter in charge that they were not harmful but could be painful. The learner/confederate said that he suffered from a slight heart condition, asked whether the shocks would be dangerous to him, and was assured that they would not be dangerous, although, again, they could be quite painful. When the experiment began, the teacher was given a sample 45-volt shock from the machine, just to add realism. Then he was taken to a position in another

²Stanley Milgram, "Behavioral Study of Obedience," *Journal of Abnormal and Social Psychology* 67 (1963), 371–78. See also his *Obedience to Authority: An Experimental View* (New York: Harper and Row, 1974).

room from which the learner/confederate was no longer visible. As the experiment proceeded, the learner/confederate repeatedly failed to answer the questions asked correctly, so that the teacher was bound to keep increasing the strength of the shocks in 15-volt increments. What were the results?

At 300 volts, the learner/confederate would pound on the wall, scream, and then, at 330 volts, stop responding. Yet most teachers continued to intensify the shocks in 15-volt increments all the way up to 450 volts. That was done by twenty-six of the forty teachers, or about 65 percent. That means that they administered *ten* further voltage boosts *after* the pounding and screaming began! This is what was done by twenty-six of the forty teachers. As for the remaining fourteen, these all went up to at least 300 volts, then stopped at somewhere between 300 volts and 450 volts. If a teacher protested to the experimenter, he got one or more of a standard set of responses, in the following sequence: "Please continue," "The experiment requires that you continue," "It is absolutely essential that you continue," "You have no other choice, you must go on." Again, *all* teachers administered shocks up to the severe 300 volts, and 65 percent went beyond that to shocks of the maximum 450-volt severity.

2. A second, much-cited experiment was conducted at the Princeton Theological Seminary.³ Seminarians were read either something about vocational choices available to them or else the parable of the Good Samaritan. They were then told to go to a certain nearby building, where they were expected to give a talk. Some were asked to hurry because they were already late, others were just told to proceed without delay, and the rest were told that they had a few minutes to spare. On the way to the next building they all came across a figure slumped over in a doorway, groaning and coughing. Whether they stopped to help was determined mainly by how much of a hurry they were in. Subjects' offers to help were positively correlated with their beliefs about how pressed they

³John M. Darley and C. Daniel Batson, "From Jerusalem to Jericho: A Study of Situational and Dispositional Variables in Helping Behavior," *Journal of Personality and Social Psychology* 27 (1973): 100-108.

were for time. Sixteen out of the forty subjects offered to help. Of the eight with some time to spare, five stopped to help. Of the twenty-two who had been told to go right over, ten stopped. Finally, of the ten who had been told they were already late, only one stopped.

True, the percentage of subjects who offered to help had some positive correlation with the passage they had just read, but the correlation was moderate: of those who had just been read the parable, 53 percent offered to help, but of those who had just been read the vocational passage, only 29 percent offered to help.

In conclusion, the experimenters and authors suggest: "Only hurry was a significant predictor of whether one will help or not." For most subjects, punctuality trumped the evident need of someone in distress.

3. Based on the results of these experiments and several others, an attack has been launched on virtue psychology and virtue ethics.⁴ The critics have raised questions of two sorts. First, they have challenged the notion that humans vary significantly in possessing traits of character—some virtuous, some vicious—*important for the explanation and prediction of human action*; second, they have challenged also the normative ideal of human virtue held up by virtue ethics for the reason that humans are so unlikely to guide their conduct by any such ideal.

Leading the situationist charge within philosophy are Gilbert Harman and John Doris, whose views we consider next.⁵

⁴To appreciate the extensive relevant literature, see the masterful book by John Doris, *Lack of Character: Personality and Moral Behaviour* (Cambridge, UK: Cambridge University Press, 2002). The three studies I cite (two here and one in the penultimate footnote to this chapter on virtue theory) are the best known and the ones I have found most striking; the philosophical issues emerge fully, as I see it, on the basis of these three studies, which are also the most cited in the relevant philosophical literature.

⁵John Doris, "Persons Situations, and Virtue Ethics," *Nous* 32 (1998): 504–30; Gilbert Harman, "Moral Philosophy Meets Social Psychology: Virtue Ethics and the Fundamental Attribution Error," *Proceedings of the Aristotelian Society* 99 (1998–99): 315–31. An earlier, softer challenge came from Owen Flanagan in his *Varieties of Moral Personality* (Cambridge, MA: Harvard University Press, 1991). See also Richard E. Nisbett and Lee Ross, *Human Inference: Strategies and Shortcomings of Social Judgment* (Englewood Cliffs, NJ: Prentice-Hall, 1980); Lee Ross

4. According to Harman, empirical testing has found no relevantly different character traits to account for behavioral differences. True, ordinarily it is *supposed* that people do differ relevantly in traits and virtues:

We ordinarily suppose that a person's character traits help to explain at least some things that the person does. The honest person tries to return the wallet because he or she is honest. The person who pockets the contents of the wallet and throws the rest of the wallet away does so because he or she is dishonest.⁶

However, people might behave differently, regularly so, without differing in character traits. The difference in behavior might, of course, derive rather from situational differences. In order to differ in character traits, people must be disposed to act differently though similarly enough situated (or similar enough in how they view their respective situations). As *ordinarily conceived*, moreover, traits are dispositions to issue the trait-relevant conduct across a *broad* range of relevant situations. True honesty, for example, requires honest conduct across a broad enough range of relevant situations. It is not enough that one is honest in forbearing to shoplift although one cheats on tests, on one's income tax returns, and in returning change.

Harman joins Nisbett and Ross, moreover, in distinguishing traits from sustained goals or strategies, and also in finding us too often guilty of the "fundamental attribution error," the error of attributing a trait based on too paltry an evidential basis. According to Nisbett and Ross, "Individuals may behave in consistent ways that distinguish them from their peers not because of their enduring predispositions to be friendly, dependent, aggressive, or the like, but rather because they are pursuing consistent goals using consistent strategies, in the light of consistent ways of interpreting their social world."⁷

and Richard E. Nisbett, *The Person and the Situation: Perspectives of Social Psychology* (New York: McGraw-Hill, 1991).

⁶Harman, "Moral Philosophy Meets Social Psychology," section 2.

⁷Nisbett and Ross, *The Person and the Situation*, 20.

Harman comments as follows on our two striking experiments:

The fundamental attribution error in [the Milgram] . . . case consists in “how readily the observer makes erroneous inferences about the actor’s destructive obedience (or foolish conformity) by taking the behavior at face value and presuming that extreme personal dispositions are at fault.”⁸

Standard interpretations of the Good Samaritan Parable commit the fundamental attribution error of overlooking the situational factors, in this case overlooking how much of a hurry the various agents might be in.⁹

And he sums up his brief:

Summary. We very confidently attribute character traits to other people in order to explain their behavior. But our attributions tend to be wildly incorrect and, in fact, there is no evidence that people differ in character traits. They differ in their situations and in their perceptions of their situations. They differ in their goals, strategies, neuroses, optimism, etc. But character traits do not explain what differences there are.¹⁰

5. Situationism for Doris involves three main commitments. Situationism’s three central theses concern behavioral variation, the nature of traits, and trait organization in personality structure:

- (a) Behavioral variation across a population owes more to situational differences than dispositional differences among persons. Individual dispositional differences are not as strongly behaviorally individuating as we might have supposed: to a surprising extent we are safest predicting, for a particular situation, that a person will behave pretty much as most others would.

⁸Harman, “Moral Philosophy Meets Social Psychology,” section 5.1. Here Harman is agreeing with Nisbett and Ross.

⁹*Ibid.*, section 5.2.

¹⁰*Ibid.*, section 8.

- (b) Empirical evidence problematizes the attribution of robust traits. Whatever behavioral reliability we do observe may be readily short-circuited by situational variation: in a run of trait-relevant situations with diverse features, an individual to whom we have attributed a given trait will often behave inconsistently with regard to the behavior expected on attribution of that trait. . . .
- (c) Personality structure is not typically evaluatively consistent. For a given person, the dispositions operative in one situation may have a very different evaluative status than those manifested in another situation—evaluatively inconsistent dispositions may “cohabit” in a single personality.¹¹

Situationism is not a Skinnerian evisceration of the person. While rejecting cross-situationally robust traits, the situationist admits local, situationally specific traits that distinguish people from one another. These traits are “local” rather than global, or frail rather than “robust”: they do not reliably result in the same trait-relevant conduct across a variety of different situations. Doris continues:

At bottom, the question is whether the behavioral regularity we observe is to be primarily explained by reference to robust dispositional structures or situational regularity. The situationist insists that the striking variability of behavior with situational variation favors the latter hypothesis.¹²

He sums up as follows:

To summarize: According to the first situationist thesis, behavioral variation among individuals often owes more to distinct circumstances than distinct personalities; the difference between the person who behaves honestly and the one who fails to do so, for example, may be more a function of situation than character. Moreover, behavior may vary quite radically when compared with that expected on the postulation of a given trait. We have little assurance that a person to whom we attributed a trait will consistently behave in a

¹¹ Doris, “Persons, Situations, and Virtue Ethics,” 507.

¹² *Ibid.*, 508.

trait-relevant fashion across a run of trait-relevant situations with variable pressures to such behavior: the putatively “honest” person may very well not consistently display honest behavior across a diversity of situations where honesty is appropriate. This is just what we would expect on the second situationist thesis, which rejects notions of robust traits. Finally, as the third thesis suggests, expectations of evaluative consistency are likely to be disappointed. Behavioral evidence suggests that personality is comprised of evaluatively fragmented trait-associations rather than evaluatively integrated ones: e.g., for a given person, a local disposition to honesty will often be found together with local dispositions to dishonesty.¹³

In brief, situationism is committed to the following three claims: *First*, behavioral variation is due more to situational variation than to trait variation. *Second*, traits are frail across situational variation, not robust. *Third*, traits do *not* integrate into coherent characters.

6. Attentive reading reveals that both Harman and Doris reject Skinnerian nihilism on behavioral dispositions. Indeed, both of them *believe that there are traits*, dispositional traits operative in human conduct generally. *What they deny is that these are traits as conceived of by the folk, or by the Aristotelian tradition of virtue ethics.* Harman does consider the possibility of rejecting traits altogether, even those that are local and frail, as opposed to the global and robust. On this more extreme view, human conduct is to be explained not by traits but perhaps through goals, policies, or strategies. But how are we to distinguish *learned* traits generally from such goals, policies, or strategies? The supposed alternative does not clearly differ more than verbally. Suppose I have a firm goal to treat others politely, and I give substance to that goal through my knowledge of what politeness requires in a broad range of situations. How importantly does this differ, if at all, from possessing a trait of politeness (of treating others politely)?

Among people, traits can be rare and distinctive or, alternatively, vulgar and widely shared. They can be local (or narrow or frail), moreover, or else global (or broad or robust). *Our traits, insofar as we have any, are said to be vulgar and local, or at least much more so than is usually supposed.* This is what situationism seems to boil down to, apart from the claim that human personality is normally fragmented and falls far short of the integration proper to Aristotelian practical wisdom.

7. Situationists do agree with the folk, and with the tradition of virtue theory, on one important point:

Variation

There is substantial evaluatively interesting variation in human behavior. People can and do behave variably regarding honesty, kindness, courage, temperance, etc.

This much is in keeping with the experiments and in line with situationist writings. What explains such variation? For the explanation of *cross-personal* variation we must invoke situational differences and distinctive traits; vulgar traits will not help. Nor can *cross-situational* variation be explained by *stably* vulgar traits. If character traits are stably vulgar, then evaluatively relevant behavioral variation must be situationally explained.

Consider now this question: What explains evaluatively interesting human behavior? This is *not* the question broached a paragraph ago. What explains cross-personal or cross-temporal variation in behavior need not be the same as what explains the behaviors severally. Compare this: The *differences* in the rolls of two round marbles cannot be explained (at all) by appeal to their roundness, because they are the same in that respect. Yet either roll might still be explicable (largely, importantly) through the roundness of the rolling object. Similarly, behavioral *variation* may not be explicable by appeal to traits, although still the behavior itself is so explicable, *even if* the traits are vulgar.

So stable vulgarity will spoil a trait for the explanation of behavioral variation, but not for the explanation of behavior itself, whether individual behaviors or behavior patterns.

What about narrowness? How, if at all, does the narrowness of a trait impair its explanatory efficacy? A narrow or local trait is one that yields its evaluatively relevant behavioral outputs in a relatively narrow or local set of circumstances. Dispositions come, of course, in degrees: not only simple dispositions, such as fragility and flexibility, but also those more relevant to ethics, such as honesty and kindness. Accordingly, it is possible to explain the breaking of a vase by appeal both to its impact and to its fragility, even if a fine wine glass is *more* fragile and would have broken not only with that impact but *also* with others that would *not* have affected the vase. Moreover, narrow, local traits of honesty (in returning change, say) may amount to ways in which one can have a low degree of honesty (because one is not also honest in filling out one's income tax forms, in taking tests, etc.). Consider the virtue psychology accepted as common sense by the folk. This is content to postulate varying degrees of its recognized virtues, though we may be surprised to see just how much we nearly all fall short, and also the ways in which we fall short.¹⁴

8. Thus far we have focused on the situationist attack on virtue ethics and psychology, led in philosophy by Harman and Doris, and on the proposed situationist alternative. For their part, advocates of virtue theory have converged on a response to that attack. Several authors have now accused situationists of adopting a crude, external, behaviorist conception of virtue psychology, one that virtue theorists reject as a caricature. Situationists are said to ignore the inner deliberative complexity so important to sophisticated virtue ethics and psychology.¹⁵

¹⁴This material was presented at the Herdecke conference on philosophy and the social sciences, where the commentator, Steven Lukes, pointed out how surprising such surprise would be given ethnic cleansing and other familiar horrors of recent history and current affairs.

¹⁵Relevant here are four illuminating articles: Rachana Kamtekar, "Situationism and Virtue Ethics on the Content of Our Character," *Ethics* 114 (April 2004): 458–91; Jonathan Webber, "Virtue, Character, and Situation," *Journal of Moral Philosophy* 3 (July 2006): 193–213; Gopal Sreenivasan, "Errors about Errors: Virtue Theory and Trait Attribution," *Mind* 111 (2002): 47–68; Rosalind Hursthouse,

We are thus presented with two conceptions of virtue psychology: (a) Crude virtue psychology (CVP) focuses directly on situation/behavior dispositions. (b) Sophisticated virtue psychology (SVP) interposes situation/attitude dispositions between situations and behavior.

