

Forthcoming in The Journal of Philosophy

REFLECTIVE KNOWLEDGE IN THE BEST CIRCLES

Ernest Sosa

Brown University and Rutgers University

A. Moore's Proof

Is the existence of external things just an article of faith? Certainly not, says G.E. Moore, and offers us a proof, thus aiming to remove Kant's "scandal to philosophy."

Moore's Proof

Here is a hand (a real, flesh and bone hand).

Therefore, there is at least one external thing in existence.¹

According to Moore, his argument meets three conditions for being a proof: first, the premiss is different from the conclusion; second, he knows the premiss to be the case; and, third, the conclusion follows deductively.² Further conditions may be required, but he evidently thinks his proof would satisfy these as well.

As Moore is well aware, many philosophers will feel he has not given "...any satisfactory proof of the point in question."³ Some, he believes, will want the premiss itself proved. But he has

¹ A simplified version of Moore's Proof.

² G.E. Moore, "Proof of an External World," Philosophical Papers (New York: Collier Books, 1962), pp. 144-5. Originally in the Proceedings of the British Academy XXV (1939).

³ *Ibid.*, p. 147.

not tried to prove it, and does not believe it can be proved. Proving that here is a hand requires proving one is awake, and this cannot be done.

Does Moore adequately answer the skeptic?⁴ Many have denied it for the reason that he fails to rule out a crucial possibility: that our faculties are leading us astray, e.g., that we are dreaming.

Aware of this objection, Moore grants, in "Certainty," that to know he is standing he must know he is awake.⁵ The point "cuts both ways," however, and he would prefer to conclude that he *does* know he is awake since he *does* know he is standing.

This has persuaded nearly no-one. On the contrary, some have thought him committed to an argument, M below, like the following:

Argument A

- A1. This map is a good guide to this desert.
- A2. According to the map an oasis lies ahead.
- A3. Therefore, an oasis lies ahead.

Argument M

- M1. My present experience is a veridical guide to reality (and I am not dreaming).
- M2. My present experience is as if I have a hand before me.
- M3. Therefore, here (before me) is a hand.

When challenged on premise A1, our desert dullard responds: "I must know A1 since the only way I could know A3 is through argument A, and I *do* know A3." Is this a just comparison? Is Moore's response to the skeptic relevantly similar?⁶

⁴ Henceforth 'knowledge' here will be short for 'plain knowledge', leaving aside Cartesian superknowledge.

⁵ Moore, "Certainty," in his Philosophical Papers, *op. cit.*

If Moore depends on argument M for his knowledge of M3, his response seems like the dullard's. The dullard is wrong to respond as he does. He must say how he knows his premiss without presupposing that he already knows the conclusion. And Moore would seem comparably wrong in the analogous response to the skeptic. In explaining how he knows M1, he must not presuppose that he already knows M3.

Does Moore depend on argument M for his knowledge of M3? There is reason to think that he does not, given his emphatic acknowledgment that he cannot *prove* M3. After all, M would seem a proof of M3 just as good as Moore's own "proof of an external world." Moore concedes, in effect, that *if* he does not know that he is not dreaming *then* he does not know of the hand before him. But that is *not* necessarily because he takes himself to know M3 only through M or any other such argument. And, in any case, even if he is relying on some such argument, which would require making that concession, the defender of common sense has other options.

One might, after all, make that concession only because of the following "principle of exclusion":

PE If one is to know that h, then one must exclude (rule out) every possibility
 that one knows to be incompatible with one's knowing that h.

As Moore grants explicitly, the possibility that he be just dreaming is incompatible with his knowing (perceptually) that he has a hand before him. And this, in combination with PE, is quite sufficient to explain his concession above.

Suppose Moore is not depending on argument M for his knowledge of M3. Although he recognizes his need to know he is not dreaming, suppose that is only because he accepts PE, our principle of exclusion. Then the sort of ridicule cast on the dullard is misdirected against Moore. What is more, it is not even clear that Moore must know *how* he knows he is not dreaming if he is to

⁶ Compare Barry Stroud, The Significance of Philosophical Scepticism (Oxford: OUP, 1984), esp. Chapters I and III.

know M3. That is not entailed by the application of the principle of exclusion. All that follows from the application of that principle is that Moore must know *that* he is not dreaming, not that he must know *how* he knows this.

In fact, however, the historical Moore *did* rely on something very much like argument M (more on this below). So is he not after all exposed to the damaging comparison with the desert dullard?

Not at all. There seems no good reason why, in responding to the skeptic, Moore must *show how* he knows he is not dreaming. Of course his response to the skeptic would be *enhanced* if he *could* show that. But it now seems *not* properly subject to ridicule even if he is not then in a position to show how he knows he is not dreaming. The question he is addressing is *whether* he knows that he is not dreaming, and, at most, by extension, what grounds he might have for his answer to *that* question, in answering which he does not, nor need he, *also* answer the question of *how* he knows himself to be awake and not dreaming.

It might be replied that one cannot know that there is a hand if one's belief rests on the unproved assumption that one is awake. According to Moore, however, things which cannot be proved might still be known. Besides, even though he cannot prove that he is awake, he has "conclusive evidence" for it. Unfortunately he cannot state his evidence, and the matter is left in this unsatisfactory state at the end of "Proof of an External World." But Moore has more to say in another paper of the period, "Four Forms of Scepticism."⁷ There he takes himself to know for sure about the hand before him, and takes this knowledge to be based on an inductive or analogical argument. We are told that introspective knowledge of one's own sensory experience can be immediate, unlike perceptual knowledge of one's physical surroundings. While agreeing with Russell that one *cannot* know *immediately* that one sees a hand, Moore thinks, *contra* Russell, that he *can* know it *for certain*. And he disagrees with Russell more specifically in allowing knowledge for

⁷ In his Philosophical Papers, *op. cit.*, pp. 193-223.

certain about his hand through analogical or inductive reasoning from premisses known introspectively.

