



BRILL

INTERNATIONAL JOURNAL FOR THE STUDY OF SKEPTICISM

XX (2025) 1–19

INTERNATIONAL  
JOURNAL  
FOR THE STUDY  
OF SKEPTICISM  
brill.com/skep

## Commentary



### So, Was Moore Basically Right?

AQ1

*Crispin Wright* | ORCID:

Division of Law and Philosophy, University of Stirling, Stirling, UK

*c.j.wright@stir.ac.uk*

Received 2 September 2025 | Accepted 5 September 2025 |

Published online XXXX

#### Abstract

Adam Leite's core claim is that there is no need for epistemological theory in addressing the challenge of external world scepticism, that our "pre-philosophical position" already contains resources sufficient to neutralise it. It is argued on the contrary that, properly understood, the best of the sceptical paradoxes arise *within* our pre-philosophical position. An "upper-level" paradox, targeting not knowledge but the notion of being in a position to claim knowledge, is developed to illustrate this, and it is suggested that the real lesson of scepticism is one repeatedly hinted at in Wittgenstein's *On Certainty*: that all enquiry rests upon framework presuppositions which are not themselves the product of enquiry, and of which we accordingly have no knowledge when that is understood as requiring cognitive achievement. The mistake of 'The Sceptic' is to interpret this point as somehow debunking our knowledge, rather than as teaching us something about its nature.

#### Keywords

G.E. Moore – scepticism – paradox – authenticity conditions – claims to knowledge

Wittgenstein simply did not know what to say about the paradoxes. I don't either. But one thing is clear: the fruitful problem is not to 'get rid of them' but to get something out of them.

KREISEL 1958: 157

## 1        **Leite's Project**

Adam Leite's (2024) book offers an exceptionally—I think an unprecedentedly—detailed working out of an outlook on the alleged challenge presented by scepticism about our knowledge of the external world which, as he explains in its early chapters, draws major inspiration from the approaches of Moore and Austin. Leite's treatment accordingly begins and ends with the idea that there is nothing fundamentally amiss with our ordinary 'commonsensical' talk of knowledge,<sup>1</sup> or with our ordinary ideas about the extent of the knowledge that we actually have. That idea is of course something which most of the philosophical work on scepticism of the last century has aspired to vindicate. But the general assumption has been that accomplishing this goal, if possible at all, will call for innovative philosophical theory. Leite holds by contrast that "our pre-philosophical position" already contains resources sufficient to rebut 'The Sceptic'. In this short assessment, I'll try to explain why I continue to prefer the recently more generally accepted, opposed idea that scepticism, at its most formidable, is not to be disposed of by appeal merely to common-sense resources and hence why, in my judgement, the general effect of Leite's discussion is to underestimate his redoubtable, if fictional, opponent.

## 2        **Our "Pre-philosophical Position"**

What does Leite intend us to understand by this phrase? Here is his initial and most extended account of it:

[W]e must begin with the relation between epistemological theorising and ordinary life. There is a familiar distinction—going back at least to Hume—between our ordinary, non-philosophical position and the

---

<sup>1</sup> For simplicity, I shall most often explicitly refer in what follows just to knowledge, omitting mention of epistemic justification. But except where otherwise stated, the reader should understand an implicit reference to epistemic justification also.

position of philosophical reflection on human knowledge. In this ordinary, non-philosophical position we are committed to all sorts of claims about the world and about what we know and have reason to believe. This position involves commitment to various practices, rules of thumb, and principles of reasoning, inquiry, and epistemic assessment, including everything involved in the conduct of science. It is hardly unreflective; reflective summaries and reflective self-correction—as in methodological discussions in science—have a central place in it. However, it doesn't involve commitment to any fully general theories about such matters as what philosophers call “the structure of empirical justification,” nor anything like a general account of the requirements one must meet in order to know things. It doesn't even involve a demand for such things. In this sense it is a *pre-philosophical* position.

LEITE 2024: 17–18

Two points are salient. One is that the pre-philosophical position (henceforward, I'll often abbreviate to “PPP”) is allowed to be to a significant degree reflective, involving thoughts about the good/bad standing of various practices and epistemic principles, including “everything involved in the conduct of science”—so it is allowed to include reflection on scientific method, embracing issues to do with conditions on the effective design of experiments, conditions on the statistical significance of data, and so on. That kind of thing may seem pretty theoretical, to be sure. What, it seems, Leite wants the PPP *not* to involve is any commitment to “fully general” ideas about knowledge or the nature of epistemic justification—the kind of generality involved in the attempt to give a constitutive analysis of knowledge, for example, or a general characterisation of the justificational architecture allegedly common to all propositions of a certain general subject matter (for instance, all propositions about the mental states of others, the past, or the external world.) So anyway, it seems, we should understand Leite's intent.