According to SVP, it is the agent's character that holistically explains his conduct. In contrast with the CVP attacked by Doris, by Harman, and by Nisbett and Ross, the character of interest to SVP is a broader whole that includes fundamental motives, desires, and even goals. The rational agent works to integrate these into a coherent whole. Given how often we are evaluatively conflicted—how often we must make hard choices—it can hardly be a surprise that we fail to be cross-situationally consistent at the level of external situations and attendant behaviors. The consistency is found, rather, internally, in the complex inner structure that is one's relevant character.

In order to possess the virtue of kindness, for example, one need not behave kindly whenever one is in a kindness-relevant situation. And the same goes for honesty and other traits. When values conflict in a situation, a value other than kindness may take priority. The virtues of SVP are rational virtues manifest primarily in right choices made through proper rational deliberation.

Here, now, is a way to develop this alternative picture: When deliberating on a yes-or-no choice, one faces a rational structure of pros and cons, of reasons for and reasons against. Here I mean *good, factive* reasons. These constitute the rational structure of the situation. One could think of this as a one-dimensional vector space, with positive and negative vectors as the pro and con reasons. Additional options, beyond our simple yes-or-no case, will of course import a more complex vector space.

Corresponding to such a rational structure more or less well is one's motivational structure constituted by positive and negative motivating reasons, reasons that psychologically attract one to a certain choice or repel one from that choice to various degrees.

"Virtue Ethics," *Stanford Encyclopedia of Philosophy*. Available at <http://plato.stanford.edu/contents.html>.

Take the case of Tom, our young man comfortably seated when the elder approaches him on the bus. Among the reasons that structure his situation rationally are her risk of falling and her evident physical and emotional suffering, whereas he is about as well off standing as sitting. Indeed, we can surely imagine the scenario so that the balance of reasons strongly favors his ceding his seat. Practical wisdom would then require that his motivating reasons reflect the objective reasons, so that the relative weights of his motivating reasons correspond to the relative weights of the relevant objective reasons. If the reasons to cede constituted by the risk and suffering of the oldster outweigh the reason not to cede constituted by the youngster's slightly greater comfort, then the motivating force of those weightier reasons should also psychologically outweigh the motivating force of the less weighty reason, and should do so by a corresponding margin. Thus, the motivational structure in the mind of the agent should reflect the rational structure of the situation faced.

One manifests practical wisdom in any given situation to the degree that one's motivational structure reflects the relevant rational structure in that situation.¹⁶ Of course, one cannot thereby *manifest* practical wisdom unless one *possesses* such wisdom, which is not something one can do only ephemerally and locally. Rather, one has true practical wisdom (a) to the degree that one

¹⁶Here I am assuming that the subject has access to the relevant plain facts. But it is a nice question how extensively factual perceptiveness constitutes practical competence and is detachable from value perceptiveness. Practical competence is a function not so much of mere factual perceptual acuity (the ability to perceive sharply) as of the foregrounding of facts that are relevant, good reasons for what is objectively required of one by the situation faced. And this is a *normatively constituted* competence: it requires systematic foregrounding of the normatively relevant as such. Here the phenomenon of inattentional blindness is highly relevant. Take the seminarians in a hurry. It is not implausible that their disregard of their fellow human in need bespeaks not so much callousness as inattention, and indeed inattentional blindness. However, what seems still an open question is the extent to which such blindness is to be classed with culpable neglect as a moral failure. And this is, again, of a piece with the question of whether failure to foreground the morally relevant is itself thus morally relevant. Compare the following: Daniel J. Simons, "Inattentional Blindness," in *Scholarpedia*, www.scholarpedia.org/article/Inattentional_Blindness.

has a stable disposition to reflect in one's motivational structures the pertinent rational structures of the various situations that one enters in the course of human relations and other events, and (b) to the degree that this disposition is robust and global. One is practically wise in proportion to how well one appreciates the rational force of the pros and cons by giving them motivationally the respective weights that they deserve.

9. Does this correction of CVE give to SVE and psychology what they need in order to repel the situationist attack? Not fully. It does so, surely, only to the extent that subjects in situationist experiments reflect in their motivational structures the rational structures of the situations faced. But it is quite implausible that they do so. Either they go wrong in giving too much weight to factors that should not have so much weight, such as punctuality, or they give too little weight to factors that should have more weight, such as whether they are inflicting severe pain: they give the wrong relative weight to such infliction by comparison with following an experimenter's instructions.

The switch from crude virtue theory to sophisticated virtue theory is nevertheless important, if only for the sake of understanding properly the subject matter of our controversy. However, it is implausible to suppose that this proper understanding immediately provides a satisfactory response to the situationist attack. It is incumbent on virtue theory to grant that the experiments do raise legitimate doubt as to how global and robust is human practical wisdom and how global and robust are its more specific component virtues, such as kindness, human decency, honesty, courage, and the rest.

That being granted, it would hardly follow that humans have no practical wisdom, none of the structure of virtues that, when properly integrated, constitutes such wisdom. This sort of invalid inference is the crucial weakness in situationism to be probed here. Indeed, probing this sort of fallacy, once spotted, deflates the situationist attack even when aimed against the crude version of virtue theory. If the attack fails even against the crude version, it will be an even worse failure against the more sophisticated variant.

B. IN DEFENSE OF VIRTUE THEORY

1. The following defense of virtue theory is based on an analogy between moral competence and driving competence. Let us define driving competence as a disposition to produce driving that is *safe* when one is at the wheel and *efficient* in routing one to one's destination upon getting directions. Recall the contrasts applied earlier to traits, between the robust and the frail (or between the broad and the narrow or the global and the local) and between the distinctive and the vulgar. These can be seen to apply with similar plausibility to driving competence. Someone's driving competence may be limited to quiet neighborhoods, for example, and may not extend to busy highways or to city driving. Such competence is, then, not as robust as it might be. The minimal driving competence required for safe and efficient driving in a sleepy village is very widely shared, moreover, and not as rare or distinctive as the physical abilities demanded by Formula One races or the navigational adroitness required by a reticulated old city.

Evaluatively relevant behavioral differences in instances of driving (one bad, one good) will not be explained by the shared vulgar competence of the two drivers. Any such behavioral difference explained by appeal to competence levels must, of course, appeal to some *difference* in such levels. Absent any such difference in competence, one must appeal to some difference in situation.

2. Consider now the factors that have been found to affect the safety or efficiency of driving:

- (a) Brightness of light, even when the road is visible.
- (b) Whether one is on a bridge when it is cold and wet (roadways on bridges being colder and potentially more slippery).
- (c) Whether one is using a cell phone.
- (d) One's blood level of alcohol.
- (e) Whether one has received directions orally or through a map.

And so on. For some or all such factors, it must at some point have been surprising that good driving does depend on them. How should one's folk theory of driving respond to such discoveries?

3. One possible response would be *driving situationism*, as follows:

- (a) Situations are dominant in the explanation of evaluatively relevant differences in driving behavior.
- (b) The robustness of our driving competence is rendered problematic: it is found to vary surprisingly with respect to previously unsuspected factors.
- (c) Personal integration in driving competence is not as widespread as one might have thought: thus, operational competence at the wheel does not necessarily go together with navigational competence.

These three theses are respectively analogous to the three listed by Doris in presenting the essentials of situationism (and quoted earlier). However, there has been no tendency to adopt driving situationism. Nor is it plausible to conclude that driving competence is just an illusion or that we make a fundamental attribution error in taking people to be competent drivers when they display a good long stretch of such driving and when we explain the quality of their driving through their relevant competence.

4. Situations *will* indeed bear on our explanations of drivers' performances. But then any competence, indeed any disposition, will issue in a certain behavior only given certain triggering conditions and given a certain shape and situation of the host of that disposition. Behavior will, in general, have a two-ply explanation, one strand being the disposition, the competence, and the other strand being the relevant triggering conditions that elicit the manifestation of the competence from the host in a given shape and situation. This is obviously true of dispositions in general. A sugar cube dissolves not just due to its solubility but also due to its insertion, while in normal shape, in a normal situation.

Of course, any factor that, to our surprise, affects the quality of one's driving shows one's driving competence not to be all we had assumed it to be. But this argues not so much for the abandonment of competence psychology as for its correction. With every such discovery we may need to change our view of the shape or situation that underlies a competence.

Sometimes we have a choice between reasonable alternatives. Take the different ways one could respond to four factors that bear on safe driving: the brightness of the light, the amount of traffic, one's blood alcohol level, and one's visual acuity. Plausibly, the alcohol level belongs with the acuity as an inner basis for safe driving. After all, one's competence can change, as can even one's acuity. By contrast, the amount of traffic is part of the situation, along with the quality of the light. How competent one is to drive on a certain road at a given time pertains to both sorts of factors, both to one's relevant shape and to one's relevant situation.

5. You might argue, as follows, that the analogy between driving and moral competence is very limited and potentially quite misleading: "You do not improve your moral competence by avoiding situations in which it will be severely tested in strict analogy to how you improve your driving competence by avoiding bridges when it's wet and wintry." Even if this point has some force, the analogy remains effective. For one thing, you need not avoid the bridges so long as you heighten your awareness of the risk and adjust your behavior accordingly. Through such heightening and adjusting, you *thereby* become more competent to then drive on such bridges. Your competence to so heighten and adjust more generally will give you a general driving competence beyond your competence to then drive on a specific bridge. Similarly, one way to improve your moral competence, surely, is through analogous competence to heighten your sensitivity to moral danger and to proceed with corresponding care.

6. Here are some lessons from our exercise.

The discovery of factors bearing surprisingly on our moral competence might more reasonably lead us to improve it than to reject its existence. So much for the normative lesson important for virtue *ethics*. (When we discover the bearing of sleep deprivation on good driving, we tend to avoid driving when sleep deprived, thus improving our competence to drive safely.)

As for the bearing of the experimental results on virtue *psychology*, what they call for is, again, correction, not rejection. The sort of practical wisdom that explains a normal human being's

ordinary behavior varies in degree from agent to agent and also somewhat in structure, given (a) how variably humans can fall short and (b) how implausible it is to postulate a single acceptable motivational structure with sharp outlines.

We have learned of factors with a previously unsuspected bearing on our morally relevant conduct, factors that dull our discernment of the moral or other practical demands in certain situations, leaving nearly all of us with less practical wisdom than we had commonsensically claimed. Similarly, driving competence is affected by certain factors in ways that once proved surprising. A driver on his cell phone while crossing a bridge in wintry twilight will still likely reach his destination without incident, but we have long known that accidents under such circumstances are significantly more probable, so that smooth driving (prior to the discovery of the factors surprisingly relevant) is then *less* explicable through driving competence than we earlier thought and more a matter of situational luck.

There are two sorts of relevant discoveries. Some factors might be thought to reduce a driver's competence temporarily: thus sleep deprivation and alcohol. Others might be viewed as conditions not covered by a driver's competence. That is to say, the competence might be viewed as more local, less robust, than we earlier thought. Thus, driving competence is now known to be affected by ambient light, by cell-phone use, and by whether one is on a wet bridge in winter, to cite just three relevant factors.

On this view, the quality of driving competence is determined by a certain inner state of the agent's and by a set of distinguished situations. Thus, the inner shape of the driver can change so as to lower the degree of situation-relative robustness of the competence.

In an alternative conception, a driver's fundamental competence does not necessarily change with the noted inner changes. In this alternative, our driving competence is viewed as stable throughout, although the difference in expectable outcomes shows it to be less robust than previously thought since the good outcomes are restricted relative to alcohol level and other inner states.

Whichever of the two views we take, we have discovered through the years that we possess neither the robustness of practical wisdom nor the robustness of driving competence that we had once

optimistically self-attributed. But we overreact if we leap from that fact to the belief that driving competence and practical wisdom are just illusions.

C. THE VIRTUE/SITUATION CLASH: GENUINE OR ILLUSORY?

Finally, the supposed clash between situationism and virtue theory turns out to be just an illusion owed to confusion and misconception.

1. Virtue theory makes no claim that normatively relevant behavioral variation is to be explained exclusively or even mainly through difference in virtues. Virtue theory is quite compatible with the view that humans are pretty much alike in the degree of virtue that they normally attain. *Compatibly with this, human conduct might still be universally explicable through the attribution of virtue to the agent.* Variability of marble rolls is not explicable through the universally shared disposition of marbles to roll, but each marble roll might still be explained in essential part through that disposition and its underlying basis, the rigid sphericity of a marble.

2. Virtue theory should accept that the situationist experiments have shown humans to be less practically wise than folk virtue theory had imagined. Our practical wisdom now seems less robust or global than we had believed. While concluding that there are “no virtues of the sort that virtue theory had imagined,” situationists do not leap all the way to a Skinnerian nihilism of behavioral dispositions. But virtue theorists, for their part, should *accept* that there are no virtues *of the sort virtue theory had imagined* since, after all, we are *less* robustly, globally virtuous than we had believed commonsensically. Situationists conclude: forget virtues, explain by situations! But wait. Remember, behavioral explanation is two-ply, requiring, when laid out fully, *both* the relevant particulars of the situation *and* the relevant non-Skinnerian dispositions. These latter look for all the world like traits, competences, or virtues, however robust they may or may not turn out to be. Neither extreme position seems acceptable: neither that

such traits explain with no situational help at all nor that such situational particulars explain with no dispositional help at all. Behavioral explanation is two-ply when laid out fully.¹⁷

3. In light of the preceding, it would seem that any remaining substantive disagreement will be over degrees: just how robust are the relevant human virtues?¹⁸

4. As for the situationist recommendation that we should assess the situations we enter for the relevant risks, this is fully endorsed by virtue theory, which takes the discernment and proper weighing of such risk to be among the most important virtues in any given domain of human performance.

5. We have been given no more reason to doubt that instilling virtue is worth the effort (on the part of parents and teachers) than to doubt that instilling driving competence, and requiring its certification, is worth the effort (on the part of relevant government authorities).¹⁹

¹⁷What is more, that *concedes* to situationism a restriction of the relevant virtues to the innermost seats of our relevant behavioral dispositions. However, common sense and virtue theory can and do recognize abilities that *encompass* the shape and situation of the agent. Thus, driving ability can be conceived to vary with changes in the shape or situation of the driver. *Competence* to drive safely on a certain road can vary with the driver's alcohol level and with the ice on the road. Corresponding varieties of competence will be distinguished in the theory of competence developed in what follows.

¹⁸And there is indeed a surprising array of factors that can influence our behavior subliminally and unexpectedly. For example, noise levels seem to affect helping behavior (Kenneth E. Mathews and Lance K. Cannon, "Environmental Noise Level as a Determinant of Helping Behavior," *Journal of Personality and Social Psychology* 32 [1975], 571–577), as can fragrances in shopping malls (Robert Baron, "The Sweet Smell of . . . Helping: Effects of Pleasant Ambient Fragrance on Prosocial Behavior in Shopping Malls," *Personality and Social Psychology Bulletin* 23 [1997], 498–503). But something similar is, again, true of our driving performance.