However, it is doubtful that any allowable form of inference—whether deductive, inductive, or analogical—will take us from the character of our experience to the sort of knowledge of our surroundings that we ordinarily claim.

Familiar skeptical scenarios—dreaming, evil demon, brain in a vat, etc.—show that our experience prompts but does not logically entail its corresponding perceptual beliefs. Experience as if there is a fire before us does not entail that there is a fire there, experience as if here is a hand does not entail that here is a hand, etc. *Perhaps* what is required for one's beliefs and experiences to have certain contents entails that these could not possibly be *entirely* false or misleading. Indeed, some such conclusion follows from certain externalist and epistemic requirements on one's justified attribution of familiar contents to one's own experiences or beliefs. But even if that much *is* right—which is still controversial—one's experience or belief that here is a hand, or yonder a fire, might still be wildly off the mark. We cannot deduce much of our supposed knowledge of the external from unaided premisses about our experience.

As for inductive or analogical reasoning, only abductive reasoning—inference to the best explanation—offers much promise, but it seems questionable as a solution to our problem.⁸ Suppose (a) that we restrict ourselves to data just about the qualitative character of our own sensory experience, and (b) that we view belief in a commonsensical external world as a theory postulated to explain the course of our experience. What exactly is the proposal? Is it proposed that when ordinarily we accept the presence of a hand before us, we *do* know, and know on the basis of an abductive inference; or is it proposed rather that in such circumstances we have resources that *would* enable us to know if only we used those resources to make effective abductive arguments?

⁸ For Russell the “common sense hypothesis” of independent physical objects is “simpler” than the supposition that life is but a dream (as he explains in chapter II of The Problems of Philosophy). For Quine the “hypothesis of ordinary physical objects” is “posited” or “projected” from the data provided by sensory stimulations. “Subtracting

The second, more modest, proposal is *too* modest, since it leaves our ordinary perceptual beliefs in a position like that of a theorem accepted through a guess or a blunder, one that we do have the resources to prove after much hard thought, but one that we have not come close to proving at the time when we are just guessing or blundering.

Even the modest proposal, moreover, seems unlikely to succeed. *Could* we form a rich enough set of beliefs purely about the qualitative character of our sensory experience, one rich enough to permit abductive inferences yielding our commonsense view of external reality? This seems doubtful when we consider (a) that such pure data beliefs could not already presuppose the external reality to be inferred, and (b) that the postulated commonsense “theory” of external reality must presumably meet constraints on abductive inference: e.g., that the postulated theory be empirically testable and also simpler and less *ad hoc* than alternatives (e.g., Berkeley's). These requirements plausibly imply that our data must go beyond detached observations, and include some acceptable correlations. Yet these correlations are unavailable if we restrict ourselves to beliefs about the character of our experience.⁹ Most especially are they unavailable, and most especially is the postulated inference implausible, when our database is restricted, as it is by Moore, to introspectively known facts of one's own *then present* subjective experience, and to *directly recalled* facts of one's own earlier experience. (If deprived of the epistemic resources of testimony and of retentive memory—except insofar as such resources can be validated by reason-cum-introspection, which is not very far if at all—then there is precious little we can any longer see ourselves as knowing, thus deprived.)

Accordingly, the skeptic has a powerful case against Moore's claim that our knowledge of the external is based on an inductive or analogical inference from such information about our experience. It is not realistic to suppose that we consciously make such inferences in everyday life.

his cues from his world view, we get man's net contribution as the difference.” (Word and Object, p. 5.) That Quine's position is deeply problematic is shown by Stroud (*ibid.*, chapter VI).

⁹ This is argued by Wilfrid Sellars in “Phenomenalism,” in his Science, Perception, and Reality (London: Routledge and Kegan Paul, 1963).

It is less implausible to conceive of such inferences as implicit and/or dispositional, but even this strains belief. Besides, even granted that we make such inferences if only implicitly, do they yield simpler and less *ad hoc* hypotheses than alternatives? That is far from clear; nor do such hypotheses seem empirically testable and credible simply as explanations of the purely qualitative character of our then present or directly recalled experience.

Having reached a dead end, let us have some second thoughts on Moore's view of perceptual beliefs as inferential. Here he joined a venerable tradition along with Russell himself. If perceptual knowledge is thus mediate and inferential, what knowledge can qualify as immediate and foundational? Modern philosophy begins with Descartes's canonical answer to this question.¹⁰

B. Descartes's Circles

Descartes had two circles, not only the big famous one involving God as guarantor of our faculties, but also a smaller one found in the second paragraph of his third meditation, where he reasons like this:

I am certain that I am a thinking being. Do I not therefore also know what is required for my being certain about anything? In this first item of knowledge there is simply a clear and distinct perception of what I am asserting; this would not be enough to make me certain of the truth of the matter if it could ever turn out that something which I perceived with such clarity and distinctness was false. So I now seem to be

¹⁰ The shift to discussion of Descartes may seem abrupt; however, what we find about the nature of immediate knowledge in that discussion has important implications for a position that Moore failed to explore. Skeptics who are willing to grant Descartes his immediate knowledge through introspection or rational intuition would need to explain exactly why perception could never yield such knowledge. (And what of memory?) The discussion of Descartes to follow is meant to highlight this issue.

able to lay it down as a general rule that whatever I perceive very clearly and distinctly is true.¹¹

And yet when he looks away from particular clear and distinct items such as the proposition that he thinks, Descartes grants that a powerful enough being could deceive him even about what seems most manifest. Descartes grants that he could be astray in his beliefs as to what he perceives or remembers, and even in taking himself to intuit something as quite clear and distinct. This doubt must be blocked if one is to attain certainty by intuiting something as clear and distinct. Accordingly, Descartes launches the theological reflections that lead eventually to his nondeceiving God.