While the intended distinction certainly could be clearer—and to the extent that it is unclear, it may be difficult to assess whether the resources Leite believes he can use to confound ‘The Sceptic’ are indeed all drawn from our PPP—the application that he goes on to make of the distinction, together with various further passages of explanation, make it clear that the PPP is to include a range of specific assumptions about the *extent* of our knowledge, assumptions like these:

- that a very great deal of our actual knowledge of the external world depends on other things we know about the world;

- that we know a good deal about the range of circumstances in which the impressions our senses give us of how things stand in the world are less than reliable;
- that we often use background information about the world to determine which are such circumstances;
- that if challenged, for example, to justify your belief that there is a pen on the desk in front of you, there is no end of information on which you can draw: for example, that you remember putting the pen there five minutes ago, that no one else has come into the room in the meantime, that you can see it there, that you feel perfectly cognitively lucid, and have no reason to think you might be hallucinating, etc. ...<sup>2</sup>

Again: Leite's central idea, emphasized throughout the book, is that in order to disarm the idea that external world scepticism presents any form of serious intellectual challenge, we need only draw on such resources and others that are available to us in the PPP.

### 3 Initial Misgivings about the Approach

There are reasons to be doubtful about this dialectical approach even before we engage with the specifics of Leite's development of it. First, our "pre-philosophical position" surely encompasses—must be allowed to encompass—a range of evaluative notions associated with the idea of epistemic *agency*: notions like epistemic responsibility, epistemic blameworthiness, and epistemic good practice. More specifically, the pre-philosophical concepts on which sceptical argument puts pressure—knowledge, justified belief, evidence, rationality, etc.—are all intuitively thought of as *normative* concepts. They are concepts which collectively implicate standards by which, so we conceive, we *ought to* operate in agentially managing our systems of belief. We hold people responsible in large measure for what they take to be true and, as attested by our use of pejoratively-flavoured concepts such as bias, prejudice, gullibility and blinkered optimism, adopt 'Strawsonian' reactive attitudes to them if they come short by the standards concerned.<sup>3</sup>

<sup>2</sup> I am paraphrasing passages from Leite (2024: 222–223).

<sup>3</sup> The reference here is to Strawson's (1963) influential discussion of free will. There are of course debatable and well debated issues close by to do with whether such attitudes can be appropriate as applied to belief unless, perhaps implausibly, believing can be conceived as some kind of state voluntarily entered into. My own view, which I will not defend here, is that there is no real tension in the vicinity. It is consistent with the appropriateness of the Strawsonian attitudes that the management of one's beliefs is subject to extensive

What follows from this reminder? Simply that, on any plausible construal of our “pre-philosophical position,” it must be allowed to include a plethora of beliefs, not just about what we normally take ourselves to know or justifiably believe—beliefs, that is, about the *extensions* of our signature epistemic concepts but *also about what constitutes compliance with the norms they implicate*. Leite’s characterisation seems to involve some element of recognition of this. But then we surely have to recognize the possibility of a mismatch between what the PPP recognises as knowledge that we actually possess and what it has to say about the methods by which knowledge can be soundly acquired. And, of course, that there *is* indeed such a mismatch is exactly the general form of the conclusion that one kind of interesting sceptical argument will purport to foist on us. Such an argument will purport to disclose a tension between what we normally take ourselves to know or to justifiably believe on the one hand and, on the other, things we also believe about the appropriate standards that attaining knowledge or justified belief ought to respect. So, any interesting sceptical argument will thus operate by playing off elements in the PPP *against each other*—and thus cannot be open to successful rebuttal if the attempt is restricted to deployment just of resources of the PPP itself.

There is anyway reason for doubt about the potential significance in this context of Leite’s core idea that there is a substantial distinction between epistemological theory and the PPP. Think about the analogous idea in the case of ethics. There too, there is something like an analogue of Leite’s idea of a pre-philosophical position: a repository of ordinary ethical opinion and practice to be distinguished from the philosophical theory of ethics. But note that it is part of being an ethical agent—contrast: an ethical theorist—that you be prepared to think about and (re-)evaluate your ethical practices and the standards they exemplify. For example, you don’t need to appeal to theoretical ethics in order reasonably to challenge someone’s apparent investment, manifest in the way they live their life, in the alleged moral importance of a distinction between the morality of doing bad things and the morality of merely permitting foreseen bad consequences of inaction. This particular challenge is something that in difficult times becomes quite urgent and affects us all; you don’t get off the hook by saying, “That is a theoretical matter, best left to the philosophers.” In such circumstances, it behoves one to think about and try to clarify and evaluate cases of the distinction. It is not

---

voluntary control even if believing and disbelieving are not voluntary actions or voluntarily accomplishable states. (A rough analogy: compare my responsibility for the state of my garden and for the flourishing, or otherwise, of its plants.)

merely an option that can be exercised by those who interest themselves in ‘ethical theory’.

It is the same, I suggest, with the epistemological case. Problems for our epistemic norms are seldom quite as urgent as those engaging our ethical standards. But I see no difference in principle. Insofar as lax, prejudicial, cynical or presumptuous habits of thought can be an instrument of real social harm, we are similarly obliged to reflect on our epistemological standards, and to that extent to be epistemological theorists.

The correct take on the best arguments of ‘The Sceptic’, I want to suggest, is that they purport to disclose incoherence precisely within and between our pre-philosophical ideas about the standards for believing that knowledgeability requires and our customary confident applications of the concept of knowledge to an agreed range of cases. Like any strong paradoxes, these arguments challenge us to make sure we confront them in their most formidable guise and then to disclose independently motivated criticism of their premises, or of their reasoning—or otherwise, even perhaps to recognise that a real glitch in our intuitive thinking has actually been disclosed.