¹⁹The third most important experiment bearing on our issues is the Stanford Prison Experiment of Philip Zimbardo, who concludes as follows in a retrospective: "The critical message then is to be sensitive about our vulnerability to subtle but powerful situational forces and, by such awareness, be more able to overcome

FURTHER READING

- Athanassoulis, Nafsika. "A Response to Harman: Virtue Ethics and Character Traits." *Proceedings of the Aristotelian Society* 100 (2000): 215–21.
A response to Harman, below.
- Besser-Jones, Lorraine. "Social Psychology, Moral Character, and Moral Fallibility." *Philosophy and Phenomenological Research* 76 (2008): 310–32.
- Doris, John M. *Lack of Character*. Cambridge, UK: Cambridge University Press, 2002.
- Harman, Gilbert. "Moral Philosophy Meets Social Psychology: Virtue Ethics and the Fundamental Attribution Error." *Proceedings of the Aristotelian Society*, new series, 99 (1999): 315–31.
- Kamtekar, Rachana. "Situationism and Virtue Ethics on the Content of Our Character." *Ethics* 114 (2004): 458–91.
- Slingerland, Edward. "The Situationist Critique and Early Confucian Virtue Ethics." *Ethics* 121 (2011): 390–41.
- Sreenivasan, Gopal. "Errors about Errors: Virtue Theory and Trait Attribution." *Mind* 111 (2002): 47–68.

those forces. Group pressures, authority symbols, dehumanization of others, imposed anonymity, dominant ideologies that enable spurious ends to justify immoral means, lack of surveillance, and other situational forces can work to transform even some of the best of us into Mr. Hyde monsters, without the benefit of Dr. Jekyll's chemical elixir. We must be more aware of how situational variables can influence our behavior." (Philip G. Zimbardo, "Revisiting the Stanford Prison Experiment: A Lesson in the Power of Situation," *Chronicle of Higher Education* 53 [March 30, 2007]: B6.) His related book is *The Lucifer Effect: Understanding How Good People Turn Evil* (New York: Random House, 2007). This critical message is one that virtue theorists can applaud. Yes, just as alcohol can deprive us of our driving competence, so group pressure can apparently deprive us of our moral competence. This would not show such competence to be an illusion, however, nor would it tend to show that belief in it is based on a fundamental attribution error.

Virtue Epistemology and a Theory of Competence

SOCIAL PSYCHOLOGY HAS now provided a critique of virtue ethics that is applicable specifically to virtue epistemology, but there is, of course, a specific defense in line with the reasoning in the preceding chapter. What follows will lay out and buttress that defense.

We first lay out a theory of competence in general and its application to epistemology. This is meant to round out our virtue epistemology, which must essentially rely on a theory of competence. Our theory of competence will then provide a basis to address the situationist critique of virtue epistemology. That critique parallels closely the corresponding critique of virtue ethics. Can our earlier response to the critique against virtue ethics be adapted to serve similarly in defense of virtue epistemology? Read on.

A. A THEORY OF COMPETENCE

1. A competence is a disposition (ability) to succeed when one tries. How are such dispositions, in general, constituted? When complete, they have a triple-S constitution. Accordingly, we can distinguish three sorts of dispositions: the innermost (seat), the inner (seat + shape), and the complete (seat + shape + situation). With regard to one's competence in driving, for example, we can distinguish between (a) the innermost driving competence that is seated in one's brain, nervous system, and body, which one retains even while asleep or drunk; (b) a fuller inner competence, which requires also that one be in proper shape, that is, awake, sober, alert, and so on; and (c) complete competence or ability to drive well and safely (on a given road or in a certain area), which

requires also that one be well situated, with appropriate road conditions pertaining to the surface, the lighting, etc. The complete competence is thus an SSS (or an SeShSi) competence.

2. Again, a competence is a disposition to succeed when one tries. So exercise of a competence involves aiming at a certain outcome. It is a competence because it is a disposition to succeed reliably enough when one makes such attempts. A competence is hence necessarily a competence to \emptyset successfully, for some \emptyset , one tied to a conditional of the following form: that *if one tried to \emptyset , one would (likely enough) succeed (given one's relevant SSS profile)*.

3. Competences come in degrees of reliability, along with a threshold. However, in order to possess a competence to \emptyset it is not enough that the following conditional be true: that one would \emptyset reliably enough if one tried to \emptyset , that one would not too easily try to \emptyset without actually \emptyset -ing. This could be true simply because, knowing one's limits, one \emptyset s rarely and only when in a narrow range in which one would indeed succeed. Thus, one might restrict one's shots as an archer to situations in which one is within two feet of the target. This might show good shot selection, assuming one is too unreliable at any longer distance, but it would not show archery competence.

4. What, then, is required for possession of a competence? Archery competence, for example, requires a "sufficient spread" of possible shots (covering enough of the relevant shapes and situations one might be in as an agent) where one would succeed if one tried, an extensive enough range across which one might easily enough try.

It might be objected that even if one is tied down, so that, knowing one's condition, one *would not* try to shoot, this does not take away one's competence. This may be thought to refute the requirement that there must be a close enough sphere of possible worlds wherein one does take shots. However, if one is tied down, that *does* surely bear on one's complete SSS competence, even when it does not affect the inner SS competence or the innermost S competence. One is, then, too poorly situated to retain the complete

competence to hit the target with one's arrows. Moreover, if indeed one is disabled psychologically by a phobia, so that one cannot so much as try to shoot, then this does take away even one's innermost skill *S*. One is now no longer so structured psychologically as to be competent to ϕ .¹

B. COMPETENCE AS A SPECIAL CASE OF A DISPOSITION TO SUCCEED

1. Every competence is, again, a disposition to succeed when one aims at a given objective, in certain (favorable enough) conditions while in (good enough) shape. But not every disposition to succeed is a competence.

Thus, take a soccer goalie who faces an opponent kicker in a tie-break attempt. Suppose the goalie anticipates with a jump to the left. In such circumstances, if the kicker attempts to score by kicking to the left, the goalie is likely to make the save. But if the goalie chooses arbitrarily between jumping to the left, jumping to the right, and staying put, then her save is not nearly as creditable to her competence as it would be if prior experience with this kicker enabled her to anticipate the direction of the kick. If the goalie chooses her direction arbitrarily, then her anticipation of the kicker's choice is insufficiently reliable to qualify as a real competence.

2. Consider also a rank beginner who receives service for the first time in a game of tennis. When the streaking serve comes at him, he swings blindly and wildly. If the coinciding trajectories of the ball and the racket produce a wonderful return, that does not reveal an unsuspected competence seated in that receiver. Yet if

¹But here we have a choice. My own preferred choice is to say that one can have a "quasi-competence" to hit the target provided one is so propertied and situated that *if* one were to try to hit the target, one would be likely enough to succeed. But a *full, proper* competence would require also that one not be so disabled that one could never even try, no matter how desirable it might be to try or to succeed in hitting (a given suitable target).

he were ever to repeat that swing in those conditions, a successful return would result. Take the vigorous enough swing with a certain arc, when the ball approaches with a certain speed and direction, so that it and the racket will happen to coincide properly. Swinging that way in the same setting will reliably produce such a wonderful return. But this does not reveal a competence, despite the fact that there is in that beginner a disposition to succeed reliably in such a situation if he swings thus. The beginner has an ability to succeed if he tries that way, because he is quite able to try that way.

3. Flukey success will generally admit description similar to that of our lucky return of service. The performer will occupy a situation wherein he is disposed to succeed if he tries in a certain way in that situation, and that is how he does try in that situation in the case at hand. There is a way of swinging available to our tennis receiver such that, relative to the specified situation of a ball traveling toward him at a certain angle and speed, if he were to swing that way, success would ensue.

What the tennis novice lacks is of a way of swinging his racket that he knows would likely enough succeed, *and such that he can likely enough produce it at will as a means to success.*

Note well, however, that even someone with such command of that way of swinging, someone thus possessed of an “ability” to succeed that way, might still lack competence to attain the relevant success. For he might *be unlikely enough ever to actually swing that way in pursuit of relevant success* and might instead systematically opt to try in ways that are *unproductive.*

4. The tennis player’s swing will manifest tennis-return competence precisely because of his disposition to swing *selectively* that way, along perhaps with competently helping to bring about the right sort of situation, for example, by positioning himself well in preparation for receiving the serve.

Nevertheless, we cannot require of a performer that he bring about every aspect of a situation that makes it an appropriate situation to constitute a competence to succeed. Thus, the tennis

player's competence is relative to certain lighting, wind, and precipitation conditions that he is not required to bring about.

Moreover, if a player tries to improve by practicing in challenging situations with high risk of failure, that cannot be held against him if those situations are not included among the situations within which his competence is supposed to guide his performance to success. Thus, a tennis player might practice with a machine that feeds him balls much faster than any he would ever encounter in an actual game, or he might practice in lighting conditions that are significantly subpar. His failure rate would, of course, rise in such circumstances, wherein he was still willing to test himself. But that need not affect his degree of competence.

Thus, suppose Improver prefers to play with far better players, whereas Dominant likes to crush weaker opponents. Dominant thus has a much higher success rate for the various performances in a tennis player's repertoire, both actually and dispositionally (given his penchant for weak opponents). But Improver may be the more competent player nonetheless.

5. Not every disposition to succeed when one tries constitutes a competence, then, although every competence will be constituted by a disposition to succeed when the agent is within certain ranges of shape and situation. A disposition to succeed is thus plausibly made a competence by some prior selection of shape/situation pairs such that one seats a competence only if one is disposed to succeed reliably enough upon trying when in *such* a shape/situation pair. Whether a particular shape/situation pair is appropriate will, of course, vary from domain to domain of performance.²

²For simplicity, I leave aside restrictions on *how you acquire* the relevant elements of competence, such as the seat and the shape, restrictions that have come to the fore with the cyclist Lance Armstrong (regarding drug-induced shape) and with the baseball player Alex Rodriguez (regarding drug-derived seat). Each of these athletes enhanced his performances by enhancing his complete SSS dispositions to succeed, but these dispositions did not remain competences once drug-enhanced. They did not remain ways to attain fully apt performance, performance creditable to (proper) athletic competence and thus creditable to the athlete's relevant competence.

6. A competence is a disposition to succeed, but it must be such a disposition properly restricted with respect to the three Ss—seat, shape, and situation. At least the shape and situation must be restricted to the appropriate. And these restrictions are imposed somehow within the domain of the relevant performances.

Again, all competences are dispositions to succeed, but not all dispositions to succeed are competences. In order to ground an innermost skill or competence, an Se state must dispose the agent to succeed broadly and reliably enough in enough of the relevant, pre-selected shape (Sh) and situation (Si) conditions. The fuller inner or complete competence is then structured by such an Se combined with Sh and Si conditions within the proper ranges. Thus, a complete SSS (or SeShSi) combination would dispose the agent to succeed reliably enough if he tried.³

7. The terminology of competence is flexible. Consider the competence to drive safely on a certain road. Take the complete competence to drive safely at a time *t* on a certain stretch of that road. This requires that the subject satisfy all three of the S requirements. She must have a seat on which is grounded the conditional that she *would* drive safely (the requisite driving skill) reliably enough. And she must be in the right shape (thus, awake and sober) and properly situated with respect to that stretch of road (so that, for example, the road is not covered with oil).

However, we also allow a competence to drive safely that the driver can keep even when asleep or drunk. This inner skill would combine with appropriate shape and situation to the effect of a complete competence. Clearly one can be a good, safe driver, in possession of a competence (an *inner* SeSh competence) to drive well and safely even if the nearby roads are all covered with oil. One's relevant status as a good driver is not beholden to the condition of nearby roads. The fact that one is not competent to pro-

³Many domains are set largely or wholly by convention, as are games, sports, and artistic domains. Other domains of human endeavor are set by our nature and needs and by evolutionary teleology. Much is set by the approval of the group or of the species. But there is surely room for the group to fall short: moral leaders, for example, can lead the way to recognition of competences previously overlooked.

duce good driving on bad, even impassable, roads does not take away one's competence as a good driver. Nor would one's skill be impaired by inebriation or sleep. None of this affects the fact that one has the competence of a good driver. What is required for this is only that one seat that competence: one must host the seat/basis that, together with *appropriate* shape and situation, would dispose one to drive well if one tried. (But we should distinguish between being a good driver and having the competence of a good driver. The latter would require only that one have the skills required for producing good driving if one tried, while the former would require also that one at least normally try when one drives.) Such an innermost competence, which abstracts from appropriate shape and situation, often constitutes a "skill."

C. COMPETENCES, EPISTEMIC AND OTHERWISE

1. We can distinguish judging that *p* from merely thinking that *p*. The latter requires that one affirm in the endeavor to get it right on the question whether *p*. But that is compatible with just psychologically guessing. Full judgment requires more: one must endeavor not only to get it right on that question but also to do so aptly enough.

2. And here we find a reason why *judgment is automatically on a reflective higher order*. For what does it mean that one endeavors to get it right "aptly enough"? Succeeding in this endeavor requires that one affirm in such a way, in such circumstances, while in possession of sufficient relevant competence that, given one's complete relevant competence, one would be likely enough to affirm correctly. That is to say, one's success rate for such attempts must be high enough. And what would go into such an endeavor? What must one consider as one aims to affirm aptly and therefore reliably enough? Must one not consider one's relevant SSS situation and how likely it is that a relevant attempt manifesting such a complete competence would succeed?

In that case one takes a higher-order attitude toward one's three options: affirming, denying, suspending. One considers which of

these offers the best risk/reward value. What one considers, then, is the reliability of affirmation or denial on the question at hand. One thus considers whether double-omission would not be preferable given the risks involved in either affirming or denying. And this is all on the second order, since one must consider one's relevant, first-order, complete competences and the first-order options of affirming, denying, and double-omitting.

3. For performances more generally, something similar is true across the board. Consider, for example, three-point shots in basketball. Even a player who overconfidently takes low-percentage shots too frequently may retain an excellent ability to sink three-point shots close enough to the three-point line. So success in such an attempt does seem properly creditable to her competence, even if she tries too often when she should not do so, when she is so far from the basket that her reliability plummets. This is especially plausible when she is *fully aware of her limits* but tries to succeed even beyond them as a deliberate risk taker. But it remains plausible enough even for someone who lacks competent and full enough knowledge of her limits and shoots beyond them only because she incorrectly takes herself to be reliable enough even then.

The difference between those two cases is as follows. In the former, the agent can still perform with reflective and full aptness when she knowingly performs within her limits. In the latter, the agent no longer performs with reflective and full aptness, at least not when too near her threshold of reliable-enough performance. Nevertheless, she can, of course, still perform with animal (and indeed superlative) aptness even so.