Even without the further boost of certainty provided by the proof of a nondeceiving God, however, Descartes takes himself to have attained some positive justification. Early in the third meditation he takes himself to perceive clearly and distinctly that he thinks, which he takes to be what gives him the certainty that he thinks. And he reasons that this clear and distinct perception would *not* give him such certainty if it were less than perfectly reliable, and apparently *concludes from this* that his clear and distinct perception *is* perfectly reliable. One could demand how he knows all these things: how he can be sure that he does clearly and distinctly perceive that he thinks, for one thing; and how he can be sure that there is nothing else in his situation that could provide the degree of certainty involved; and how he can be sure that the clarity and distinctness of his perception could not possibly provide that degree of certainty unless it were infallible. What could he say in response? Descartes might well have a uniform response to all such questions: in each case he might appeal once again to clear and distinct perception, each of the things in question being something we are assured of by our clearly and distinctly perceiving *it*.

About the *cogito*, I wish to highlight, not Descartes's answers to such questions, however, but the inference that he draws: *So I now seem to be able to lay it down as a general rule that whatever I*

¹¹ The Philosophical Writings of Descartes, eds. J. Cottingham, R. Stoothoff, and D. Murdoch (Cambridge:

perceive very clearly and distinctly is true. Just what is Descartes's argument in support of this general rule? Would his reasoning take the following form?

1. Datum: I know with a high degree of certainty that I think.
2. I clearly and distinctly perceive that I think, and that is the only, or anyhow the best account of the source of my knowledge that I think.
3. So my clear and distinct perception that I think is what explains why or how it is that I know I think.
4. But my clear and distinct perception could not serve as a source of that knowledge if it were not an infallibly reliable faculty.
5. So, finally, my clear and distinct perception must be an infallibly reliable faculty.

The move from 1 and 2 to 3 is an inference to an explanatory account that one might accept for the coherence it gives to one's view of things in the domain involved. Descartes does elsewhere appeal to coherence at important junctures.¹² So he may be doing so here as well, although questions do arise about how Descartes views coherence. Does he accept the power of coherence to add justified

Cambridge University Press, 1975), Vol. II, p. 24.

¹² In his Principles of Philosophy (Part IV, art. 205) for example, he notes that if we can interpret a long stretch of otherwise undecipherable writing by supposing that it is written in "one-off natural language," where the alphabet has all been switched forward by one letter, etc., then this is good reason for that interpretation. There he also argues for his scientific account of reality in terms of certain principles by claiming that "...it would hardly have been

certainty, and, in particular, would he claim infallibility for (sufficiently comprehensive and binding) coherence as he does for clear and distinct intuition?¹³ In any case, the comprehensive coherence of his world view would be enhanced by an explanation of how clear and distinct perception comes to be so highly reliable, even infallible. And this is just what Descartes attempts, through his theological and other reasoning. Descartes can see that reason might take him to a position that is sufficiently comprehensive and interlocking—and thereby defensible against any foreseeable attack, no holds barred, against any specific doubt actually pressed or in the offing, no matter how slight. Unaided reason might take him to that position. Need he go any further? What is more: Might one reach a similar position while dispensing with the trappings of Cartesian rationalism?

C. Circular Externalism

Compare now how Moore might have proceeded:

1. Datum: I know with a high degree of certainty that here is a hand.
2. I can see and feel that here is a hand, and that is the only, or anyhow the best

possible for so many items to fall into a coherent pattern if the original principles had been false." (From p. 290 of J. Cottingham, R. Stoothoff, and D. Murdoch, *op. cit.*)

¹³My attribution to Descartes is tentative because of the enormous bibliography on the "Cartesian Circle." In deference to that important tradition of scholarship, I do no more than *suggest* that there is logical space for an interpretation of Descartes that is perhaps more complex than many already tried, but that seems coherent and interesting. (I am myself convinced that this *is* Descartes's actual position, and defend this more fully in "How to Resolve the Pyrrhonian Problematic: A Lesson from Descartes," forthcoming in Philosophical Studies. In "Mythology of the Given," forthcoming in The History of Philosophy Quarterly, I argue for the relevance of this Cartesian strategy to issues of empirical foundations that have divided philosophers since the Vienna Circle, and have pitted, for example, Sellars, Rorty, and Davidson, on one side, against C.I. Lewis, Hempel, and Chisholm, on the other. My warm thanks to Lex Newman and James Van Cleve for helpful discussion of, and further references relevant to, this way of viewing Descartes.)

account of the source of my knowledge that here is a hand.

3. So my perception that here is a hand is what explains why or how it is that I know (with certainty) that here is a hand.
4. But my perception could not serve as a source of that degree of justified certainty if it were not a reliable faculty.¹⁴
5. So, finally, my perception must be a reliable faculty.