Consider, for example, how an analogue of Leite’s approach plays out in response to a Sorites paradox for ‘looks red’, developed with respect to a series of colour patches transitioning apparently seamlessly from a scarlet-looking patch at one end to a vivid cobalt blue at the other. “Seamless” here involves that there is no discernible difference in hue, however good the lighting, however sharp one’s vision, between any pair of adjacent patches. The major premise for the paradox states that the meaning of phenomenal colour predicates is such that of any pair of surfaces whose respective shades of apparent colour cannot be told apart, no matter how good the lighting conditions or acute the vision of the observer, any correct description of the phenomenal colour of one must be a correct description of the other. *Phenomenal* colour concepts, the thought is, cannot distinguish between things that look exactly alike.

Well, that can seem *evidently* correct as a reflection of the meanings of phenomenal colour predicates: surely it is part of our ‘pre-philosophical’ understanding of the relevant predicates. However, our pre-philosophical thinking about the patches at the ends is that one definitely looks red and the other definitely looks blue, So—the Leite-ish move—that gives us all we need to reject the major premise. No need for any ‘theory’. But *of course* there is such a need. If we are going to reject that premise, we owe an improved account of the meanings of predicates of phenomenal colour, and indeed of phenomenal predicates in general, under the aegis of which we can explain away the temptation of pre-philosophical thought to accept the premise—or

else to recognise, as Michael Dummett (1975) once suggested, that predicates expressive of phenomenal properties are actually semantically inconsistent.

#### 4 In Mitigation

A reader with whom these misgivings resonate might be tempted to conclude that Leite's book will not be worth pressing on with. That, though, would be a mistake. For one thing, Leite is not committed to the view that philosophical paradoxes *in general* should be, as it were, "hissed out" in a Moorean spirit. He can agree that that is not an appropriate response to philosophical paradoxes without exception, while still maintaining that in the special case of external world scepticism, some of the assumptions at the root of 'The Sceptic's' arguments belong to a level of epistemological theory much loftier and more rarefied than—and indeed at odds with—anything incumbent on an ordinary, responsibly reflective epistemic agent. So we should read on and see. And for a second thing, *How to Take Scepticism Seriously* is anyway full of arguments, objections and counter-objections whose interest does not, or not always, depend upon sympathy with Leite's overarching approach. His discussion is often, in fact, actually much less Moore-ish than his official advertising conveys, and the moves that he makes in the attempt to see off 'The Sceptic' merit careful consideration whether or not the resources they require do indeed draw only on aspects of our PPP.

#### 5 Leite's Sceptic

Leite's response to scepticism is almost exclusively addressed to what we might call Scenario Scepticism. We are, familiarly, invited to consider a situation (the Sceptical Scenario)—stock examples include experiencing a persistent, perhaps perpetual lucid dream; suffering a persistent, global, and coherent hallucination caused by Descartes' powerful Evil Demon; or being trapped in a context in which we have been "envatted"—where we are radically but undetectably mistaken about the *bona fides* of all our apparently perceptual experience, and almost nothing is as our experience represents it. An imaginary mischief-maker then denies that we know that such a Sceptical Scenario does not obtain, on the ground that everything would seem just the same if it did, and, on that basis, challenges all the knowledge of the external world that we take ourselves to have.

In thinking about the issues raised by this train of thought, we do well to usher ‘The Sceptic’ off the stage. As the way I have been expressing matters hitherto will have made clear, my suggestion is that we are dealing with a paradox, not with an imaginary philosophical opponent whom we need to confound, or at least silence, by argument.<sup>4</sup>

So, how exactly does the challenge to our presumed perceptual knowledge issue, according to the Scenario Sceptical argument? It depends in part on the detail of the Sceptical Scenario—specifically, whether like the perpetual dream, the Scenario calls into question your proper perceptual function, or whether more specifically it is meant to be inconsistent with the truth of all or the majority of your quotidian perceptual beliefs. But either way, the sceptical thought runs, you do not—cannot—know that the Sceptical Scenario does not obtain. And that ignorance, one way or another, surely undermines the status as knowledgeable of all your quotidian apparently perceptual beliefs.

The details of how exactly the undermining is supposed to work don’t matter at this point if—Leite’s first response—we are allowed deploy resources from the PPP to repudiate the claim that we do not know that the Sceptical Scenario does not obtain. In Leite’s deployment, the kind of repudiation thereby made possible goes like this. I *do* know that I am not now suffering a persistent lucid dream because I know all kinds of things about what dream experience is phenomenologically like, and how dissimilar it is to waking experience, and my experience right now is paradigmatically of the latter sort. Moreover, I also know that there is no Evil Demon—that notion is the merest fantasy of the mediaeval mind and even if there were such a creature, there is no way except by magic that it could in the manner suggested take comprehensive and continuous control of my experience: moreover, we know, alas, that there is no real magic. I also know that the technology necessary to preserve my disembodied brain in a vat of nutrients and to monitor and to regulate every aspect of its phenomenally conscious life, does not and quite likely never will exist.