4. But what is the aim of a basketball shooter? Does she aim merely to get the ball in the hoop? Normally performers do not aim just to succeed in the sense of attaining their basic aim no matter how unreliably. Normally they aim to succeed aptly enough while avoiding too much relevant risk.

5. Suppose that the three-point shooter's reliability from a certain distance is above the relevant threshold but *indiscernibly so*

to her. A statistician-coach-observer might know perfectly well that the player's reliability is still barely above the threshold by use of a device that measures with exactitude her distance from the basket. That way he can tell that she is still reliable enough at that distance. But the player knows no such thing. As she shoots from that distance, however, she may attain her basic aim: namely, to sink that basket. And her shot may be apt by succeeding in a way that manifests a competence the statistician knows her to possess even at that distance. Her shot is not just successful, not just competent, but also apt.

6. Although she attains her first-order objective, and does so aptly, however, the player does not aptly attain her reflective aim of succeeding aptly enough. She does succeed reliably enough, but her attaining *this* objective is a matter of luck. The full success of her shot (its going through the hoop reliably, aptly enough) is not *fully* creditable to her, given the important element of luck involved. So her first-order success will be apt, but it will not be meta-competent and hence not meta-apt, and so it will not be fully apt.

D. MORE ON THE NATURE AND EPISTEMOLOGY OF COMPETENCES

1. Competences are a very special case of dispositions. Suppose a solid iron dumbbell would shatter upon hitting a certain surface only because a hovering fiend is determined to zap it iff it hits the surface, and just as it does so. This would make the relevant conditional true: the dumbbell *would* shatter on that occasion, but this does not make it fragile. For the latter to be the case, one would need, at a minimum, for the dumbbell to shatter whenever relevantly impacted in a relevant combination of condition and situation.

2. Just what makes a combination of skill, shape, and situation *thus* "relevant" is an interesting, and neglected, question. Such combinations constitute the innumerable competences of interest

in the many performance domains recognized by human communities. Such domains—whether athletic, artistic, intellectual, medical, scientific, or legal—contain performances aimed at certain distinctive aims, along with the competences whereby performers might succeed more or less competently. When a success manifests such competence, then it is apt, and only then. The SSS profile that underlies any such competence is not always subject to full and detailed linguistic formulation. But in that respect it joins much else of substantial human interest and importance, such as when conduct counts as polite.

What conduct does count as polite? This is much easier to discern by perception in a particular situation than it is to formulate in full detail. This is so for politeness in general, as well as for particular matters of politeness, such as the distance appropriate for normal conversation in person and the tone and volume of voice that counts as rude, to take just two out of many instances.

Moreover, that is not to say that rules of polite conduct are just a myth. Not all rules need to be formulated linguistically. If what is polite is defined by human convention, however, then there must be rules, in some broad sense, that communities agree upon ahead of time. There need not be a constitutional convention to institute those rules. Their institution may be more organic, less artificial, than that. Moreover, the agreement within the community will not require so much as linguistic communication. It may be instituted more implicitly than that through persistent profiles of approval/disapproval, praise/blame, and systematic and implicit manifestation in individual and social conduct.

3. Returning to our own issues, this pattern of normativity is applicable also to competences. Thus, the SSS profiles of the competences that acquire salience for communities may be determined not by antecedent, linguistically formulable convention but rather by persistent implicit profiles of the sort that determine the content of politeness for that community.

What is more, the relation is not just a matter of coincidence or mere analogy. There is, after all, such a thing as competence in etiquette, and even in morality, so that these domains might be

conceived either in terms of rules implicitly operative or in terms of competences determined by proper community interest.⁴

E. COMPETENCE, SAFETY, AND RELIABILITY

1. In order to possess a performance skill (or the seat/basis of a competence), one need not satisfy any general reliability requirement. We have seen how one can exercise one's skill too often in inappropriate (and irrelevant) shape or situation so that it would fail with an extremely high frequency while, of course, remaining in place (to be thus frequently exercised).

2. Skill might be present and exercised, leading to success without being manifest in that success. Thus, a skilled archer might shoot skillfully, and the arrow might be deflected by an unexpected gust. However, a hovering guardian angel might happen by, determined to correct the trajectory of any well-shot arrow so that it will hit the target when (but only when) it would have done so but for the unlucky intervention of a gust. In this case, the archer's skill is exercised, and in the situation as pictured the skill does partly underlie a disposition to succeed. But this disposition does not amount to a true competence, since it is situationally based essentially on the angel. And the angel's presence is not among the relevant background conditions. What makes that innermost archery seat a true skill is that it would combine with appropriate shape and situation so as to dispose the archer to succeed reliably enough if she tried to hit a target. However, in the actual situation the archer's seat/basis fails to combine with an appropriate situation. The angel is an ad hoc situational factor not included among those that agents in that domain can rely upon by default. No

⁴Of course, that is not to say that all such domains have their normative source or basis in human convention. Nor need we suppose that convention-based domains lie beyond objective assessment in global or specific respects. They may still be subject to such assessment on how well they further any value they may be designed to secure.

such angel figures in the situations of interest in the domain of archery, the situations relative to which we assess true archery competence and achievement. What makes that innermost archery seat/basis a true skill is that it would combine with appropriate shape and situation so as to dispose the archer to succeed reliably enough if she tried to hit a target.

Consider the archer who shoots with the unexpected gust about to cross the field while, unbeknownst to him, a guardian angel is ready to intervene. This archer does not earn proper credit for his success, which does not really manifest competence. And the reason for this, I suggest, is that he does not shoot when in appropriate shape, in an appropriate situation. The appropriate shape and situation for archery competence cannot depend essentially on interventions by an ad hoc guardian angel that could not have been relied upon by default. The appropriate shape and situation must be taken for granted as background for archery shots. Relevant credit will accrue to the archer's skill only if the success of the shot does manifest that skill. Such credit will normally accrue only if the agent's skill produces success with the agent in appropriate shape and situation.

3. Skills do come in associated clusters. An archer who knows of the guardian angel, with her power and intentions, might have an enhanced archery skill that takes that into account. Such an archer might indeed manifest her enhanced archery skill when the gust comes along and the angel puts the arrow back on track. The success of her shot will thus be properly credited to that archer's enhanced skill and thus properly credited to her. (Compare the conductor who makes skilful use of orchestra players in pursuit of his proper ends.)

Similarly, an archer could deliberately frustrate the attempt of a hovering demon to take his arrow off course. Here, again, an enhanced competence would be in play, to which the shot's success could be credited properly.

4. In summary, the success of a performance in attaining its aim is properly credited to a skill—and to its grounding innermost seat of competence—only if that skill counts as a skill relative to pre-

selected shape/situation combinations of interest in the domain of that performance and the performance's success manifests the skill. The skill must produce the success reliably in combination with the agent's relevantly appropriate shape and situation.

F. THE SITUATIONIST CRITIQUE OF VIRTUE EPISTEMOLOGY

1. Situationists have argued that virtue epistemology, whether responsibilist or reliabilist, faces serious problems posed by experimental results accumulated over years. These are said to be analogous to the problems uncovered by situationists for virtue ethics. What exactly are these problems for virtue epistemology?

One problem is supposed to be that we are surprisingly likely to be less reliable than we had thought and that we are made less reliable by influences astoundingly *irrelevant* to the truth of the beliefs that they nevertheless do influence. So the bottom line here is that we are *less reliable* epistemically than we had implicitly supposed all along.

2. There are two serious gaps in that critique, the first of which (though not the second) is independent of any particular theory of competence. We next consider how damaging these are.

In the first place, in order to show that we are lacking in epistemic competence, it would need to be shown that the ways in which we form beliefs are *unreliable* or *insufficiently reliable*. All that has been made plausible, however, and even this only for *some* domains of belief formation, is that we are *less reliable than had been thought* and *less reliable in ways and for reasons that are surprising, sometimes even astounding*. But from none of this does it follow that there is *any* recognized domain in which we form beliefs in some distinctive way(s), wherein we are *unreliable* or *insufficiently reliable to count as competent* (and as competent enough to regularly attain knowledge through the exercise of our competence in that domain).

Suppose, however, that we *have* been shown to be *insufficiently* reliable in *some* domain where we had taken ourselves to be reliable enough to attain knowledge competently. Even if that is so

for some specific domain, even if we are indeed unreliable believers in a particular set of circumstances distinctive of that domain, it would be fallacious to infer from this that we are more generally unreliable in some way that puts our common-sense attributions of knowledge seriously and generally in doubt.

Compare the many perceptual illusions that psychologists have detailed. These conclusively establish how seriously unreliable our perceptual systems are in certain situations. But this is not a good basis for any general rejection of our perceptual competences.

3. So much for the first supposed problem. A second problem is laid out as follows. First, it is pointed out that in a certain situation X our supposed competence to \emptyset regularly fails us. And it is inferred that we do not really have any such competence. What we have is rather something like this: we have a “competence to \emptyset when not in X.”

The fallacy in that form of reasoning is exposed if we apply it to almost any athletic competence. Take a basketball player who has a highly prized competence to sink three-point shots. Placed far enough back from the basket, at center court, his reliability will plummet. So what should we conclude from that? Must we say that, strictly speaking, he lacks competence as a three-point shooter? Has he at most a competence to sink three-point baskets “provided he is not at center court”? That is not how we normally speak, or even think, of a competence. A competence does come with implicit reference to levels of pre-selected shape and situation within which it might be exercised properly. So yes, when we invoke a competence, or a virtue, there is an implicit relativization to shape and situation. But we have seen that to be not so much a problem as a feature.

4. As we have seen, a competence is a disposition to succeed if one tries. But a competence is more than a mere ability. Your *ability* to \emptyset is a disposition to \emptyset successfully if you try in a way that is then available to you. But a *competence* requires more than that, as is brought out by a tennis beginner’s disposition to succeed if he swings a certain way when the ball approaches in a corresponding way. So, in a particular situation, the player may be disposed

to return serve well if he swings in way W, given that the ball approaches in way W', as in fact it does. This is a disposition grounded in a certain combination of seated skill, shape, and situation, true enough. But it is not a disposition that qualifies as a competence.

What is required for possession of a given competence is that one be disposed to succeed if one tries to ϕ , which requires not only that one have an ability to ϕ but also that one would employ that ability (or some other in a restricted ability range) when one tried to ϕ .

Moreover, what is required for a competence to ϕ is only a disposition based on a certain pre-selected range of shape/situation combinations. These are the combinations within which it is of interest that one succeed if one tries. The selection is made by the relevant community (or even by the relevant species, for species-wide competences and aims). The fact that one does not succeed when one tries in a certain specific sort of situation may thus have no bearing at all on whether one enjoys a certain competence to succeed by ϕ -ing when one tries to do so. That competence may, after all, involve a pre-selected shape/situation range that excludes that particular sort of situation. In sections A–E of this chapter I have tried to elucidate the nature of such competences.

5. We may thus conclude that neither situationist objection withstands a closer look. An interesting conclusion is implied by our approach: namely, that knowledge (and other human accomplishments or achievements) are inherently normative in ways previously unremarked. They are normative *in a way*, in that they are successes that manifest competences, where competences are dispositions involving pre-selected shape/situation combinations wherein specifically human accomplishment is prized (or otherwise of special interest).⁵

⁵This (along with the foregoing) sketches a theory of competence meant to help explain the nature of human accomplishment, including the epistemic accomplishment of the knower. A fuller account can be found in my *Judgment and Agency* (Oxford, UK: Oxford University Press, 2015).

FURTHER READING

- Alfano, Mark, and Abrol Fairweather, eds. *Epistemic Situationism*. Oxford, UK: Oxford University Press, forthcoming.
- Dickie, Imogen. "Skill before Knowledge." *Philosophy and Phenomenological Research* 85 (2012): 737–45.
- Miller, Christian. "Moral Virtues, Epistemic Virtues, and the Big Five." In *Naturalizing Epistemic Virtue*, ed. Owen Flanagan and Abrol Fairweather, 92–117. Cambridge, UK: Cambridge University Press, 2014).
- Olin, Lauren, and John M. Doris. "Vicious Minds: Virtue Epistemology, Cognition, and Skepticism." *Philosophical Studies* 168 (2014): 665–92.
- Pavese, Carlotta. "Skill in Epistemology." In *Philosophy Compass*. Forthcoming.
- Stanley, Jason, and Timothy Williamson. "Skill." *Noûs*, forthcoming (published on Early View in 2016).

Knowledge and Justification

THIS CONCLUDING CHAPTER sketches some main components of the epistemology laid out in the earlier chapters while providing further historical context. The chapter has three main parts. Part A briefly reprises our account of *knowledge as action* using the notion of epistemic competence, then connects this with central ideas of Aristotle's ethics and Descartes's epistemology. That account then illuminates epistemic justification in part B and radical skepticism in part C.

A. KNOWLEDGE

1. *Knowledge is a form of action, to know is to act, and knowledge is hence subject to a normativity distinctive of action, including intentional action.* Some basic elements of this account will be sketched here in part A, setting the stage for the main theses to be considered in parts B and C.

2. We begin with a notion central to Aristotle's virtue ethics, a notion of "apt performance" (for a handy label). Here is a quotation from the *Nicomachean Ethics*:

It is possible to do something that is in accordance with the laws of grammar, either by chance or at the suggestion of another. A man will be a grammarian, then, only when he has both done something grammatical and done it grammatically; and this means doing it in accordance with the grammatical knowledge in himself. (Aristotle, EN II 4, 1105a22–6)

In order to be "in accordance with" the grammatical know-how in a speaker, an utterance must attain its grammatical success

under the guidance of the speaker's own competence, so that it does not succeed just "by chance."

That introduces the notion of *aptness*, of success attributable to the competence of the agent, so that it is not just "by chance" (in a distinctive way). Whether grammatical or otherwise, a success will be apt not just because it happens to satisfy some proper standard of quality but rather because of the agent's relevant *competence to bring about* the satisfaction of that standard. This success is, then, more deeply "in accordance with" the relevant know-how or competence in the agent.

So far aptness has been attributed *no normative import*. But that changes dramatically in the fuller development of Aristotle's account.

3. How central such a concept of aptness is in Aristotle's ethics may be seen in the following passage: "Human good turns out to be activity of soul in accordance with virtue, and if there are more than one virtue, in accordance with the best and most complete" (Aristotle, EN I 7, 1098a16-17).

Just as the grammaticality of an utterance can be in accordance with the grammatical knowledge or competence in an agent, so the good quality of an action or an activity can be in accordance with the virtue seated in the agent. We are thus told by Aristotle what in his view human good *is*, how it is constituted.