Moore could of course go on to say more about the nature of the perception that assures him about the hand. He might still say that such perception involves an implicit inference from what is known immediately and introspectively, perhaps an inductive or analogical inference of some sort. And that might make his view more comprehensively coherent, but we have already seen reasons why postulating such an inference is questionable. So we focus rather on a second alternative: Moore might well take perceiving to involve no inference at all, not even implicit inference, but only transfer of light, nerve impulses, etc., in such a way that the character of one's surroundings has a distinctive impact on oneself and occasions corresponding and reliable beliefs. This might also amount eventually to a comprehensively coherent view of one's knowledge of the external world. *And its epistemologically significant features would not distinguish it in any fundamental respect from the procedure followed by Descartes.*

The theme of accidentally true belief has loomed large in the epistemology of recent decades. The Gettier problem, for example, is posed by a justified belief true for reasons far removed from whatever causes it to be held and justified. Externalist conceptions of propositional

¹⁴ Here one would reduce Descartes's requirement of *infallible* certainty.

knowledge focus on this theme, as do one offered by Peter Unger (nonaccidentally true belief) and one offered by Alvin Goldman (belief caused by the truth of its content). And Nozick's tracking account is also a conception of this sort: S knows that p if and only if S believes correctly that p, and also (in the circumstances): *both* it would have been true that p only if S had believed it, *and* if it had not been true that p, then S would not have believed it.

Why are these conceptions of knowledge of special interest to us here? Because each offers a way to explain how one can know that p without reasoning from prior knowledge. The key idea exploited is this: you can know something noninferentially so long as it is no accident or coincidence that you are right.

Both the tracking and the causal accounts defensibly require a special nonaccidental connection between the belief and the fact believed. Nevertheless, in each case other levels of accidentality remain. Suppose I fancy myself a connoisseur of tomato ripeness, but suffer from a rare form of color blindness that precludes my discerning nearly any shade of red except that displayed by this particular tomato. Therefore my judgments of tomato ripeness are in general apt to be right with no better than even chance. But when it's the particular (and rare) shade of red now displayed, then I am nearly infallible. Oblivious to my affliction, however, I issue judgments of tomato ripeness with abandon over a wide spectrum of shades of red. Assuming that, unknown to me, the variety of tomato involved always ripens with this shade of red, then my belief that this tomato is ripe *is* in step with the truth, and arguably satisfies the requirements of Unger, Goldman, and Nozick. But, again, it is nevertheless in some relevant sense or respect only an accident that I am right in my belief.¹⁵ We need a clearer and more comprehensive view of the respects in which one's belief must be nonaccidentally true if it is to constitute knowledge.

¹⁵ For an early statement of this sort of problem, urged against Nozickian tracking, see Colin McGinn's "The Concept of Knowledge," Midwest Studies in Philosophy IX (1984): 529-554.

Unaided, the tracking or causal requirements proposed suffer from a sort of tunnel vision. They permit too narrow a focus on the particular target belief and its causal or counterfactual relation to the truth of its content. Just widening our focus will not do, however, if we widen it only far enough to include the process that yields the belief involved. We need an even broader view.

D. Virtue Epistemology

[When] ... thought is concerned with study, not with action or production, its good or bad state consists [simply] in being true or false. For truth is the function of whatever thinks.... *Aristotle, Nichomachean Ethics*¹⁶

Hence the function of each of the understanding parts is truth; and so the virtue of each part will be the state that makes that part grasp the truth most of all.

*Aristotle, Nichomachean Ethics*¹⁷

Virtue epistemology is distinguished by its emphasis on the *subject* as seat of justification. In order to qualify as knowledge a belief must be "apt," epistemically so, in a strong sense that goes beyond its being just a belief that coheres well within the subject's perspective. The "tracking account" (Nozick) sees here little more than a claim about that belief's counterfactual relation to the truth of what is believed. "Reliable indicator" accounts require rather that the belief itself or the reasons for it have properties nomically sufficient for its truth (Armstrong, Swain). "Reliable process" accounts focus instead on the cognitive process, beneath the skin, that yields the belief, and on the truth ratio in the products of that process, actual and counterfactual (Alvin Goldman).

¹⁶ Aristotle, *Nichomachean Ethics*, 1139a27-30; translated by T. Irwin.

¹⁷ *Ibid.*, 1139b11-13.

It is rather the subject and her cognitive virtues or aptitudes that hold primary interest for virtue epistemology. Consider the athletic virtues of a tennis champion. When we say that a shot is not just a winning shot but a skillful one, we imply a comment on the player as shotmaker. Suppose a tyro wields a racquet on a court and, unaware of the approaching ball, issues what amounts by luck to a stylish and effective backhand stroke. Such a shot might be an unreturnable winning shot, but it would manifest no real skill.

Why are we unwilling to admire a performance as "skillful" if it manifests only a fleeting, or even an instantaneous state of the agent's? Skills, abilities, competences, aptitudes, prowess—these come and go, true enough, but they do not flit by instantaneously. Why not? Why do we tend to define these concepts so as to require such stability? We might have defined similar concepts without requiring stability. Why do we define these concepts as we do? Why have we adopted these and not others? Should one not expect that, other things being equal, the more clearly useful a concept is to us, the more likely it is that we shall retain it? People need to know who are dependable members of their group—this is a *kind* of thing we need to monitor in a great variety of contexts, with a great variety of objectives. Cooperative success depends on the group's ability to monitor people's aptitudes and ineptitudes. So it is no surprise that the sorts of aptitudes (skills, competences, virtues) that we recognize and admire are those that linger stably.¹⁸

To praise a performance as skillful or an action as right, or a judgment as wise or apt, accordingly, is to assess not only the action or the judgment, but also the reflected aptitude or character or intelligence. This is a distinctive view with versions both in epistemology and in ethics. It is distinctive in that the rightness of an action (or a choice) and the aptness (or positive epistemic status) of a belief would involve not just whether the performance is *optimific* (if an action) or *true* (if a belief); nor just whether a good enough procedure was followed, *perhaps accidentally*, in arriving at

