

UC Irvine

UC Irvine Previously Published Works

Title

Sosa on scepticism and the background

Permalink

<https://escholarship.org/uc/item/2sf437tj>

Author

Pritchard, Duncan

Publication Date

2024

DOI

10.1007/s11098-024-02203-w

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NoDerivatives License, available at <https://creativecommons.org/licenses/by-nd/4.0/>

Peer reviewed



Sosa on scepticism and the background

Duncan Pritchard¹ 

Accepted: 29 July 2024

© The Author(s), under exclusive licence to Springer Nature B.V. 2024

Abstract

Sosa's influential work on virtue epistemology includes an intriguing proposal about background commitments, which he in turn relates to the Wittgensteinian notion of a hinge commitment. A critique is offered of Sosa's proposal, particularly with regard to how he aims to apply it to the problem of radical scepticism. In light of this critique, an alternative conception of hinge commitments is offered that enables them to play a very different role in our treatment of radical scepticism.

Keywords Epistemology · Hinge commitments · Radical scepticism · Sosa · Ernest · Virtue epistemology

1 Introductory remarks

It is always a delight to engage with Ernie Sosa's work. His contribution to contemporary philosophy, and to epistemology in particular, is unique in its ambition and scope. One always comes away from reading his writing with a fresh appreciation of the issues and often also with a new slant on debates the main contours of which had previously seemed settled. There is a lot within Sosa's epistemology that I am broadly congenial towards, but I want to focus here on one aspect of his view where I think he has, unusually, made a misstep. This concerns his treatment of radical scepticism. I will be engaging with this aspect of Sosa's work by considering what he says in his most recent monograph—Sosa (2020)—about what we might broadly term our 'background' commitments. As we will see, Sosa offers an innovative version of a relevant alternatives line in this regard, one that he thinks has application to the problem of radical scepticism. I am not so sure. I will also be considering what Sosa says in this work about Wittgenstein's (1969) discussion of the commonsense certainties that are held to play a 'hinge' role in our rational practices. According to Sosa, at least some of these hinge certainties—the ones most relevant to the problem of radical scepticism—should be construed as background assumptions. In response,

✉ Duncan Pritchard
dhpritch@uci.edu

¹ University of California, Irvine, USA

I will be arguing that this is not the right way to understand hinge commitments. Indeed, I will be suggesting that Wittgenstein's remarks in this regard are especially important to this debate because they hold the key to showing where Sosa's line on radical scepticism goes awry.

2 Sosa on the background

Let's start with Sosa's account of knowledge and the envisaged role of background commitments within that account. As will be familiar to anyone working in contemporary epistemology, Sosa understands knowledge in terms of what he terms *aptness*. Roughly, a performance is apt ('accurate because adroit') when one's success in the target endeavour is properly attributable to one's manifestation of relevant skill. As applied to the epistemic realm, we thus get the idea that knowledge is apt belief—i.e., one knows when one's cognitive success (true belief) is properly attributable to one's manifestation of relevant cognitive agency.¹

Sosa makes the novel claim that apt performances can legitimately presuppose background assumptions that the agent might not know are true—indeed, in some cases, which cannot be known at all—and which might in fact obtain simply by luck. In support of this contention, he gives the example of the performance of a nighttime baseball fielder who presupposes that the lighting in the ballpark is working properly. (Sosa, 2020, 160 & ff.) Sosa argues that this background condition can be non-negligently assumed to be in place (at least if one is given no explicit reason to consider it), even if its obtaining is just a matter of luck. More specifically, its obtaining can be *unsafe*, in the sense that such a presupposition could have very easily been false (i.e., in close possible worlds the lighting fails). Nonetheless, Sosa claims that even when unsafe such a presupposition can be entirely legitimate. Indeed, he claims that for the baseball fielder to concern himself with the obtaining of these conditions, and thereby weaken his focus on the skillful task in hand, would be negligent.²

Transposed to the epistemic realm, the corresponding idea is that so long as one's background assumptions are held non-negligently (e.g., one is not ignoring some specific reason to doubt them), then they can be operative in one's epistemic performance without this thereby undermining its aptness. In particular, it doesn't matter whether the assumptions are unknown or even that they are unsafe. So long as they are true, then one can make use of these assumptions in forming apt belief and

¹ For the key developments of Sosa's virtue epistemology, see Sosa (1991, 2007, 2009, 2015, 2020). Note that these days Sosa (e.g., 2015, 2020) fleshes out apt belief in terms of what he calls 'complete competence', which he understands in terms of 'seat', 'shape', and 'situation'. Take the competence to drive. Seat is the innermost competence (which you have even when indisposed, and thus in poor shape, such as when drunk or asleep). Then there is shape, such that seat+shape would be a competence manifest when one is not drunk, asleep, and so forth. Finally, the outermost competence is situation, as when driving conditions are appropriate for the manifestation of that competence. Complete competence is thus seat+shape+situation.