Next we turn to epistemology, first to a distinction between functional belief and judgmental belief. Although our account is intended to apply to both, we focus here on the *judgmental* sort of belief, defined through the notion of an affirmation (an affirmative saying). More specifically, the notion of judgmental belief is defined through the notion of an *alethic* affirmation (an affirmation that aims at getting it right, at truth). A judgmental belief that *p* is either a disposition to affirm alethically that *p* if one tries to get it right on the question whether *p* or it is a ground of that disposition.

4. *Aptness figures crucially not only in Aristotle's ethics but also in Descartes's epistemology*. For Descartes, falsity is sufficient for error but not necessary. One can still be in error with a true belief

so long as its truth is not attributable to one's clear and distinct enough perception. That belief is then true *by accident*.

Here is a revealing passage to that effect, from the *Meditations*:

[If] . . . I simply refrain from making a judgment in cases where I do not perceive the truth with sufficient clarity and distinctness, then it is clear that I am behaving correctly and avoiding error. But if in such cases I either affirm or deny, then I am not using my free will correctly. If I go for the alternative which is false, then obviously I shall be in error; if I take the other side, then *it is by pure chance that I arrive at the truth*, and I shall still be at fault. . . . In this incorrect use of free will may be found the privation which constitutes the essence of error. (Descartes, MIV.10; emphasis added)

In this passage, the *essence* of error is in effect equated with *inaptness* of alethic affirmation, which may or may not attain truth but in any case falls short of aptness. When such inapt affirmation attains truth, this is just a lucky accident that does not manifest appropriate epistemic competence.

And here is another relevant passage:

It is also certain that when we assent to some piece of reasoning when our perception of it is lacking, then either we go wrong, or, if we do stumble on the truth, it is *merely by accident*, so that we cannot be sure that we are not in error. (CSM I: 207, emphasis added; from the comment on Principle 44 of Descartes's *Principles*)

In line with Descartes, a contemporary virtue epistemology features a concept of "error"-avoiding *aptness*.

5. Concerning attempts (whether consciously intentional or functional/teleological), virtue theory distinguishes (a) accuracy or success, where the aim of the attempt is attained, (b) adroitness or competence, and (c) aptness, where the accuracy manifests the adroitness. This is thus a AAA account of performance normativity.

6. A competence is a disposition to succeed when one tries. Whether one is competent to perform on a certain occasion depends not

only on one's pertinent skill but also on the shape one is in and on how favorable is one's relevant situation, which gives competences a triple-S structure.

7. *Surprisingly, aptness depends not at all on the safety of one's attempt.* How apt one's performance is does not depend, *without qualification*, on how likely one is to succeed if one tries.

Suppose an archer in a windy environment retains his high level of skill and good shape. Suppose his arrow in fact goes straight to the bullseye. However likely a gust may have been, none in fact intervenes. In that case, the archer does deserve credit, surely, no matter how *likely* a spoiler gust may have been.

The arrow's *actual* trajectory from bow to target is such that at no point in it does a gust impinge. And this is a nonmodal property of the trajectory. However, at many points, maybe at all points, it may still be very likely that a gust will take the arrow off course. My thought is that the latter modal property—the pervasive *danger* of diversion by a gust—is irrelevant to the evaluation of the shot. The shot is apt, no matter how unsafe. What is relevant is only that no gust *in fact* intervenes, so that the speed and direction of the arrow can fully enough manifest archery competence without intervening luck.

So safety of situation is inessential for aptness of performance. Neither safety of skill nor safety of shape is essential, either. What matters for aptness is that the relevant skill, shape, and situation be *actually* in place, no matter how safely or unsafely.

B. JUSTIFICATION: THE THINKER WITH AN ENVATTED BRAIN

1. *In what follows I argue for a knowledge-friendly epistemology: not knowledge-first but knowledge-friendly. Our aim is a theory of knowledge that will fit a unified and wide-scope epistemology.* The account to be offered is meant to cover not only knowledge and belief but also epistemic justification, without disjoining knowledge from justification.

Some think that knowledge and justification *are* disjoined conceptually and metaphysically and they focus on one or the

other side of that divide. But it is not easy to understand *epistemic* justification independently of knowledge.

We seek instead an account of epistemic justification as a particular normative status of a belief, one that enables it to constitute knowledge. Knowledge is thus constituted by belief with that epistemically favored status. This is where our account of knowledge as a form of action will be essential. Only with understanding of how knowledge is constituted can we properly seek the place of epistemic justification in that constitution.

What, then, is the place of epistemic justification in the constitution of human knowledge? How do knowledge and justification join together for an overall virtue theory?

In order to answer these questions we will need to understand how skill is manifest in a skillful performance. And in order to gain this understanding it will be helpful to consider a familiar thought experiment. I mean the envatted brain experiment, which we next try to accommodate within our virtue framework. This is the example in which a subject's brain is envatted while his course of experience is seamless, since his brain is then stimulated directly, giving the subject a stream of experience of a sort that he might naturally have. First, however, let us return to our archery example.

2. *Our archer's shot can be skillful (or adroit) without being apt.* The shot is skillful if and only if it manifests skill. When the shot is also apt, then its *success* manifests skill. But a *shot* can manifest skill without its *success* doing so.

The arrow might leave the bow headed straight to the bullseye, absent spoiler gusts. Despite the skill that *the shot* manifests because of that initial orientation, its *success* might fail to manifest skill nonetheless because of spoiler gusts that turn up, even when those gusts combine so as to yield success anyhow, if only by luck. It is by manifesting a skill that a shot gets to be "a skillful shot," because some quality of it manifests that skill.

3. *Take next an archer who not only wishes or hopes but tries to release her arrow here and now. The trying might be constituted by, or might directly yield, a certain brain state that satisfies the following condition: If the subject were in good shape, that state would*

reliably and generally enough lead to good orientation and speed for the arrow leaving the bow.

Something accidental (gustlike) might intervene *within the archer's efferent nerves*, however, affecting the outcome at her limbs, so that the shot fails. That subject then manifests her skill in a certain feature of her attempt, despite the lapse in her inner shape. The relevant feature of her attempt is that it is constituted by a certain brain state, *which would normally yield success.*

4. *Suppose a subject with his brain in a vat (BIV) wishes or hopes to raise an arm here and now, and even decides to do so forthwith, yet akratically fails to try. Of two BIVs, if one tries while the other does not, that can matter for responsibility and proper blame.* So we need to distinguish *trying* to do something from *wishing* or *hoping* that one will do it, and even from *deciding* to do it forthwith.

When an agent performs a physical act by intentional design, there is an initial physical state (however complex) that would normally bring about the intended outcome. And that brain-involving state is something the agent brings about. She does not bring it about *intentionally as such, by design*. She is unlikely to be *able to specify* that particular brain state. But an agent can bring about something *attributable to her as her doing* even without bringing it about *by intentional design*. That doing is then a *deed*. A “deed” is a doing attributable to the agent as her own doing.

For example, when signing a form at a government office, I may bear down so as to sign in one doing several carbon copies of the form (in an example owed to Donald Davidson). I do not know that I am signing the fifth copy (or even the fourth, for that matter), as I think there are only three. So I may take myself *not* to sign a fifth copy, and yet may do so anyhow, as my own doing, attributable to me. I sign that fifth copy even if I do so *unintentionally (not by design)*, unaccompanied by any corresponding intention, either concurrent or aforethought. That doing of mine is then an attributable doing, despite being unintentional (not by design). It is not a *mere* “doing,” as is my squashing of a rabbit by falling on it when pushed unconscious off a cliff.

5. *Perhaps, then, I can bring about a certain brain state, one normally sufficient for the rising of my arm, even if I do not bring it about by intentional design, either concurrent or aforethought.* What is more, my doing so may be a deed, an attributable exercise of competence and even skill, despite not *being an intentional doing (by design)*.¹

Exercise of skill requires only that one manifest a disposition to succeed, which one might do through a *deed*, even one that is not *intentional* (by design). Here is another example. A pianist in performance may press a certain key with the third finger of her right hand. She does not, then, pick out that particular action as such for that very instant. So she does not press that key with that finger at that moment *by intentional design, either concurrent or aforethought*. And yet it is something she does, the doing of which is attributable to her as her doing. Moreover, that deed of hers manifests her *skill* as a great pianist.²

Similarly, we might manifest our competence to raise our arm, and indeed our skill for so doing, when we try to do so through a certain brain state whereby one normally brings about the rising of one's arm. And that is then a way in which a BIV might manifest the skill of a great pianist, even if she fails to exercise any fuller pianistic competence.

6. *An analogy to the BIV who manifests virtuosity in her perceptual judgments is now attractive. That believer seems "justified" in so believing, even if her belief is radically false. She still believes in an epistemically appropriate way, and the belief is still recognizably skilful.*

¹But the skill exercised need not be the skill to bring about that particular brain state. We need not commit at all on that. The skill exercised may be an ordinary physical skill: to play a certain concerto, say, or to tie a knot.

²The case of the pianist and that of the brain performance reveal a distinction between intentional and ontological "by" relations. Once the pianist knows that in playing a certain phrase she will be hitting a certain key with a certain finger, she can intentionally, by design, bring it about that she does hit that key with that finger. She brings that about intentionally by intentionally playing the relevant phrase. But *ontology* reverses the "by" direction. It is in part by hitting that key with that finger that the pianist ontologically brings about her intentional playing of that phrase. This seems also relevant to the whole "X first" controversy. X could be "first" in ways that do not affect the ontological order of the relevant domain and so do not affect proper *metaphysical* explanation in that domain.

Consider a skill-like competence to perceive that one faces something red and round. A certain visual impact on one's eyes leads through the afferent nerves to a corresponding brain-involving state that grounds its visually seeming to the subject as if he faces something red and round. Brain states in a certain range respond respectively to the relevant range of inputs from the vision-involving afferent nerves. Entering such a state is then something agents do as their own attributable doing, even if this outcome is not *consciously intentional*.

Having entered that afferent state, suppose the agent then *tries* to get it right on the question whether he faces something red and round. The agent might try judgmentally (intentionally) to get it right on that question. Alternatively, the agent might *aim* to get it right just functionally through psychological teleology, with no conscious intention. Let us focus, in any case, on the brain state that constitutes visual experience as if one faces something red and round. On the basis of such visual experience, one can then properly *try* to get it right, which trying might take either a functional (teleological) form or a consciously intentional form. And this *attempt* (functional or judgmental) can then be justified by being skillfully truth-conducive. Such "skillful justification" is hence compatible with failure, since one can exercise one's skill while in poor shape or in a poor situation. Clearly, such skillful justification might reside in a BIV. It is required only that the given perceptual afferent brain state lead competently to the relevant efferent brain state. This is the given afferent/efferent function that, when normally occurrent in a properly encased brain, leads reliably enough to true belief.

7. In our proposal there is thus a way for a BIV to be epistemically justified. What is that way? It is for his belief to constitute a skillful attempt to get it right, one that can fail miserably to be right while still being a highly skilled attempt. This is because the attempt resides in a deed that is normally very likely to bring success but fails to have its expected effect only because of the agent's defective shape or situation.

Here I have worked with a *broad* sense of epistemic justification so as to make room for externalist epistemologies that require no strict deontic sense applicable only to free and voluntary agency.

In our broader understanding, a belief can be epistemically justified even if it is not an exercise of free agency, so that it attains no deontic status. Our acceptance is in line with a conception of *competence* that is, of course, applicable to free attempts but also to *deeds*. A deed is, again, a doing attributable to a subject as *his own doing*, one that may or may not manifest competence, despite falling short of free, active agency. Our perceptual competence, for example, resides largely in perceptual mechanisms whose exercise yields passive functionings rather than free actions. When things seem a certain way perceptually, this is not often—or ever—the direct outcome of any choice or intentional design.

In a narrower sense, only free actions can be justified or unjustified. This, too, can be understood in terms of competence. But now the competence would involve a disposition to succeed when one freely tries intentionally (by design).

My account of skillful justification is meant to be compatible with the broad sense of justification that admits passive functionings within its scope, and compatible also with the narrower sense restricted to a deontic status that applies only to free actions.

So far we have discussed knowledge in part A and justification in part B. Next we turn to skepticism.

C. SKEPTICISM

1. *Aptness and Modal Safety in Epistemology*

Suppose that human beings pursue their daily routines some fine day based on the ordinary knowledge that normally enables us to do so. What has escaped nearly everyone, however, is that a rogue state led by a maniac has a hydrogen bomb sufficient to destroy all life and normality on Earth in extremely short order. And the maniac is about to flip a coin that will determine our fate. Accordingly, at that point people lack knowledge that even their extremely short-term predictions *would* turn out right. Suppose the coin lands favorably, however, so that Planet Earth escapes that dire fate and everything goes on normally, with no sign of trouble, actual or potential. Would that matter to the assessment of human performance within that time span?

It *would*, of course, matter if the relevant quality of one's performance required proper trust in absolute conditionals and risk assessments. Take someone who, at the time of the coin flip, trusts that by operating the starter she would start her car. This is the sort of assumption needed for successful means/end actions. If such assumptions had to be absolute, however, then at the time of the coin flip, human performances across the globe would fall short seriously. They would be based essentially on luck rather than knowledge.

Here I assume that the bomb is so powerful and fast-acting that its blast would obliterate all life and normality instantaneously. So if the coin flip had turned out unfavorably, the causal chain from starter to motor would have been cut short immediately, making the driver's prediction that the car would start a false prediction. Our driver assumes that, if she were to operate the starter, she would thereby start the car without delay. However, given the momentous coin flip, that implicit assumption is false. *Rather*, if she were to operate the starter, her car would start *provided the coin landed favorably*. If, on the contrary, she were to operate the starter *while the coin landed unfavorably*, then *there would be no car start, no car, and no driver*. So at the time of the coin flip, our driver is wrong to assume that if she were to operate the starter, then the car *would* start. Yet full credit for her success is dependent on the epistemic standing of her guiding belief that she would be able to start the car that way. What should we conclude? Because of the coin-flip setup, does our driver fall short of fully creditable action in starting her car? Does she fall short *even when the coin lands favorably*?

Anyone who draws that conclusion is focused on the epistemic shortfall that derives from the coin-flip setup. The epistemic shortfall is thought to affect not just epistemic performances but human performances generally. Performances are very generally based on guiding means/end beliefs and assumptions, and the quality of such a performance will depend on the epistemic quality of those beliefs and assumptions. But what *is* that required epistemic quality?

It seems implausible that human performance would be spoiled across the globe when the maniac's coin flip could so easily have gone the other way. Human performance seems unaffected by any

such element of luck, epistemic performance included.³ So it cannot be that our performance quality at that fateful moment depends on *knowledge* that relevant background conditions do hold.