¹⁸ The social utility of concepts is invoked occasionally to defend a proposed account of knowledge in my [Knowledge in Perspective](#) (Cambridge: Cambridge University Press, 1991); e.g., on pp. 27 and 275. The form of argument involved is used insightfully in E.J. Craig's [Knowledge and the State of Nature](#) (Oxford: Clarendon Press, 1990), where it takes center stage.

that choice, or whether a good enough cognitive process chanced to lead to the belief; nor even just whether a rule somehow in effect demands that choice or that belief in those circumstances. Our virtue epistemology and virtue ethics focus rather on the agent and cognizer. When the agent's actions are said to be *right* and the cognizer's beliefs *knowledge*, we speak implicitly of the virtues, practical or intellectual, seated in that subject, which (a) give rise to that action or belief, adding to the subject's worth as agent or cognizer, and which (b) make him reliable and trustworthy over an interesting spread of possible choices or beliefs, and circumstances.

E. Virtue and Coherence

Can we explain what distinguishes a system of beliefs (and experiences) that, internally regarded, is intellectually virtuous and admirable? Presumably our explanation would involve the system's explanatory coherence, its overall simplicity and lack of adhoc epicycles, and so on. Philosophical mythology contains creatures who excel in all such respects, however, though their beliefs fall short of being knowledge even when true. Take the brain in a vat, for example, or the victim of the evil demon. An adult recently envatted, or victimized by the demon, can be indistinguishable from the best of us in respect of the comprehensiveness and coherence of their beliefs and experiences. Even when right about environing objects and events, such a victim's beliefs are far from being knowledge.

Might such comprehensive internal coherence exhaust all cognitive or intellectual virtue, at least when it comprises not only beliefs but also experiences? To assent here would be overbold even for a rationalist. Yet coherence is of course valued not only by philosophers but by the reflective more generally. One also wants faculties and virtues beyond reflective, coherence-seeking reason: perception, for example, and memory. Equally, internal coherence goes beyond such faculties, and requires reason, which counts for a lot in its own right. But why should that be so, if comprehensive coherence is no guarantee of truth, if the internal coherence enjoyed by the envatted yields little if any truth?

Compare first reason with memory. Input beliefs are required for retentive memory and inferential reason, which then yield beliefs as outputs. Retentive memory yields again the input belief itself, while inferential reason yields a new belief. Even the most excellent transmission faculty will not guarantee the truth of its output, which will depend not only on the transmission but also on the inputs. But our transmission faculties are valuable even so, if only because they *combine* with other faculties to increase vastly the total yield of true beliefs.

How does internal coherence, of little significant epistemic value in itself, become more valuable when combined with external aptness? Coherence-seeking inferential reason, like retentive memory, is of epistemic value when combined with externally apt faculties of perception, because when so combined it, like retentive memory, gives us a more comprehensive grasp of the truth than we would have in its absence.

Good perception is in part constituted by certain transitions from experiences to corresponding beliefs—as is the transition from the visual experience characteristic of a tomato seen in good light to belief in the tomato. Other such transitions help constitute good introspection, as when one's headache prompts awareness of it as a headache. Finally, if the comprehensive coherence of one's system of beliefs is at least in part responsible for its constitution and persistence, it thereby manifests a virtuous faculty of reason. Such comprehensive coherence is not just mechanical, but must reflect appropriate sensitivity to factors like adhocness, simplicity, and explanatory power. And it must include, not only belief/belief connections, but also experience/belief connections constitutive of good perception, and conscious-state/belief connections constitutive of good introspection. This broader conception of the coherence of one's mind involves not only the logical, probabilistic, and explanatory relations among one's first-order beliefs, but also coherence between these beliefs and one's sensory and other experiences, as well as comprehensive coherence between first-order experiences, beliefs, and other mental states, on one side, and on the other beliefs *about* first-order states.

We may well ask about certain aspects of broad coherence—e.g., the experience/belief transitions as well as the enumerative and abductive inferences involved—why these should be

viewed as adding to the subject's intellectual worth or merit. "Because it is truth conducive," or at least in good measure for that reason, we are told, "because it increases the likelihood that the subject will have true beliefs and avoid false ones." But that is obviously false of victims in skeptical scenarios, who nevertheless are internally coherent, and even epistemically justified.

Although that seems undeniable, we can perhaps understand it comfortably if we distinguish two sorts of epistemic justification: (a) S is "same-world justified" in believing P in world W iff S believes P in W in virtue of a faculty that *in W* is truth conducive; and (b) S is "actual-world justified" in believing P in world W iff S believes P in W in virtue of a faculty that *in our actual world* is truth conducive.

Such relativizing and contextualizing is familiar enough in ordinary thought and speech. Here it enables us to combine the following theses: (a) our broad coherence is necessary for the kind of reflective knowledge traditionally desired; and (b) such broadly coherent knowledge is desirable because in our actual world it helps us approach the truth and avoid error. This is not to deny that there is a kind of "animal knowledge" untouched by broad coherence. It is rather only to affirm that beyond "animal knowledge" there is a better knowledge. This reflective knowledge does require broad coherence, including one's ability to place one's first-level knowledge in epistemic perspective. But why aspire to any such thing? What is so desirable, epistemically, about broad coherence? Partly for the integrated understanding that it can help us to attain, and for which it is requisite. Partly because it is truth conducive, *even if in a demon world broad coherence fails this test, and is not truth conducive*. Even so, *we* can still regard broad coherence as intellectually valuable and admirable so long as we do not regard *our* world as such a world.