I find this aspect of Leite’s response uninteresting for the reasons elaborated above. From a perspective in which what we are dealing with is a paradox which challenges elements in our PPP, it is totally intellectually unsatisfying merely to confront it with aspects of that position. It is as if we were to respond to

---

4 This shift in perspective has an effect on the rules of discussion. It has the effect that your antagonist in debate with ‘The Sceptic’ is in effect yourself—at once drawn by the assumptions and line of argument in question, but repelled by its consequences. In dialogue with a real opponent, you can win by trapping them into self-contradiction. But if your opponent is yourself, that will not be an intellectually satisfying position in which to leave matters. Cf. Wright (1991: 88–89).

Zeno's paradoxes of motion by simply flatly asserting that motion is evidently possible, or to the Sorites considered above by simply flatly asserting that, of course, looking blue and looking red are coherent but mutually exclusive concepts, so that the major premise of the soritical reasoning involving them has to be plain false and there's an end on't.<sup>5</sup> What is needed is diagnosis of the attraction and a principled dissolution of the premises or reasoning of the paradox.

This, though, is an instance where Leite might seem to have a better point to offer, one that is independent of his overall Mooreanism. He makes the apt observation that the various Sceptical Scenarios standardly canvassed in developments of the paradox are *epistemically asymmetric*. That, if the Scenario obtained, we would not be able to tell—would have no evidence, one way or the other—does not entail, he asserts, that we have no knowledge that it does not obtain *if indeed it does not*. On the contrary, if I am not envatted, nor duped by the demon, nor dreaming, then surely there is no end of information, Leite's thought is, that I have or can get that rules these hypotheses out. So, if the sceptical reasoning relies on the premise that there is no knowing *whether or not* the scenario obtains, it relies on a thesis that is true only if we actually are in the sceptical scenario. In that case, a key premise in the Scenario Sceptical paradox would be unmotivated.

On reflection, though, matters are not so clear. We have a conditional: *if* the Sceptical Scenario does not obtain, and we inhabit a world much as we normally conceive it, then we have lots of quotidian perceptual knowledge that is inconsistent with the supposition that the Sceptical Scenario obtains.<sup>6</sup> But to *utilise* this conditional to facilitate a knowledgeable repudiation of the Sceptical hypothesis, we need first to know its antecedent, do we not? So, if the idea is somehow to marshal the perceptual knowledge that we will have *provided* the Sceptical hypothesis is false against the Sceptical hypothesis, there is the clear threat of a circle.

I think that if we pursue this suspicion, a descendant of the Scenario paradox—though perhaps sufficiently different to be rated a different paradox—does indeed emerge. I will outline how in Section 7. However, doing so will require some additional pieces of conceptual kit, as follows.

---

<sup>5</sup> Notwithstanding that there are philosophers who have argued for exactly this.

<sup>6</sup> Actually, it is not clear that we do have even this conditional, unless some kind of externalist conception of knowledge is assumed. Leite, notably, does not explicitly endorse such a conception, and presumably would anyway not want to include any such general epistemological thesis as part of our PPP.

## 6 Some Additional Resources for The Sceptic

A. First, we need the distinction between *possessing* knowledge and *being in position to claim* that one knows. Such a distinction is, intuitively, underwritten by many externalist conceptions of knowledge. It is for example evident that I may arrive at a true belief, B, by what is in fact a *reliable method*, or that I may arrive at B *safely*, or *sensitively* in Nozick's (1981) sense, while remaining unaware of such facts about the provenance of the belief. In that case, while the process whereby I adopt B has putatively relevant externalist credentials sufficient for knowledge, that this is so is, and hence I am in no position to claim to know—according to these conceptions of knowledge. To be sure, those three particular historical externalist proposals about the constitution of knowledge have been found variously wanting. But that's beside the present point, which is that to the extent that the factors which underwrite the knowledgeability of a true belief are held to consist in matters external to the reflective awareness of the subject, the subject may not be in position to claim that they obtain.<sup>7</sup>

What is interesting is that, intuitively, even on externalist accounts of knowledge in general, the possibility of such a gap between knowing and being in position to claim knowledge does not arise in every case. Some kinds of knowledge of some kinds of propositions may be achieved in such a way that they are, as a recently introduced terminology enables us to express the matter, *luminous*:<sup>8</sup> if you get to know a suitable such P by one of a number of suitable methods, then you will be in position to be reflectively aware that exactly that is what you have done. Plausible examples include many kinds of psychological self-knowledge, based purely on awareness of occurrent mental states, knowledge achieved by proof in mathematics—if you confidently and competently construct a proof, you will be in position to be aware that you have done exactly that—and perhaps also some basic non-inferential *a priori* knowledge (knowledge of certain kinds of elementary analytic propositions, for example). But perceptual knowledge is, of course, not of this luminous character. If it were, then the distinction between genuine perception and the kind of delusion highlighted in the Sceptical Scenarios would be a distinction one could draw purely on the basis of reflective awareness. The paradox derives its apparent pulling power from the fact that it is not. And so much is implicit

<sup>7</sup> Note that the issues here are orthogonal to the question of the externalistic validity of the KK principle, which requires, of course, that the outer 'K' also be externalistically understood.