We properly assume, as a default, that relevant background conditions *will actually hold*. By contrast, we need not rule out the mere *danger* that a background condition *might fail to hold*. Although background conditions must actually hold, they can hold by luck, despite grave danger. Of course, not just *any* kind of luck affects the quality of an attained success, of an achievement. It may be lucky that the achiever retains his skill, is in good shape or well situated, or is even alive. Many essential elements of achievement can depend on luck that does not diminish achievement. A case in point, in the present proposal, is that of background conditions assumed to hold. A background condition might be essential to the success of a performance while holding only by luck and beyond the ken of the performer. But *that particular sort of luck* may have no effect on the quality of that achievement.

2. Generalizing to All Performance Domains

The phenomenon we have just encountered can be found across the spread of performance domains. Consider, for example, athletic and artistic domains. Performers in such domains properly assume by default that relevant background conditions hold as a matter of course.

When an athletic event takes place at night, all present assume that the stadium lights will not fail midperformance, causing the athletes to perform poorly. If in a certain game the lights in fact stay on only by chance because the lighting system is likely to fail, this casts no reflection whatever on the quality of athletic performance in that game. Nighttime athletes do need to presuppose that their SSS competence is adequate for a sufficiently safe attempt. What they need to presuppose is that, *relative to their SSS competence*, their attempt is likely enough to succeed. They need to presuppose that, if they made their attempt and their SSS

³This is the element of luck that is in play when “background conditions” in fact hold by luck.

conditions were still satisfied, then they would succeed. This competence is thus determined by a combination of skill, shape, and situation that is good enough to ensure that the agent's attempt would likely enough succeed.

We properly assume favorable background conditions across domains of human performance. Agents perform based on an assumption, normally implicit, that their relevant SSS competence makes their attempt safe enough. In other words, they assume that, *given their SSS profile*, they will succeed safely enough.

3. *Definition of a Background Condition*

A nighttime athlete has a likelihood of success determined by his skill, shape, and situation. Moreover, his attempts are well selected only if the agent is likely enough to succeed given his SSS profile and given that he assesses relevant risks well enough. The athlete needs to consider various shape and situation factors: how tired he is, for example, how far from the target, and so on, for the many shape and situation factors that can affect performance. But there are many factors that he need *not* heed. It is no concern of *an athlete as such* whether an earthquake might hit, or a flash tornado, or a hydrogen bomb set off by a maniac leader of a rogue state, and so on. As an athlete, he is not negligent for ignoring such factors.

Background conditions come in three sorts, corresponding to the three parts of the SSS structure of competence. So they can concern the situation, the shape, or the skill of the performer. A background condition is a condition that must hold if the relevant S is to be in place at the time of performance. Thus, the presence of the pertinent skill, shape, or situation will entail respective background conditions that must then hold. What makes such a condition a mere *background* condition is that, although it must hold, the performer need not *know* that it will hold. Nor need it hold *safely*. Reducing the safety of that condition would not reduce or in any way affect performance quality.

4. *Radical Skepticism and Relevant Alternatives*

Against the skeptic's use of radical skeptical scenarios, it is often objected that, being so remote, they can safely be assumed to

be false. Our account enables an objection independent of any such appeal to remoteness. As we have seen, inability to rule out even very close danger might have zero bearing on quality of performance.

There is hence an independent and still telling objection to the skeptic: namely, that his possibilities are like the possibility that the lights will go out in a night game. Even present danger that the lights will go out need not affect the quality of performances on the field. *Nor need the players pay any heed to that present danger*, since the mere *danger* is irrelevant to the athletic quality of their performances. As we have seen, all they need is a *correct* assumption that the danger will not in fact be realized. Analogously, then, epistemic players may have no need to rule out the irrelevant alternatives posited by the skeptic.

Consider the BIV skeptical scenario. If that scenario were actual, that *would*, of course, affect our epistemic performances, just as the lights going out would affect the performance of players in a night game. But the *merely modal* fact that some such scenario is a close and present danger would not affect the quality of our epistemic performances, just as the merely modal fact that the lights might go out affects not at all the quality of athletic performance on the lighted field. Nor need such a skeptical scenario be ruled out *knowledgeably*, just as the lights' going out need not be ruled out knowledgeably by athletes performing at night.

Accordingly, perhaps we need not rule out radical dangers in order to sustain the quality of our epistemic performances. Our response to this skeptic is that he has mistaken what is required for the epistemic quality of our ordinary judgments and beliefs. His mistake is like disparaging a superb baseball catch as due essentially to credit-reducing luck because, unbeknownst to the fielder, the lights could so easily have failed. After all, no one in our imagined stadium could rule out such a failure *knowledgeably*: by hypothesis, the lighting system is extremely likely to fail, *unbeknownst to anyone*.

5. A Response by the Skeptic

Granted, there is a potentially important difference between fully apt athletic performance by athletes performing at night and fully

apt epistemic performance by an ordinary believer. Nighttime athletes might still be thought *not* to be threatened by the iffiness of the required background conditions, even when these cannot be known to remain probably enough in place. *Less* plausible, however, is the thought that epistemic players can similarly perform perfectly well despite a similar *epistemic* defect concerning *their* required background conditions. Can we really know about hands or fires if we cannot rule out that we are victims of evil demons or envatters? Surely we cannot just ignore the epistemic standing of our underlying common-sense framework of beliefs.

That does seem plausible, granted. Fortunately, we do enjoy full epistemic standing for trust in our Moorean common sense.⁴ For one thing, the cases of iffy background conditions we have been considering are quite different from the familiar skeptical scenarios: radical skeptics do not suggest that our common-sense framework is actually unsafe. In any case, I take that to be a separate issue to be tackled separately, as I have tried to do in separate writings.

Even bracketing that, however, we still have a considerable response to the radical skeptic by pressing the analogy between the following two things:

- (a) how the quality of a great catch under the lights is not affected by any need for the fielder to trust *knowledgeably* that the pertinent background conditions were likely enough to hold, and
- (b) how the quality of a brilliant belief by Sherlock, or by a lookout at the top of a mast, would be similarly unaffected by any similar need to trust *knowledgeably* that *his* pertinent background conditions were likely enough to hold.⁵

⁴Or so I have argued at length in *Reflective Knowledge* (Oxford, UK: Oxford University Press, 2009).

⁵Our approach suggests a promising way to understand the appeal to “relevant alternatives” seen in responses to skepticism of years ago. At that time the notion of relevance was left in some darkness, whereas the present approach may now shine its light. Our *background conditions*, and their generalization to human performance generally, seem also interestingly related to Wittgenstein’s *hinge propositions*, though this must be left here as a topic for later study. For a discussion of the relevant alternatives approach to skepticism and to the analysis of

This is the analogy that is made plausible by examples like that of the maniac with the bomb and the coin flip.

True enough, there is an important difference. Again, we *can* trust Moorean common-sense beliefs knowledgeably, whereas our baseball spectator cannot trust knowledgeably that the lights will stay on. But we now have another way to resist the radical skeptic anyhow; we can oppose him on the alternative basis that presses our analogy between 5a and 5b.

6. Knowledge Closure versus Justification Closure

In rejecting knowledge transmission and even knowledge *closure*, however, we make no such claim about justification. Our stance concerns *knowledge* specifically. It has no immediate bearing on the closure or transmission of *justification*.⁶ Suppose that something defeats the spectator's *justification* for believing that the lights will stay on. This will surely impact his justification for believing that the fielder is likely to succeed and indeed *will* succeed (and therefore impact also his knowledge, if any, to that effect). So

knowledge, see section 5.3 of the article on the analysis of knowledge by Matthias Steup and Jonathan Ichikawa in the *Stanford Encyclopedia of Philosophy (SEP)*, available at <http://plato.stanford.edu/contents.html>. (References may be found there to the relevant work of Stine, Goldman, and Dretske.) Our present approach also entails a crucial distinction between safety and relevance. The rapprochement between those two approaches suggested in the *SEP* entry is in conflict with our approach. For us, as we have seen, safety is not required for knowledge or for apt performance generally. A performance of a nighttime athlete can be apt even if extremely unsafe (because of the unsafe standing of the lighting system). And a similar conclusion may be drawn about epistemic performances, in which again great danger, great risk of failure, seems quite compatible with success that is perfectly apt. Only the lack of safety involving "relevant" alternatives can deny us knowledge.

⁶Issues of transmission versus closure and of default justification are featured in writings of Crispin Wright. See, for example, "Warrant for Nothing (and Foundations for Free?)," *Proceedings of the Aristotelian Society* 78 (2004 suppl.): 167–212. And there is an extensive literature on the importance of transmission in epistemology, and of distinguishing it from closure. See the *Stanford Encyclopedia* entry, "Transmission of Justification and Warrant," by Luca Moretti and Tommaso Piazza, and the *Internet Encyclopedia* entry by Chris Tucker on "Transmission and Transmission Failure in Epistemology."

we can accept justification closure even while questioning knowledge closure.

If justification is closed under competent deduction, however, this reveals a limit of our alternative defense against the radical skeptic. For we still have to tackle the traditional skeptical problematic based precisely on the proposition that we lack justification for rejecting the skeptical scenario. After all, given justification closure through competent deduction, this would make it very hard to see how we could possibly be so much as justified in our ordinary beliefs about hands and fires. Accordingly, more work would be needed in order to resist the radical skeptic's justification-focused reasoning. That is a job taken on in earlier chapters.⁷

FURTHER READING

Dutant, Julien, and Clayton Littlejohn. *The New Evil Demon Problem*. Oxford: Oxford University Press, 2016.

Kvanvig, J., and C. Menzel. "The Basic Notion of Justification." *Philosophical Studies* 59 (1990): 235–61.

Littlejohn, Clayton. "The New Evil Demon Problem." *Internet Encyclopedia of Philosophy*. Available at <http://www.iep.utm.edu/evil-new/>.

Poston, Ted. "Internalism and Externalism." *Internet Encyclopedia of Philosophy*. Available at <http://www.iep.utm.edu/int-ext/>.

Further Reading Related to This Book

This list includes readings on the topics, broadly interpreted, taken up in this book.

Alfano, Mark. *Character as Moral Fiction*. Cambridge, UK: Cambridge University Press, 2013.

———. *Current Controversies in Virtue Theory*. New York: Routledge, 2015.

Alfano, Mark, and Abrol Fairweather, eds. *Epistemic Situationism*. Oxford, UK: Oxford University Press, 2015.

Alston, William P. *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca, NY: Cornell University Press, 1989.

Anscombe, G.E.M. *Intention*. 2nd ed. Oxford: Blackwell, 1963.

Appiah, Kwame Anthony. *Experiments in Ethics*. Cambridge, MA: Harvard University Press, 2008.

Audi, Robert. *The Structure of Justification*. New York: Cambridge University Press, 1993.

Baehr, Jason S. *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. New York: Oxford University Press, 2011.

Becker, Kelly. *Epistemology Modalized*. New York: Routledge, 2007.

Becker, Kelly, and Tim Black, eds. *The Sensitivity Principle in Epistemology*. Cambridge, UK: Cambridge University Press, 2012.

Bonjour, Laurence. *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press, 1985.

Bonjour, Laurence, and Ernest Sosa. *Epistemic Justification*. Oxford, UK: Blackwell, 2003.

Brandt, Robert. *Reason in Philosophy*. Cambridge, UK: Harvard University Press, 2009.

Conee, Earl, and Richard Feldman. *Evidentialism: Essays in Epistemology*. Oxford, UK: Oxford University Press, 2004.

Craig, Edward. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford, UK: Clarendon, 1990.

Davidson, Donald. *Essays on Actions and Events*. Oxford, UK: Oxford University Press, 2001.

DePaul, Michael, and Linda Zagzebski, eds. *Intellectual Virtue: Perspectives from Ethics and Epistemology*, 135–54. Oxford, UK: Oxford University Press.

Further Reading

- DeRose, Keith. *The Case for Contextualism*, volume I. Oxford, UK: Oxford University Press, 2009.
- Doris, John. *Lack of Character*. New York: Cambridge University Press, 2002.
- Dutant, Julien, and Clayton Littlejohn. *The New Evil Demon Problem*. Oxford, UK: Oxford University Press, 2016.
- Fairweather, Abrol, and Owen Flanagan, eds., *Naturalizing Epistemic*. Cambridge, UK: Cambridge University Press, 2014.
- Fairweather, Abrol, and Carlos Montemayor. *Epistemic Dexterity*. Cambridge, UK: Cambridge University Press, 2017.
- Fantl, Jeremy, and Matthew McGrath. *Knowledge in an Uncertain World*. Oxford, UK: Oxford University Press, 2009.
- Fernandez, Miguel Angel. *Performance Epistemology: Foundations and Applications*. Oxford, UK: Oxford University Press, 2016.
- Fogelin, Robert. *Pyrrhonian Reflections on Knowledge and Justification*. Oxford, UK: Oxford University Press, 1994.
- Foley, Richard. *When Is True Belief Knowledge?* Princeton, NJ: Princeton University Press, 2012.
- Fumerton, Richard. *Metaepistemology and Skepticism*. Lanham, MD: Rowman and Littlefield, 1995.
- Goldberg, Sanford. *Relying on Others*. Oxford, UK: Oxford University Press, 2010.
- Goldman, Alvin. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press, 1986.
- Greco, John. *Achieving Knowledge*. Cambridge, UK: Cambridge University Press, 2010.
- , ed., *Ernest Sosa and His Critics*. Oxford, UK: Blackwell, 2004.
- Harman, Gilbert. *Thought*. Princeton, NJ: Princeton University Press, 1973.
- Hawthorne, John. *Knowledge and Lotteries*. Oxford, UK: Oxford University Press, 2004.
- Huemer, Michael. *Skepticism and the Veil of Perception*. Lanham, MD: Rowman and Littlefield, 2001.
- Hyman, John. *Action, Knowledge, and Will*. Oxford, UK: Oxford University Press, 2015.
- Klein, Peter. *Certainty*. Minneapolis: University of Minnesota Press, 1981.
- Kornblith, Hilary. *On Reflection*. Oxford, UK: Oxford University Press, 2012.
- Kvanvig, Jonathan. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge, UK: Cambridge University Press, 2003.