We are now, it seems to me, in just the position of arch-internalist Descartes. Consider the following passage:

The fact that an atheist can be "clearly aware that the three angles of a triangle are equal to two right angles" is something I do not dispute. But I maintain that this awareness of his [*cognitionem*] is not true knowledge [*scientia*], since no act of awareness that can be

rendered doubtful seems fit to be called knowledge [*scientia*]. Now since we are supposing that this individual is an atheist, he cannot be certain that he is not being deceived on matters which seem to him to be very evident (as I fully explained). And although this doubt may not occur to him, it can still crop up if someone else raises the point or if he looks into the matter himself. So he will never be free of this doubt until he acknowledges that God exists.¹⁹

Descartes considers reasons to doubt, not only one's faculties of perception, memory, and introspection, but even one's faculty of intuitive reason, by which one might know that $3+2=5$, that if one thinks one exists, and the like. And he defends against such doubts by coherence-inducing theological reasoning that yields an epistemic perspective on himself and his world, in terms of which he can feel confident about the reliability of his faculties, *including* the very faculties employed in arriving, *via a priori* theological reasoning, at that perspective on himself and his world, the perspective that enables him to see his world as epistemically propitious.²⁰

In *structure* virtue perspectivism is thus Cartesian, though in content it is not. Radical rationalism admits only (rational) intuition and deduction (along with memory) as its faculties of choice (or anyhow of top choice) and wishes to validate all knowledge in terms of these faculties. Thus the Cartesian grand project. Virtue perspectivism admits also perception and introspection, along with intuition and deduction, as well as inductive and abductive reasoning. Gladly using all such faculties, through testimony it accepts also the aid of one's epistemic community. Fortunately, the overview thus attained inspires confidence in the means used.

¹⁹ This passage is from the Second Set of Replies as it appears in [The Philosophical Writings of Descartes](#), eds. J. Cottingham, R. Stoothoff, and D. Murdoch, [op. cit.](#), Vol. II, p.101. However, I must add that where this translation says that an atheist can be "clearly aware," Descartes's Latin is "clare cognoscere."

²⁰ Although unremarked by Descartes, the role of dreams in his perception skepticism is analogous to a role assignable to paradoxes and aporias in a parallel skepticism vis-a-vis rational intuition.

Rejected as viciously circular by Descartes's critics, and by many today, our procedure does present a troubling aspect of circularity. However, a closer look may show this to be only an illusion.²¹

F. Epistemic Circularity: What's the Problem?

"I think, therefore I am," says Descartes, adding: "Here at last is something I really know. But what is it about this knowledge that makes it knowledge? As far as I can see, it is knowledge because it is a clear and distinct intuition. But it would not be real knowledge unless such intuition were reliable. So I can already lay it down as a general rule that clear and distinct intuition *is* reliable."

"Here is a hand," says G.E. Moore, adding: "Here is something I really know. But what gives me this knowledge? As far as I can see, it is knowledge in virtue of being a deliverance of perceptual experience. But it would not be knowledge if I were dreaming. So I can already conclude that I am *not* dreaming."

Descartes goes on to buttress the reliability of his rational intuition by developing a theology through vigorous use of that very rational intuition. And Moore can similarly appeal to what he knows about his reliable senses on the basis largely of those very senses.

But isn't any such reasoning circular? Yes, circular it does seem to be, "epistemically circular," let us say. But is it *viciously* circular? Skeptics through the ages *have* attacked it as such. Sextus Empiricus already uses the tropes of Agrippa in order to develop the so-called *diallelus*, or "problem of the criterion." And many have followed his lead in a long tradition. Today skepticism cum relativism has spread beyond epistemology and ethics, beyond philosophy, and even beyond the academy, and its champions often wield circularity as a weapon. But, again: Is such circularity vicious? To say that it is vicious, in the present context, is to say that it is somehow bad, intellectually bad, that it puts us in a situation that is somehow intellectually unsatisfactory. When

²¹ See my "Philosophical Skepticism and Epistemic Circularity," Proceedings of the Aristotelian Society Supplementary Volume (1994); and compare Barry Stroud's response, "Scepticism, 'Externalism', and the Goal of Epistemology," *ibid.*

we ask how the circularity is vicious, therefore, what we want to know is just how it puts us in an unsatisfactory state: when we reason in the way alleged to be viciously circular, wherein lies the defect in our reasoning or in the resulting state?

Largely through the use of rational intuition, Descartes supports the view that rational intuition is reliable and that through its exercise he knows that he thinks and exists.

Largely through the use of perception, Moore could support the view that perception is reliable, that he is *not* misled by a dream, and that through the exercise of perception he knows of the hand before him.

If a crystal ball claims itself to be reliable, then, largely through the crystal ball, a crystal-ball gazer could support the view that such gazing is reliable, that it is rarely misleading, and that through the crystal ball he can foretell the future.

Epistemic circularity is vicious, it might be said, because it would make the gazer as well justified as Descartes or Moore. *Since* there is no way to support adequately the view that intuition is reliable, or that perception is reliable, without employing those very faculties; and *since* the same goes for memory, deduction, abduction, and testimony; *therefore* there is no way to arrive at an acceptable theory of our knowledge and its general sources.

Perhaps that shows only how defective is the attempt to develop such a general theory of one's knowledge and its sources. There *is* an easy way to avoid the intellectual discomfort of having to use a faculty in answering the question whether that faculty is reliable: namely, not to ask the question. Call this the avoidance strategy.