<sup>8</sup> I believe it was Timothy Williamson (1996) who first introduced "luminous" as philosophical jargon in this sense. Williamson argues that almost nothing is luminous. His argument, though widely influential, seems to me to be mistaken (see Wright 2023: 39–42).

in Leite's willingness to allow that, if one is the Sceptical Scenario, one will have no way of knowing that one is.

These reflections invite us to be alive to the possibility of a sceptical paradox that challenges not our perceptual knowledge, but our being in position to claim it. Since I am confident that Leite does want us to be in position to claim extensive knowledge of the external world, we should try to see what form such a paradox might assume, and how Leite's approach might be able to tackle it.

B. A second, additionally needed bit of kit is the familiar notion of an *undermining*<sup>9</sup>—contrast: an overriding—defeater for a belief. Suppose we are in the business of trying to determine whether or not a proposition P is true. Let E be the evidence so far gathered and suppose we incline to regard E as defeasibly sufficient for P. Then U is a potential underminer of E's support for P if (we are of the view that) if U is true, *then while that has no direct bearing* on whether P itself is true, E should no longer be regarded as sufficient for accepting P.<sup>10,11</sup>

C. Third, we need a notion of *open-mindedness*. Your open-mindedness about P, in the sense that concerns us here, requires that

(i) you have no opinion whether or not P, but

9 John Pollock says "undercutting." *Locus classicus*: Pollock (1974).

10 I imagine these ideas are pretty familiar for most readers, but for the benefit of any for whom they are not, here's an example. Let P be that you—a resident of Indiana—have, as you suspect may be true after a gathering in which many were infected, contracted COVID 19 but are so far asymptomatic, and let E be the positive outcome of a lateral flow test that you have just taken using a kit purportedly supplied by a well-known pharmaceuticals manufacturer based in San Diego, CA. We can count it as an additional part of your evidence that the kit is advertised at approximately 98% positive accuracy—false positives occur on average only in 1 of 50 tests. Now suppose you learn—this is U—that the South Bend depot in which the kits are stored for distribution in Indiana has inadvertently taken in a large number of defective kits. Presumably, this information should rationally be received as to some significant, maybe decisive, degree compromising E's support for P, but has no direct bearing on whether or not P is true.

11 Note that whether something is an underminer of the significance of a given item of evidence may vary with collateral information. It may only be in the context of background theoretical knowledge that a given kind of evidence for a certain belief should be regarded as losing (some of) its supportive force. For example, up on the main ridge of the Cuillin Hills on the Isle of Skye in thick mist, I get a steady compass signal (E) indicative, in conjunction with the map, that I am indeed on track and can safely continue on in the same direction. Someone says, "You know the rock around here is predominantly gabbro?" Innocent of the powerfully magnetic properties of gabbro, I reply "Yes," and carry on (perhaps tragically) trusting my compass, undaunted. In what follows, we can harmlessly prescind from this kind of information-relativity in the notion of an underminer.

- (ii) you are nevertheless not *agnostic* about P in the sense of “agnostic” frequently associated with the stance of religious agnosticism, that is, a stance which holds, e.g., that the existence of God is a matter on which no rational subject can in any circumstances consider that they have a knowledgeable or even a justified view. Your open-mindedness about P, by contrast, involves readiness to accept or reject P, should appropriate evidence turn up but also, as we shall here understand it,
- (iii) your being of the opinion that *so far* you have no basis for a view, again while not ruling out the possibility of new information that might mandate taking one.

With these last two notions—open-mindedness and undermining defeat of evidence—in place, we can elicit an evident but interesting

*Corollary:* Suppose you are presented with evidence E, for P, and suppose you are open-minded about whether something you recognise as a specific potential underminer, U, for E obtains. (Again, note that we are not supposing that you are *sceptical* about U—doubtful that U is true—but merely that you consider that you so far have no basis for a view about whether or not U is true.) Well, if U does obtain, then, as you recognise, E should not be regarded as supportive of P. So if you are open-minded about U, you ought rationally to be open minded about *whether or not E should be regarded as supportive of P*.

More explicitly, here is the reasoning that underlies that thought. U, if true, is as you recognise an underminer of E’s evidential force for P. So if E is supportive of P, U is false. So, if your present state is one of open-mindedness whether U is false or not, it cannot be that you are in epistemic position to regard E as supportive of P.

There is, to be sure, an implicit appeal here to closure for whatever relevant epistemic notion is supposed to be at odds with open-mindedness. But that seems perfectly intuitive. Consider an example. You are running a schoolroom physics experiment in which it is important that there be no leakage of pressure from the apparatus. If pressure is lost, = U, the significance of the findings for the hypothesis under test, = P, will be severely compromised. Suppose you are open-minded about whether pressure *was* lost; members of the class attached the clamps, say, but you haven’t checked their work and reckon it quite possible it was slapdash. However you run the experiment, take the final readings, = E, and announce that they do indeed confirm P. Isn’t that irrational? Can you rationally profess to have confirmed P while remaining open-minded about whether pressure was indeed lost?