- Lackey, Jennifer. *Learning from Words*. Oxford, UK: Oxford University Press, 2008.
- Littlejohn, Clayton. *Justification and the Truth Connection*. Cambridge, UK: Cambridge University Press, 2012.
- Luper-Foy, Steven, ed. *The Possibility of Knowledge: Nozick and His Critics*. Totowa, NJ: Rowman and Littlefield, 1987.
- Lyons, Jack. *Perception and Basic Beliefs*. New York: Oxford University Press, 2009.
- McCain, Kevin. *Evidentialism and Epistemic Justification*. New York: Routledge, 2014.
- McDowell, John. *Mind and World*. Cambridge, UK: Harvard University Press, 1994.
- Montmarquet, James. *Epistemic Virtue and Doxastic Responsibility*. Lanham, MD: Rowman and Littlefield, 1993.
- Moore, G. E. *Philosophical Papers*. London: Allen and Unwin, 1959.
- Morton, Adam. *Bounded Thinking: Intellectual Virtues for Limited Agents*. Oxford: Oxford University Press, 2013.
- Moser, Paul, ed. *The Oxford Handbook of Epistemology*. New York: Oxford University Press, 2005.
- Nozick, Robert. *Philosophical Explanations*. Cambridge, MA: Harvard University Press, 1981.
- Plantinga, Alvin. *Warrant and Proper Function*. New York: Oxford University Press, 1993.
- Poston, Ted. *Reason and Explanation*. Houndmills, UK: Palgrave Macmillan, 2014.
- Pritchard, Duncan. *Epistemic Angst*. Princeton, NJ: Princeton University Press, 2015.
- . *Epistemic Disjunctivism*. Oxford, UK: Oxford University Press, 2012.
- . *Epistemic Luck*. Oxford, UK: Oxford University Press, 2005.
- Reid, Thomas. *Essays on the Intellectual Powers of Man*. In *Reid: Inquiry and Essays*, ed. R. E. Beanblossom and K. Lehrer. Indianapolis, IN: Hackett Publishing, 1983.
- Ross, Lee, and Richard Nisbett. *The Person and the Situation: Perspectives of Social Psychology*. London: Pinter and Martin, 2011.
- Roush, Sherry. *Tracking Truth*. Oxford, UK: Oxford University Press, 2005.
- Russell, Daniel. *Practical Intelligence and the Virtues*. Oxford, UK: Oxford University Press, 2009.
- Sinnott-Armstrong, Walter, ed. *Pyrrhonian Skepticism*. Oxford, UK: Oxford University Press, 2004.

Further Reading

- Sosa, Ernest. *Judgment and Agency*. Oxford, UK: Oxford University Press, 2015.
- Stanley, Jason. *Knowledge and Practical Interests*. Oxford, UK: Clarendon, 2005.
- Stroud, Barry. *The Significance of Philosophical Scepticism*. Oxford, UK: Oxford University Press, 1984.
- Turri, John, ed. *Virtuous Thoughts: The Philosophy of Ernest Sosa*. Dordrecht: Springer, 2013.
- Unger, Peter. *Ignorance: A Case for Scepticism*. Oxford, UK: Oxford University Press, 1975.
- Williams, Michael. *Unnatural Doubts*. Princeton, NJ: Princeton University Press, 1996.
- Williamson, Timothy. *Knowledge and Its Limits*. Oxford, UK: Oxford University Press, 2000.
- . *The Philosophy of Philosophy*. Oxford, UK: Blackwell Publishers, 2007.
- Zagzebski, Linda. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. New York: Cambridge University Press, 1996.
- Zalabardo, José. *Scepticism and Reliable Belief*. Oxford, UK: Oxford University Press, 2012.

Index

- accuracy, adroitness, and aptness (AAA) structure, 72, 112, 146, 209–11
- action, 69, 71–86, 125–31; and accuracy, adroitness, and aptness (AAA structure), 72, 112, 146, 209–11; attempts vs. achievements, 71–72; causal accounts of action, 125–31; deeds, 212–13, 215; judgment as, 71, 130; knowledge as, 71–86, 207, 211; negligent actions, 103–4 (*see also* negligence); and risky attempts, 77–78; and virtue theory vs. situationism, 171–90; and waking state, 28. *See also* affirmation; aptness; attempts; competence; intention; performance
- affirmation: alethic affirmation, 82, 84–86, 89–90, 100, 152n11, 153nn12,13, 155, 208–9; aptness of, 82, 89–93; and epistemic normative hierarchy, 100; forms of, 155–56; and intention, 83, 152–53; and judgment, 92–94, 208; public assertion vs. private affirmation, 153n12; quasi-affirmation, 35–36, 38; and safety, 89, 90; skepticism and the refusal to affirm, 155n14; and voluntary/involuntary beliefs, 83–84
- agency: and competence virtue epistemology, 140–45; conscious epistemic agency, 155; and deeds, 212–13, 215; and epistemic justification, 214–15; and motivation and responsibilist virtues, 148–50; and reflection, 100–101, 155–56
- aiming, 193–95; action, perception, and knowledge constituted by aimings, 130; beliefs as intellectual shots, 109–13; and competences and their manifestations, 129; and guessing, 113; shot at the beach, 108–10; and true aim, 108–9. *See also* archery; baseball; basketball; soccer; tennis
- airplane (thought experiment), 167–68
- alethic affirmation, 82, 84–86, 89–90, 100, 152n11, 153nn12,13, 155, 208–9
- animal beliefs (implicit or involuntary beliefs), 12–13, 83–84, 102, 151–53, 208
- animal knowledge, 45–46, 91, 141; and bilevel epistemology, 45–46; distinguished from reflective knowledge, 18, 103; endorsement of, 154; and importance of reflective knowledge, 95–100
- aptness, 207–8; of affirmation, 82, 89–93; and archery shots, 72, 112, 211; Aristotle and, 207–8; and avoiding error, 18; of awareness, 96–100; and basketball shots, 73–78, 99, 199; and Cartesian epistemology, 18, 40, 94, 208–9; causation of, 129–30; and epistemic normative hierarchy, 100; full aptness in epistemology, 79–83, 89–92; of judgment, 40, 82, 89–90, 197; knowledge as apt belief, 126, 129; and morality, 146; and perception, 129; of performance, 71–78, 96–99, 125, 129, 198–99, 207, 210, 219–20; and safety, 210, 215–17; successful, competent, and apt attempts, 73; and threshold of reliability, 73–78; and value of performances, 112; and value of truth, 115. *See also* competence

Index

- archery: and accuracy, adroitness, and aptness (AAA structure), 72, 112, 210, 211; and brain state, 211–12; and competence, 192, 201; success by neophyte archer, 113n3; and value of performance, 111–12; wind-aided shot, 125, 201–2, 210
- Aristotle, 117, 118, 120, 140, 171n1, 178, 207–8
- arithmetic, 9–10, 43–45, 53
- assassin (thought experiment), 150–51
- athletic performances, 78–79n6, 111–12, 134, 193, 204; nighttime athletes, 217–20, 221n5. *See also* archery; basketball; soccer; tennis
- attempts, 71–72, 209; and epistemic normative hierarchy, 100; platitudes about, 114; risky attempts, 77–78, 97; safety and aptness of, 210; and skillful justification, 214; successful, competent, and apt attempts, 73; trying vs. wishing or hoping, 212
- Augustine, St., 29n11
- Baehr, Jason S., 148n9, 151n10
- bank account (thought experiment), 165–67
- barn façade (thought experiment), 55, 87–90, 99
- baseball (nighttime athletics), 217–20, 221n5
- basketball, 73–78, 97–99, 198–99, 204
- behavior. *See* situationism; virtue ethics
- beliefs: as acts of judgment, 6–7, 156 (*see also* judgmental belief *under this heading*); and agency, 100–101; animal beliefs vs. conscious reflective beliefs (voluntary vs. involuntary, implicit vs. intentional), 12–13, 83–84, 102, 151–53, 208; based on perception, 5–6; and begging the question, 46; beliefs that are false despite being competent or justified, 64–65, 65–66n13, 69, 107, 213 (*see also* Gettier problem); beliefs that are true by accident, 15, 124–25, 208–9; and brain in a vat, 214–15; conditions for beliefs to constitute knowledge, 22–24, 56–62, 106–7, 124–25; and credence vs. judgment, 6–9; and Descartes's virtue epistemology, 1–20; and diachronic reliability, 102, 164–65; and dogmatism, 157–60, 165–69; as domain of epistemic performance, 73, 110–13; and dream state, 26–31, 37–38; endorsement of, 4–6, 16–17, 102–3, 154–56, 161 (*see also* affirmation); and epistemic competence, 106–7, 122, 163, 203–4; and epistemic justification, 214–15; epistemic status of beliefs, 162–63; implicit awareness of beliefs, 17–18; as intellectual shots, 109–13; judgmental belief, 13, 84, 152–53, 156; justified true belief, 62–70, 87n; knowledge as superior to true belief, 108, 113–16, 145; make-belief, 6–7, 153; platitudes about, 114; rejecting beliefs/suspending judgment, 3–6, 22n2; and responsibilism, 151–53; safety and sensitivity requirements, 22–24, 56–62; skillful belief, 212; and status of targets determining value of shots aimed at target, 109; and synchronism, 160–65; true beliefs falling short of knowledge, 55–57, 87–92, 107, 124–25; truth and knowledge as aim of, 108–10. *See also* *cogito* beliefs; doubt; justification; skepticism
- blindsighters, 80, 81n9, 88, 91, 102
- bootstrapping. *See* circularity
- brain in a vat (thought experiment), 23, 24, 57–58, 60, 211, 214–15, 219–20
- Buckwalter, Wesley, 87n2, 89n2

- carbon copies, signing (thought experiment), 212
- Carroll, Lewis, 63n7
- Cartesian epistemology: Cartesian error, 14–18; and certainty, 4–6, 11–12, 16, 18–19, 40, 42; and circularity, 18–19; and credence vs. judgment, 6–9; distinction between *cognitio* and *scientia*, 18; and doubt, 1–9, 16; and endorsement of beliefs, 4–6, 16–17; and epistemic competence, 17; project described, 17–18
- cat-on-the-mat, analysis of statement, 120–21, 128
- causation, 123–37
- certainty, 4–6, 11–12, 15–19, 40, 42, 52, 94. *See also* dogmatism
- character traits, 141–47, 171–83
- chicken sexers, 80, 81n9, 88, 91
- circularity, 18–19, 39–49, 103
- clairvoyants, 80–81n9, 88–91
- clarity and distinctness of Cartesian perceptions/seemings, 7, 8, 15, 18, 40, 52
- close-mindedness. *See* dogmatism
- cogito* beliefs, 9–10, 32–34, 43n7, 44–45, 52–53, 85, 94. *See also* thinking
- cognitio*, 9, 18–19, 41, 43n7. *See also* animal beliefs; animal knowledge
- coin-flipping maniac (thought experiment), 215–17, 221
- Colaco, David, 87n2
- common sense, 19, 57, 58–61, 132–34, 135n14, 180
- competence, 146, 167n5, 191–206, 195n2; and archery shots, 72, 112, 125, 192, 201; and basketball shots, 97–98; and beliefs, 124–25; beliefs that are false despite being competent or justified, 64–65, 65–66n13, 69, 107, 213 (*see also* Gettier problem); competent deduction, 57, 58, 222; and degrees of reliability, 192; and disposition to succeed if one tries, 191, 193–97; and driving, 131–32, 136, 184–90, 191, 196–97; epistemic competence, 17, 106–7, 122, 163, 203–5; and epistemic normative hierarchy, 100; first- and second-order competence, 77, 93–95; fundamental vs. non-fundamental competences, 49–51; and higher-order phenomena, 93–95, 197–99; and judgment, 93–94; manifestations, 128–30, 132–37; and memory, 163; and morality, 184–88; and negligence, 95, 96n11, 104; quasi-competence, 193n1; reflective competence, 94–95; and responses to situationist critique of virtue ethics, 184–90; and safety and reliability, 201–3; and search for a legitimizing account of knowledge in general, 48; and situationist critique of virtue epistemology, 203–5; skill/shape/situation (SSS) conditions, 77, 131–36, 146, 167–68, 191–203, 217–18; as special case of a disposition to succeed, 191, 193–97, 199, 209–10; successful, competent, and apt attempts, 73; and synchronism, 160–65; and “the right way,” 123, 127–30; virtue epistemology and a theory of competence, 191–206. *See also* aptness
- competence virtue epistemology, 140–45, 156
- conceptual analysis, 120–22, 127–28, 137
- conceptual innovation, 137–38
- Confessions* (Augustine), 29n11
- confidence, 81–82, 161, 198. *See also* certainty
- conscientiousness, 145–46, 160
- consciousness, 25, 101–2. *See also* reflective knowledge; waking state
- credence, 6–9, 16, 17, 19

Index

- cubical box (thought experiment), 66n14
- curiosity, 109–10
- Davidson, Donald, 123, 125, 126, 128–29, 212
- deduction, 8, 21, 41, 42, 54, 57, 58, 69, 103
- deeds, 212–13, 215
- Descartes, René: apple-basket metaphor, 2–4; and apt belief or judgment, 40, 208–9; Cartesian error, 14–17; and circularity, 43–46, 49; and *cognitio* and *scientia*, 40 (see also *cognitio*; *scientia*); and daily life guided by resultant seemings, 8–9; and dream scenario, 11, 24–25, 31–34, 43n7; and external world, 41; and judgment, 152, 155; and mathematics, 43–45, 52–53; and the method of doubt, 1–9; principle of clarity and distinctness, 7, 8, 15, 18, 40, 42, 52; and Pyrrhonian virtue epistemology, 1–20; and reliabilism, 40–43. See also Cartesian epistemology; *cogito* beliefs
- dispositions, 132–36, 191–97, 199, 204–5, 209–10. See also character traits
- dogmatism, 157–60, 165–69
- Doris, John, 174–78, 180, 185
- doubt, 1–9, 16, 31
- dream state, 24–38; beliefs and intentions in, 26–31; Descartes and, 11, 24–25, 31–34, 43n7; dream-shadow believing, 37–38; and personal responsibility, 29–30; and phenomenal experiences, 30–31; and quasi-affirmation, 35–36, 38; and skepticism, 31–34; and thinking, 32–34
- driving, 103–4, 131–32, 136, 184–91, 196–97
- dumbbell, shattered (thought experiment), 199
- Earman, John, 21n1
- Empiricism, 101
- envatted brain experiment. See brain in a vat
- epistemic circularity. See circularity
- epistemology: bilevel epistemology, 40, 45–46; and conceptual innovation, 137–38; domains of epistemic performance, 73, 110–13; epistemic competence, 17, 106–7, 122, 203–5; epistemic justification, 106–7, 210–15; epistemic negligence, 103–4, 165–69; epistemology of the reflective, 92–95; full aptness in, 79–83, 89–92; and natural externalism, 39; search for a legitimatizing account of knowledge in general, 46–49; and semantics, concepts, and extralinguistic/extracognitive entities, 122, 137; and transmission, 221n6. See also Cartesian epistemology; virtue epistemology
- error, 14–15, 18, 208–10. See also falsity
- ethics. See virtue ethics
- evil demon (thought experiment), 11, 24, 57–58, 220
- “failed threat” cases, 89n2
- falsity: beliefs that are false despite being competent or justified, 64–65, 65–66n13, 69, 107, 213 (see also Gettier problem); and circularity, 43–44; and dream state, 11, 25, 29n11, 35–36; falsity as sufficient but not necessary for error, 14–15, 18, 208; and method of doubt, 1–4, 6, 11; and negligence, 166; and perceptual beliefs, 26, 39n1, 40; plausible false things, 60; and sensitivity and safety requirements for belief to constitute knowledge, 22–23, 56, 59, 61–62. See also error; illusions
- “firsters,” 127–28