Of course, we will hardly lack company if we avoid philosophy because we find it frustrating. But the avoidance strategy that I wish to consider is not just a rejection of what seems too difficult for one's own intelligence. The implication of the avoidance strategy is not that there is something lacking in one's intelligence but that there is something wrong with the questions avoided.

Much might indeed be wrong with our very general, philosophical questions. Many find them too abstract, too impractical, too useless, and so on. But these are *not* the concerns of *our*

avoidance strategist. He is after all a *philosopher*. His concern is not that the questions are just too hard for his intelligence, nor is it their abstractness, impracticality, or uselessness. He would hardly have gone into philosophy with *such* concerns, nor are they his concerns now. Difficulty, abstractness, impracticality, and uselessness are not in his view disqualifying drawbacks.

Why then should one as philosopher avoid questions of epistemology, such as those about the reliability of one's faculties? These questions become pressing with the realization that only if they reliably yield truth can our faculties yield knowledge. This is not just a commitment peculiar to contemporary reliabilism. Indeed, it is found already in Descartes, who, as we have seen, also stresses that intuition (and clear and distinct perception) yields knowledge only if reliable.

Consider again our principle of exclusion:

PE If one is to know that p then one must exclude (rule out) every possibility that one knows to be incompatible with one's knowing that p.²²

(By 'excluding' here I mean 'knowing not to be the case'.)

On the basis of PE, we can see that, in order to know that p, one must know that the faculties employed in arriving at one's belief that p are reliable faculties. After all, just consider the possibility that one's operative faculties be unreliable. That is surely a possibility generally known to be incompatible with attaining knowledge through them. Unreliable mechanisms of belief acquisition will not yield knowledge.

If the principle of exclusion is right, therefore, one cannot possibly know that p unless one knows that the faculties involved are reliable. But this is just the sort of knowledge that we seem able to attain only through epistemically circular reasoning.

²² Chapter 1 of Stroud's *The Significance of Philosophical Scepticism* is an illuminating discussion of this principle and its importance for understanding philosophical skepticism. In "How to Resolve the Pyrrhonian Problematic: A Lesson from Descartes" (*op. cit.*) I suggest how to derive it from other principles with independent plausibility.

One might of course question the principle of exclusion.²³ One might hold that in order to know that p one's pertinent faculties need only *be* reliable; one need not *know* them to be reliable.

One might for example appeal to a conception of knowledge as mere tracking. One might grant Rorty that causation should not be confused with justification, while joining Nozick in taking tracking as the essence of knowledge. "To know *is* just to mirror (or to track) nature. *Justification* is quite another matter. Justification of some sort may well require the principle of exclusion. Thus it may be that in order to be *justified* in believing that p, one must exclude every possibility one knows to be incompatible with one's knowing that p. But such justification is *not* required for simple knowledge."

That response seems essentially right. What is more, even Descartes would agree. For Descartes, you will recall, our knowledge that our faculties are reliable, *even* our faculty of reason, depends on our knowledge of God's epistemic good will. Yet Descartes grants explicitly that the atheist mathematician *can* know some mathematics—as we have seen.

The knowledge of an atheist is said to be *cognitio*, however, a second-class accomplishment by comparison with *scientia*. *Scientia*, by contrast, *does* require relevant knowledge of one's reliability. Only thus can one repel doubts about the possible unreliability of one's faculties. Only thus can one exclude a possibility evidently incompatible with one's knowing that p: namely, the possibility that only unreliable faculties yield one's belief.

By analogy we can more generally distinguish *animal* knowledge, which requires only that one track nature, on one hand, and on the other *reflective* knowledge, which requires also awareness of *how* one knows, in a way that precludes the unreliability of one's faculties. Unlike Descartes's

²³ For one thing, as it stands it leads, apparently, to a vicious regress. But that is an illusion. After all, what PE requires one to rule out is, *not* every possibility incompatible with one's knowing that p, but rather every possibility *known* to be thus incompatible. Since, for one thing, knowledge requires belief, the regress is hence not infinite, nor does it seem vicious.

cognitio and *scientia*, our more general animal and reflective knowledge do not require *infallible* reliability, but only a high level of reliability.²⁴

The avoidance strategy now has not only the cost of suppressing philosophical curiosity about knowledge. We can now see how it *also* precludes first-level *reflective* knowledge, and of course *scientia*.

Given these costs, what again counts in favor of avoidance? So far we have been told that we must avoid epistemic circularity because it entails arriving at a generally positive view of one's faculties only by use of those very faculties. But why should that be frustrating when it is the inevitable consequence of its generality. So far the answer is only that the superstitious crystal-gazer could reason analogously and with equal justification in defense of his own perspective. How damaging is this?

Suppose we grant the gazer epistemic justification and internal coherence equal to our own. Still, internal coherence is clearly insufficient. Isn't that obvious in view of paranoia, hypochondria, and similar psychoses? Logical brilliance permits logical coherence but does not even ensure sanity, much less general epistemic aptitude. There are faculties other than reason whose apt functioning is also crucial to the subject's epistemic welfare.

In light of that result, why not distinguish between the gazers and the perceivers in that, although both *reason* properly and attain thereby coherence and justification, only the perceivers are epistemically apt and attain knowledge?