I have introduced the relevant options here in relation to cases where what is at issue is the suasive force of defeasible evidence. But for the purpose of

what is to follow, it is vital to note that the basic framework applies equally to cases where P is based on some form of (putatively) *direct* cognition and reliance on defeasible evidence is not involved: perception (on some views), episodic personal memory (on some views), awareness of certain of one's own mental states, or (perhaps) non-inferential intellectual intuition, for example. Even so, cognitive faculties of the relevant direct kind have their own widely understood conditions of effective operation. Accordingly, a claim to knowledge based on the presumed operation of such a faculty may be undermined by finding that these conditions were compromised in one way or another. Perhaps the conditions under which the faculty was operating were unsuitable for its effective function. (Suppose, for instance, you were in Barn Façade County.) Or perhaps the faculty itself was suffering some kind of interference or internal malfunction. (Perhaps you had been slipped an hallucinogenic substance of some kind.) In such circumstances, once again, you cannot rationally advance a claim, P, based on the presumed effective operation of such a faculty, and at the same time profess open-mindedness about whether some condition, U, obtains which you acknowledge would be a potential underminer of the operation of the faculty—a condition such that, if it obtained, the effective functioning of the faculty concerned would indeed be likely to be compromised.

D. We still need one more notion. Where U is any underminer of the support of evidence E for P, let us call the negation of U an *Authenticity Condition* for E with respect to P. Likewise, where U is an underminer of the capacity of faculty F to reliably determine whether P, let the negation of U be an Authenticity Condition for F with respect to P.

In Wright (2022), I have deployed these notions to characterise the idea of what it is for a valid argument to be *cogent*—contrast: *non-transmissive*. Say that an argument  $\{P \Rightarrow Q\}$  is potentially cogent for a rational thinker who is *either* presented with evidence E for P, *or* who achieves a putative verification of P (“There’s a shire horse in that field”) by employing faculty F (looking and seeing), just if *acknowledgement of the evidential force of E for P, or of the competence of F in the circumstances to determine whether P, is rationally consistent with open-mindedness about Q*. Thus intuitively, a *cogent* valid argument is one where acceptance of the grounds offered for its premises, or acceptance of the competence in the circumstances of the relevant direct cognitive faculty, is consistent with anterior open-mindedness about its conclusion. Within the class of deductively valid arguments, it is only the cogent arguments, thus characterised, that are at the service of extending our knowledge. Those that fail are non-transmissive.

## 7 Basing and Support: The Arguable Dialectical Irrelevance of Supportive Holisms

Transmission-failure is not here my primary concern. But the notions latterly introduced do bear directly on one recurrent and important aspect of Leite's discussion. Leite repeatedly stresses that while there will usually be, for example, no good sense in which his belief, e.g., that he is a resident of Indiana is *based on* his belief, e.g., that he has an Indiana driving license, or his belief that the mailman just delivered a state tax demand addressed to him in an envelope bearing a Bloomington address and a logo of the Indiana State Tax dept., there is nevertheless a clear sense in which these latter considerations *support* and may be used to attest to the truth of his belief. He also points out, quite correctly, that the relevant support relation can obtain symmetrically whereas epistemic basing cannot, and that it can and does obtain very widely so that we can have a holism of mutually supportive beliefs none of which are clearly based on any of the others. Leite suggests that large classes of the beliefs that belong with the PPP have this character whereas—he thinks—scepticism gets off on the opposed notion that if our beliefs are rational at all, they have to allow of the disclosure of a determinate, as it were, linear epistemic architecture of basing relations whereby it can be made clear how they might in principle be rationally acquired from a position of ignorance or local scepticism.

These ideas require a separate detailed discussion for which I have no space here, but two observations are relevant. First, the point that certain subsets of our beliefs have this mutually supportive structure does not do anything, I suggest, to discredit the thought that, for any belief that P which is rationally held, there must at least be a *possible model* of how a thinker might justifiably move to it from a position of antecedent agnosticism about P. I think most would accept that idea from a pre-philosophical position: it is just the idea that knowledgeable belief has *always to be earned*. That idea does not exclude the possibility that the movement in question might encompass, draw on and issue in many beliefs simultaneously—some of which may be of the same broad subject matter—nor that the process may result in placement of the new belief in a holistic network of mutual support. But the thought is, I believe, very deep in our intuitive thinking about knowledge and justified belief that they are always specific cognitive achievements—earned either by individual enquiry or by (ancestral) inheritance and education or by other forms of testimony, each of which it takes the successful execution of specific processes to accomplish.

So what? Well, I have the impression, which I will not attempt to corroborate in detail now, that Leite assumes that 'The Sceptic' makes a stronger assumption: something like that for each belief knowledgeably or justifiably

held, there must be an epistemic highway leading directly via a chain of basing relations from antecedent agnosticism about it to that belief in particular and specifically. He critiques that assumption. But so far as I can see, ‘The Sceptic’ need make no such assumption. What I think Scepticism does assume is that, for any belief that is knowledgeably or justifiably held, an account—a model—has to be possible in principle of how that belief might be knowledgeably or justifiably acquired. No assumption up front need be made about the allowable structure or presuppositions of such a model—though of course, depending what is offered, challenges may arise.