- first-order phenomena, 14–18; first-degree reflection, 101; first-order competence, 77, 163, 167–68; first-order judgments and error, 18; first-order performance, 96–99, 98nn13, 14; first-order safety, 76. *See also* animal beliefs; animal knowledge
- freedom of indifference, 83
- free will, 83–84, 100–101, 209
- Fumerton, Richard, 46n11
- fundamental attribution error, 175, 185, 190n19
- geometry, 9–10, 43–45, 53
- Gettier, Edmund, 62, 87n1
- Gettier problem, 62–69, 80, 87–92, 96, 99, 106, 125
- gold-in-the-dark (thought experiment), 98n14
- Goldman, Alvin, 87n2
- Good Samaritan experiment, 173–74, 182n16
- Grice, Paul, 124, 125, 126, 128
- Grundmann, Thomas, 65n12
- guessing, 79, 92, 113
- Harman, Gilbert, 159, 174–78, 180
- honesty, 175, 177–78, 180, 181
- Horvath, Joachim, 65n12
- human flourishing, 115–18, 134, 150
- Hume, David, 51
- hydrogen bomb (thought experiment), 215–17, 221
- Ichikawa, Jonathan Jenkins, 56n3
- illusions: cognitive illusions, 45, 60, 62; illusory clash between virtue ethics and situationism, 185, 188–89; perceptual illusions, 9, 53n13, 60, 156, 204. *See also* dream state
- imagination, 35, 37–38
- inattentional blindness, 182n16
- intellectual virtues, 140–46, 156
- intention, 123–24; affirmations as intentional acts, 83; and causal accounts of action, 125–31; and deeds, 212–13, 215; and dream state, 26–31; intention to close one's mind, 157–59; judgment and knowledge as intentional action, 71 (*see also* action); waiter knocking over dishes example, 123, 128–29
- intuition, 34n15, 49, 62–70, 103; and blindsighters, chicken sexers, 91–92; and conceptual innovation, 138; Descartes and, 7, 41, 42, 49; epistemic structure of rational intuition, 34n15; and Gettier problem, 62–63, 65–66, 68, 99; and negligence, 167n5; real vs. pseudo-intuitions, 41n4; and reliabilism, 42–43, 45
- judgment, 96n11; affirming, denying, or suspending judgment, 93, 197–98 (*see also* suspending judgment *under this heading*); and agency, 100–101; and alethic affirmation, 153nn12, 13, 208; aptness of, 82, 89–90, 93, 197; and competence, 93–94; vs. credence, 6–9, 19; distinguished from guessing, 92; and dogmatism, 159–60; as domain of epistemic performance, 73; and error, 14–17; first-order judgment, 93–94; as form of action, 71, 130; and higher-order phenomena, 92–96, 101, 163–64, 197–98; judgmental belief, 13, 84, 152–53, 156, 208; and reflective knowledge, 95; safety of, 90; suspending judgment, 3–6, 93, 96n11, 101, 154; and synchronism, 163–64; and truth, 15–16
- justification, 210–15; and alethic affirmation, 84–86; and dogmatism, 159–60; epistemic justification, 106–7, 210–15; and evidence, 51–52; extrinsic vs. intrinsic reasons, 84; justified true belief, 62–70,

Index

- justification (*cont.*)
87n; knowledge closure vs. justification closure, 221–22; and search for a legitimating account of knowledge in general, 48–49; skillful justification, 214; and synchronism, 160–65
- kindness, 171, 180, 181
- Knobe, Josh, 87n2
- knowledge: as apt belief, 126, 129; and bivel epistemology, 40, 45–46; causal accounts of, 124–31; and common sense, 57; and competence virtue epistemology, 141–45; and competent deduction, 58; conditions for beliefs to constitute knowledge, 22–24, 56–62, 106–7; constitution of (definition problem), 55–70, 106–8, 141–42, 147; and dogmatism, 157–60, 165–69; epistemic normative hierarchy, 100; failure to know full well, 87–92; as form of action, 71–86, 207, 211 (*see also* action); and Gettier problem/cases, 62–69, 87–92, 106, 107; and guessing, 79–82; and higher-order phenomena, 92–100 (*see also* second-order phenomena); and justification, 210–15 (*see also* justification); and justified true belief, 62–70, 87n (*see also* Gettier problem); knowledge closure vs. justification closure, 221–22; metaphysics of, 62–69; platitudes about, 114; safety and sensitivity requirements, 22–24, 56–62; as special case of performance, 111–13; as superior to true belief, 108, 113–16, 145; and time, 157–70; value of, 69–70, 106–19, 141, 147; varieties and levels of, 87–105 (*see also* animal knowledge; reflective knowledge)
- Kripke, Saul, 104, 160–65
- Kuhn, Thomas, 21n1
- language, 66, 120–22
- Larissa, road to, 107, 147n5
- light switch, flipping (thought experiment), 85–86
- lottery tickets (thought experiment), 57
- luck, 56n2, 77, 98n14, 113n3, 115, 122, 126, 130, 199
- Lukes, Steven, 171, 180n14
- Machery, Edouard, 87n2
- manifestations of competence, 128–30, 132–37
- mathematics, 8–10, 43–45, 52–53
- Matrix, the, 24, 117, 119n9
- Meditations on First Philosophy* (Descartes), 1–20, 40, 43, 94, 101, 209.
See also cogito beliefs
- memory, 17, 51, 80n7, 102, 161–63
- Meno* (Plato), 106, 141, 147n5
- metaphysical analysis, 67, 92n9, 120–21
- metaphysics of knowledge, 62–69
- metaphysics of persons, 66–67, 120–21, 137
- Milgram, Stanley, 172–73
- mind-world relations, 100, 120–39;
causal accounts of action, 123, 125–31; causal accounts of knowledge, 124–31; causal accounts of perception, 124–31; and words, concepts, and extralinguistic/extracausal entities, 120–22. *See also* action; knowledge; perception
- Moore, G. E., 61
- morality, 104, 114–15, 140, 145–46, 184–88, 200. *See also* virtue ethics
- motivation, 111, 148, 150–54
- Müller-Lyer illusion, 9, 53n13, 156
- mystery box (thought experiment), 143, 145
- negligence, 95, 96n11, 103–4, 114n6, 145–46, 160, 165–69
- Nicomachean Ethics* (Aristotle), 207–8
- nighttime athletes, 217–20, 221n5
- Nisbett, Richard E., 175

- Objections and Replies* (Descartes), 19
 obsessiveness, 143, 144
 open-mindedness, 141, 142, 145–46, 169
- perception: as apt perceptual experience, 129; beliefs based on, 5–6; causal accounts of, 124–31; and circularity, 48, 49; clarity and distinctness of Cartesian perceptions, 7, 8, 15, 18, 40, 42, 52; and dream scenario, 25, 31–32, 34; illusions, 9, 53n13, 60, 156, 204; independent checks on, 49–50; and manifestations of competence, 130; seemings, 7–9, 154; Zagzebski on, 148
- performance: and accuracy, adroitness, and aptness (AAA structure), 72, 112, 146, 209–11; aptness of, 71–78, 96–99, 125, 129, 198–99, 207, 210, 219–20; attempts and achievements as, 71–72; and background conditions, 217–20; beliefs as intellectual shots, 109–13; and coin-flipping maniac example, 216–17; knowledge as special case of performance, 111–13; and negligence, 103–4; professional performance (medicine, law), 77, 134, 145, 150; and risky attempts, 77–78, 97; value of, 111–12. *See also* aiming; competence
- personality. *See* character traits
- petitio*, 46
- piano playing (thought experiment), 213, 213n2
- piloting a plane in bad weather (thought experiment), 167–68
- Plato, 106, 141
- plausibility, 60, 62. *See also* Gettier problem
- pleasure, 115–18
- politeness, 131, 200
- psychology, 66, 122; virtue psychology, 171–73, 180–81, 186–87 (*see also* virtue ethics)
- Pyrrhonism, 1–20, 22n2, 40, 152, 154–56
- Pythagorean theorem, 8–9
- quasi-affirmation, 35–36
- reason: and apt performance, 99; armchair reason, 45, 50–51, 103; and daily life guided by resultant seemings, 8–9; epistemic rationality, 35; knowledge and competent deduction, 57–58; models of inductive reasoning, 21n1; and project of the *Meditations*, 10–11; and reflective knowledge, 45–46. *See also* circularity; common sense; deduction; intuition
- reductio*, 58
- reflection, 101–3
- reflective knowledge, 92–100, 153–54; and agency, 155–56; and bilevel epistemology, 45–46; and circularity, 18, 103; distinguished from animal knowledge, 18, 103; and endorsement of animal beliefs, 154; epistemology of the reflective, 92–95; as important but subsidiary, 95–100; reflective competence, 94–95; and the subconscious, 102–3; and synchronism, 160–65
- regress skepticism. *See* circularity
- reliabilism, 40–43, 41n4, 45, 141, 149–54, 203. *See also* competence virtue epistemology
- reliability, 73–78, 81, 97–99, 192, 198–99, 201–4
- Replies to the Second Set of Objections* (Descartes), 44
- responsibility, 29–30, 142, 145–54, 203
- risk, 77–78, 81, 97, 101, 102, 145–46
- road to Larissa, 107, 147n5
- Ross, Lee, 175
- Russell hypothesis, 51
- safety, 23–24; of affirmation, 89, 90; and aptness, 210; aptness and modal

Index

- safety (*cont.*)
safety in epistemology, 215–17; of belief that we are awake, 24; and competence, 201–3; distinguished from relevance, 221n5; first-order and second-order safety, 76; of judgment, 90; safety requirement for belief to constitute knowledge, 23–24, 57–62
- Schwitzgebel, Eric, 37n18
- scientia*, 18–19, 40, 43n7, 102
- second-order phenomena, 19; and beliefs, 102; and Descartes's method of doubt, 4; and endorsement of beliefs, 17–18, 102, 103; and epistemology of the reflective, 92–100; and importance of reflective knowledge, 95–100; and judgment, 92–96, 96n11, 101, 163–64, 197–98; and negligence, 104; and risk assessment, 101; second-degree reflection, 101; second-order awareness, 98; second-order safety, 76
- seemings, 7–9, 154
- semantic analysis, 62, 67, 120, 122, 127–28, 137
- sensitivity requirement for belief to constitute knowledge, 22–24, 56–62
- situationism, 171–90; central theses, 176–79; experiments, 172–74, 176, 182n16, 186–87, 189–90n19; illusory clash with virtue ethics, 188–89; response to situationist critique of virtue ethics, 184–90; and theory of competence, 191, 203–5; and two-ply behavioral explanations, 185, 188–89
- skepticism, 21–38, 218–22; academic skepticism, 21–22; and aptness and modal safety, 215–17; and brain in a vat, 57, 219–20 (*see also* brain in a vat); and circularity, 39–49; and *cogito* beliefs, 9–10; and Descartes's virtue epistemology, 1–20; and dream scenario, 24–38; global skepticism, 22; and the method of doubt, 1–9; radical skepticism and relevant alternatives, 218–22; and reflective knowledge, 103; and refusal to affirm, 155n14; and sensitivity and safety requirements for belief to constitute knowledge, 22–24; underdetermination-based arguments, 51. *See also* Pyrrhonism
- skill/shape/situation (SSS) conditions for competence, 77, 131–36, 146, 167–68, 191–203, 217–18, 218
- sleepwalking, 28n10
- soccer, 193
- Socrates, 24n5
- Stanford Prison Experiment, 189–90n19
- Steup, Matthias, 56n3
- Stich, Stephen, 87n2
- Stoicism, 118
- subconscious, 12–13, 101–3
- success, 209; accidental success, 193–94, 204–5; and competence, 191, 193–97, 199, 209–10; and epistemic normative hierarchy, 100; and full aptness, 71–78, 125, 129; platitudes about, 114; and value of performance, 111–12
- synchronism, 160–65
- targets. *See* aiming
- tennis, 112–13, 193–95, 204–5
- Theaetetus* (Plato), 106, 141
- theology, 11, 19, 49, 52, 53
- thinking: certainty that one is a thinking being, 40; critical assessment, 101–2; degrees of reflection, 101; and dreaming, 32–34; fundamental human cognition as not independently checkable, 50; reflective thought, 100–103 (*see also* reflective knowledge); thick thought

- and thin thought, 32–33. *See also cogito* beliefs; consciousness; reason; reflective knowledge
 time, and knowledge, 157–70
 Truetemp (thought experiment), 80–81n9, 88–91
 truth: beliefs that are true by accident, 15, 124–25, 208–9; Cartesian endorsement of a belief as doubtless true, 5–6; and clarity and distinctness of Cartesian perceptions, 42; and judgment, 15–16; justified true belief, 62–70, 87n; platitudes about, 114; true beliefs falling short of knowledge, 55–57, 87–92, 107, 124–25; truth and knowledge as aim of belief, 108–10. *See also* alethic affirmation; error; falsity; reliabilism
 Turri, John, 87n2, 88–89n2
 virtue epistemology, 140–56; and circularity, 18–19; competence virtue epistemology, 140–45; and conscientiousness, 145–46; constitutive vs. auxiliary, 140–45; Descartes's Pyrrhonian virtue epistemology, 1–20, 41–42 (*see also* Cartesian epistemology; Descartes, René); and epistemic competence, 122; and epistemic justification, 210–15; and fully apt performance, 98–99; responsibilism as a kind of reliabilism, 149–54; responsibilist virtue epistemology, 142, 145–49; and theory of competence, 191–206; and value of human knowledge, 113–16; virtue theory vs. situationism, 171–90
 virtue ethics, 171–83; Aristotelian virtue ethics, 207–8; crude vs. sophisticated virtue theory, 180–83; experiments, 172–74, 176, 182n16, 186–87, 189–90n19; illusory clash with situationism, 188–89; responses to critiques, 180, 184–88; situationist critique of, 172–83
 virtue perspectivism, 19
 virtue psychology, 171–73, 180–81, 186–87
 vision test, 78–81, 91
 waiter knocking over dishes (thought experiment), 123, 128–29
 waking state, 28–29
 Williamson, Timothy, 63–65, 78n5
 wine glass mimicking fragility (thought experiment), 135–36, 180
 Wright, Crispin, 221n6
 Zagzebski, Linda, 147–48
 Zimbardo, Philip, 189–90n19