On this view, the crystal-gazers differ from the perceivers in that gazing is not reliable while perceiving is. So the theory of knowledge of the perceivers is right, that of the gazers wrong. Moreover, the perceivers *can* know their theory to be right when they know it in large part through perception, since their theory *is* right and perception *can* thus serve as a source of knowledge. The gazers are by hypothesis in a very different position. Gazing, being *unreliable*, cannot serve as a

²⁴ This distinction figures in my Knowledge in Perspective (Cambridge: Cambridge University Press, 1991); see, e.g., pp. 240 and 282.

source of knowledge. So the perceivers have a good source or basis for their knowledge, but the gazers, lacking any such source or basis, lack knowledge.

Still one might insist that the perceivers should not be so smug. They should still feel acute discomfort and intellectual frustration. This I find a very widely shared view, in epistemology and, *mutatis mutandis*, far beyond. According to Barry Stroud, the perceivers can at best reach a position where they can affirm the conditional proposition that *if* their perception is reliable, *then* they know.²⁵ And he has recently reemphasized what is essentially the same thesis as follows:

... Sosa's "externalist" could say at most: "If the theory I hold is true, I do know or have good reason to believe that I know or have good reason to believe it, and I do understand how I know the things I do." I think ... we can see a way in which the satisfaction the theorist seeks in understanding his knowledge still eludes him. Given that all of his knowledge of the world is in question, he will still find himself able to say only "I might understand my knowledge, I might not. Whether I do or not all depends on how things in fact are in the world I think I've got knowledge of".²⁶

However, it is not easy to understand this position. If our perceivers believe (a) that their perception, *if* reliable, yields them knowledge, and (b) that their perception *is* reliable, then why are they restricted to affirming only the conditional, *a*, and not its antecedent, *b*? Why must they *wonder whether* they understand their relevant knowledge? Indeed, to the extent that they are really convinced of *both a and b*, it would seem that, far from being logically constrained *to* wondering whether they know, they are, on the contrary, logically constrained *from* so wondering. After all, first, if you are really *certain* that *p*, then you cannot well consider whether you know it without

²⁵ See "Understanding Human Knowledge in General," in Knowledge and Scepticism, ed. by M. Clay and K. Lehrer (Boulder: Westview, 1989), p. 47.

²⁶ B. Stroud, "Scepticism, "Externalism", and the Goal of Epistemology," *op. cit.*, pp. 303-304.

thinking that you do. Moreover, second, isn't it incoherent to be convinced that *p* and yet wonder whether *p*?

In sum, I see no sufficient reason to settle *either* for irresolvable frustration *or* for the avoidance strategy. The main argument we have seen for that depends on the claim that if we allow the circular defense offered for externalist epistemology, then the gazers turn out no less epistemically justified than the perceivers. In a sense that is true: but then in a sense they *are* equally internally justified, equally coherent. Nevertheless, they are *not* equally apt in all epistemically relevant respects. Perception is of course reliable while gazing is not. Therefore, the perceivers are right and apt both in their particular perceptual beliefs, at least generally, *and* in their theory of knowledge—for it all rests in large measure on their reliable perception. By contrast, the gazers are wrong and inapt both in their particular gaze-derived beliefs and in *their* theory of knowledge—for it all rests on their *unreliable* gazing. Moreover, I see no reason why the perceivers must be restricted to affirming only the conditional that *if* perception is reliable then they know. I see no reason why they cannot also affirm the antecedent, why they cannot believe, both rationally and aptly, that perception *is* reliable and does enable them to know.

G. Circles Beyond Belief

Why require the appeal to comprehensive enough coherence for justification, an appeal that I have attributed, tentatively, to Descartes, as part of what justifies his recourse to theology in accounting for true knowledge (*scientia*)? Why not say that what justifies is that one's beliefs be caused by the gods? And if the question arises, why not add that *this* belief itself is justified because it is itself caused by the gods? We *could* of course proceed in this simplified way without worrying about coherence or about the source of these beliefs beyond attributing them to divine agency. But that is not the way we are built, most of us: we just do not acquire such beliefs the way we do acquire beliefs willy-nilly when we open our eyes in good light. But what if we *were* built that way? Would we then be justified in having such beliefs, and in explaining our justification for having them, by their origin in divine agency? Would we then be justified to the degree and in the way in which

Descartes is justified or in the way in which our imagined Moore would be justified through his appeal to a more ordinary reliabilism than that of Descartes? Internally regarded, the structure of beliefs would share prominent features in all three cases. Of course, from our Moorean, commonsense position we can object both to Cartesianism and to the invocation of the gods. These views are internally coherent, but we might still reject them as wrong. And we might be able to explain what is wrong with them, from our point of view, especially if our point of view rules out their leading ideas. But they can, for their part, return the favor. Besides, we can anyhow imagine someone brilliant but insane, who weaves a system of immense interlocking complexity, but one wholly detached from reality as we know it commonsensically. Such a madman could object to our commonsense beliefs in a way that would seem relevantly analogous to the way in which we would object to his mad beliefs.

What all of that shows, it seems to me, is nothing more than that knowledge does not live by coherence and truth alone. Knowledge requires truth and coherence, true enough, but it often requires more: e.g., that one be adequately related, causally or counterfactually, to the objects of one's knowledge, which is not necessarily ensured by the mere truth-*cum*-coherence of one's beliefs, no matter how comprehensive the coherence. Madmen can be richly, brilliantly coherent; not just imaginary madmen, but real ones, some of them locked up in asylums. Knowledge requires not only internal justification or coherence or rationality, but also external warrant or aptness. We must be *both* in good internal order *and* in appropriate relation to the external world.²⁷

²⁷ I am pleased to acknowledge helpful discussion, first here at Brown, in a faculty discussion group and in my graduate seminar; then at the University of California, Santa Barbara; at the Oxford Philosophical Society; at Cambridge University; and at the University of Granada.