Second, and by way of illustration of such a challenge, it’s relevant to note that the phenomenon of transmission failure is not restricted to inferences that purport to disclose possible *bases* for belief in their conclusions but can contaminate alleged support relations as well. Here is a toy example:

*Town clock*

You are at the Urbino Summer School in Epistemology, sitting out in the town square during the mid-morning break and wondering if you have time for another delicious Italian coffee before the 11.00 am session, which takes place in a seminar room just two minutes away. You look across at the town hall clock and note that it says 10.51. So, there is just about time. Can you trust it to be accurate though? (You know it is functioning—it chimed a few minutes ago.) You reason as follows:

E: The clock says 10.51

P: So the time is 10.51

Q: So (from E and P together) the clock is accurate on this occasion.

Thus (idiotically) reassured, you catch the barista’s eye .....

In the presence of the datum E, obtained by looking across the square, you could certainly treat P and Q as mutually supportive were you to adopt either. If you adopt either, you’ll have one of the pairs of beliefs, E + P or E + Q, that you can rationally cite in support of the third. But can you justifiably adopt either P or Q? Maybe, if you have some further collateral information of some relevant sort. But without such, it would be egregiously irrational to cite E as justification for both of the, in the presence of E, mutually supportive pair of beliefs P, and Q. There is no small (two-membered) justificatory holism here. Rather, Q is an authenticity condition for treating E as evidence for P, and it is only if you are independently justified in accepting Q that you can rationally treat E as evidence for P. Likewise and conversely, P is an authenticity condition for accepting E as evidence of Q, and it is only if you are independently justified in accepting P that you can rationally regard E as corroborating Q.

Moral: the point that there can be mutual support of the kind Leite envisages among a cluster of beliefs does not somehow pre-empt the need for a coherent model of the justified acquisition of (some) of the mutually supportive beliefs individually. And in particular, where a set of supportive beliefs involves some that are authenticity conditions for the evidential/cognitive force of procedures leading to knowledge of others, Leite's observation that there may be no determinate overall well-order of justificatory basing to be found is no assurance against the possibility of epistemic circularity. The demand for a model of how first-time knowledge of any proposition, P, can be acquired will generally require independent warrant for some collateral proposition, Q, which, in the context of one's evidence, and depending on the detail of the case may then, once accepted, be mutually supportive of P.

## 8 The Upper-Level Cartesian Paradox

The following paradox is 'upper level' in so far as it targets not knowledge and justification but whether we are in position rationally to lay claim to them. For simplicity, I focus on the case of a direct realist conception of perceptual knowledge. Any routine case of such—say glancing at my watch to check the time—will introduce a plethora of authenticity conditions which I will be required to endorse if I am rationally to lay claim to knowledge of the time as a product of the glancing. The conditions will include, for instance, many applying to the watch itself: that it has not stopped, that its hands have been rotating at their relative respectively appropriate rates, that it was correctly set at some recent relevant time, that it has not been reset or otherwise interfered with since it was correctly set. There will also many such conditions on me, the glancer—that my eyes are functioning well enough, that conditions are suitable for their effective function, that I have not unknowingly been transported to a different time zone, that I can correctly remember the significance of the position of the hands on the watch, that I am otherwise generally cognitively lucid ... Then there are external conditions that underpin the whole institution of time-keeping: that the earth has maintained a 24-hour rotational cycle, and more generally *that there is an external world at all whose local characteristics are by and large amenable to perceptual recognition.*

Here is the sting. The Corollary we noted above was that if a *rational* subject is justifiably to make any claim to knowledge via a particular method, they cannot be agnostic about but must embrace and endorse anything they recognise as an authenticity condition for accomplishment of the knowledge in question

by the relevant method in the context concerned. Some of these conditions will be ones for which they may have independent evidence. However, in every case, there will be some which, according to orthodox ideas, they can simply “take for granted,” as we are wont to say. And in laying claim to any specific, quotidian item of perceptual knowledge, you are already committed to what you will doubtless presume to be an epistemically entitled acceptance of the ‘metaphysically heavyweight’ condition expressed in italics in the paragraph immediately preceding this one.

“So,” asks our fictional Sceptic, raising one eyebrow, “how did you get to be so entitled?” Well, obviously, not by appealing, Moore-style, to its being supported by any of the legion of other specific, perceptually based claims about the external world which you consider that you know since, so the Sceptic smilingly points out, there could have been no rationally claiming knowledge of any of those, at least as the culmination of a fully reflective process, unless you could already rationally consider that you were entitled to the authenticity conditions of the relevant mode of access. And the metaphysical heavyweight will be, every time and inescapably, one of those.

Leite writes:

The fundamental point that Wright has missed is this: it is part of our ordinary, pre-philosophical position that we can sometimes have decisive evidence against a hypothesis even if we couldn’t acceptably utilise that evidence to reason our way to the falsehood of the hypothesis in conditions of explicit agnosticism, neutrality, lack of background information, or reasonable suspicion the hypothesis is true. We may therefore vindicate a claim to know something, such as the falsity of the sceptical hypothesis, by citing evidence that decisively supports the truth of what we claim to know, even though we could not arrive at that knowledge for the first time through an inference from evidence in conditions in which we are antecedently neutral, agnostic, or in doubt.

2024: 214

The fundamental point that Leite has missed is that the final sentence in that passage is simply wrong. You *cannot* rationally vindicate a claim to know the truth of the negation of the kind of sceptical hypothesis here concerned by any route involving the presumed knowledgeability of ordinary claims about the perceptible external world, no matter how holistically interlinked and non-linear the route. Specifically, you cannot rationally vindicate a claim to have come to know anything, P, including the negation of the relevant kind of sceptical hypothesis, by a route that involves citing presumed knowledgeable

support, Q, for P in any case where the negation of the sceptical hypothesis is an authenticity condition for the knowability of Q. For then you have to think that you were justified in rejecting the sceptical hypothesis *before* you started on the journey.

It's an oversight of this that spoils Moore's 'Proof'. It's also an oversight that persists in Leite's whole way of thinking about the sceptical problem. And *were* he right that the idea promoted in the quoted passage is part of our pre-philosophical position, then so much the worse for that position.

## 9 Getting Something Out of the Sceptical Paradoxes

Kreisel had the paradoxes of set-theory and the Liar, etc., in mind when he wrote the remark I quoted in opening, but his suggestion is salutary, I would suggest, for our thinking about the paradoxes of scepticism too. On Leite's approach, were it cogent, there would be nothing much to expect to learn from them except the costs of overlooking the resources of the PPP. What, if anything, more valuable or more interesting can we get out of them? Something along the following lines, I suggest.

First, that all the manifold achievements of knowledge and justified belief to which we can rationally lay claim rest on the possibility of rational acceptance of a ramified network of authenticity-conditions, collectively comprising a 'world view' both of the kinds of things there are to be known and of the kind of ways such knowledge may be accomplished.

Second, that this world view is not merely something passed on to us by our cognitive ancestors but is, at its most basic and general, necessarily *unevidenced*. Our epistemic right to it is not the product of (ancestral) cognitive achievement.

This can prove a tough point to integrate into one's thinking. As Wittgenstein expressed the predicament in *On Certainty*: "The difficulty is to realise the groundlessness of our believing" (Wittgenstein 1991: §166).

It's a difficulty because of our assumption, essential to the dialectical routines of 'The Sceptic' and, I would say, certainly part of the PPP, that all epistemically justified belief has ultimately to be earned by direct or evidential specific cognitive accomplishment. What we can "get out of" the paradoxes is the insight that that idea is actually a great mistake. If knowledge is to be possible at all, and we restrict the application of the term to true beliefs arrived at by reputable cognitive process, its achievement has to be underwritten by framework acceptances which are not knowable.

10      **So, Back to the Titular Question: Was Moore Basically Right?**

Indeed, he was not. Robert Graves memorably wrote that

To bring the dead to life  
Is no great magic.  
Few are wholly dead:  
Blow on a dead man's embers  
And a live flame will start.

GRAVES 2003: 363

It's a romantic notion and in some cases, maybe, it is true. In Moore's case, however, notwithstanding Adam Leite's strenuous pumping of the bellows, the embers are, in my opinion, just too cold.<sup>12</sup>

### References

- Bar-On, Dorit & Wright, Crispin. 2023. *Expression and Self-Knowledge*. Hoboken, NJ: John Wiley and Sons.
- AQ2 Dummett, Michael. 1975. "Wang's Paradox," *Synthese* 30 (3-4): 201-232.
- Graves, Robert. 2003. *Robert Graves: The Complete Poems*, Penguin Modern Classics series, London: Penguin Books Ltd.
- Kreisel, Georg. 1958. "Wittgenstein's Remarks on the Foundations of Mathematics," *British Journal for the Philosophy of Science* 9 (34): 135-158.
- Leite, Adam. 2024. *How to Take Skepticism Seriously*. New York: Oxford University Press.
- Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge: Harvard University Press.
- Strawson, Peter. 1963. "Freedom and Resentment," *Proceedings of the British Academy* 48: 187-211.
- Williamson, Timothy. 1996. "Cognitive Homelessness," *Journal of Philosophy* 93 (11): 554-573.
- Wittgenstein, Ludwig. 1991. *On Certainty*. Edited by G.E.M Anscombe & George Henrik von Wright. Oxford: Wiley-Blackwell.
- Wright, Crispin. 1991. "Scepticism and Dreaming: Imploding the Demon," *Mind* 100 (1): 87-116.
- Wright, Crispin 2022. "Closure and Transmission Again." In Matthew Jope and Duncan Pritchard (eds.), *New Perspectives on Epistemic Closure*, 163-189. New York: Routledge.

12 I am very grateful to Giacomo Melis and Duncan Pritchard for helpful critical comments.

# AUTHOR QUERIES

## AUTHOR PLEASE ANSWER ALL QUERIES

AQ1— Please provide missing ORCID value.

AQ2— Please note that the reference “Bar-On and Wright (2023)” is not cited in the text. Please provide an in-text citation for the reference.