² In terms of Sosa's 'triple-S' account of complete competence (see endnote 1), the thought is thus that the background conditions that can be assumed when manifesting apt belief with complete competence are not part of the relevant features of the situation.

thus acquiring knowledge. Given that one's operative assumptions could easily be false, it thus follows, as Sosa acknowledges, that apt belief, and hence knowledge, doesn't entail the corresponding safety principle for knowledge (roughly, that one's true belief, so formed, couldn't have easily been false).³

As Sosa (2020, 126–27) notes, the conception of knowledge that results is essentially a version of relevant alternatives theory whereby in order to know we don't need to exclude all possibilities of error but just the *relevant* ones. The irrelevant error-possibilities can be legitimately assumed to be false (i.e., without one having any specific reason to think they are false).⁴ That certainly sounds right. Knowledge can be fallible, after all (i.e., acquired via fallible processes), and hence why would it be required for knowledge that all possibility of error be excluded? If that's correct, however, then the view seems to have immediate anti-sceptical import, since if any error-possibility looks like it would be irrelevant in the target sense of the term it is surely radical sceptical scenarios. Accordingly, our everyday knowledge of shoes, ships, sealing wax, cabbages and kings can be compatible with our failure to know the denials of radical sceptical hypotheses. Sosa makes the same move in applying his proposal to the sceptical problematic, except that of course he goes further to argue that we not only don't need to know our assumption that radical sceptical hypotheses don't obtain, but that this assumption can also be unsafe (i.e., it is compatible with our widespread apt belief, and thus knowledge, that such hypotheses could very easily obtain). As he puts it:

“Conscientiously enough, without negligence or recklessness, we normally assume ourselves free of skeptical scenarios. And this assumption is proper even on the rare occasions when it is true but not known to be true and even quite unsafe.” (Sosa, 2020, 159).

Notice that in taking this line Sosa embraces what is regarded by many as an unfortunate consequence of (a straightforward version of) relevant alternatives theory, which is the denial of the closure principle for knowledge.⁵ In its most plausible form, this is the principle that one knows the competently deduced entailments of one's knowledge, at least insofar as one's belief in the entailment is based on a competent deduction from one's knowledge.⁶ According to Sosa's proposal, however,

³ Sosa is, of course, one of the foremost exponents of safety and its relevance for epistemology, though as we will see he ultimately doesn't regard it as necessary for knowledge. As we will discuss presently, this is because his virtue-theoretic account of knowledge can deal with standard Gettier-style cases, and he wants to argue that non-standard Gettier-style cases, like the barn façade case, are genuine instances of knowledge (even though they also involve veritic luck), at least when construed from the virtue-theoretic conception of knowledge that he advocates. See, for example, Sosa (1999). For some other key discussions of safety-type principles, see Sainsbury (1997), Williamson (2000, *passim*), and Pritchard (2005, *passim*; 2007).

⁴ For some non-contextualist versions of relevant alternatives theory, see Dretske (1970, 1971), Stine (1976), and Nozick (1981, part 3). For some contextualist versions, see Lewis (1996), DeRose (1995), and Cohen (1999).

⁵ In previous work—e.g., Sosa (1999)—Sosa had defended the closure principle (at least in its core form anyway), so this is a departure from his earlier views. I am grateful to an anonymous referee from *Philosophical Studies* for reminding me of this fact.

⁶ See, for example, Williamson (2000, 11) and Hawthorne (2005, 29).

one can know everyday claims, be fully aware that they entail the denials of radical sceptical hypotheses, and yet be merely assuming that these denials are false. Moreover, Sosa grants that one cannot simply come to know these anti-sceptical assumptions via closure-style inferences from one's everyday knowledge, as that would be bootstrapping. Accordingly, closure has to go.⁷

I think Sosa is quite right that apt performance doesn't entail safety; indeed, I take this to be a deep insight of Sosa's work. As we can put the point, achievements (i.e., successes that are because of ability, or apt performances in Sosa's terminology) can be modally fragile (i.e., the target success could have easily been a failure). And Sosa has put his finger on a core reason why this is the case: all that matters for apt performance is that the legitimate background assumptions are true—in particular, it doesn't matter whether they are luckily true (much less whether they are known). That's why the fielder's apt performance is compatible with the fact that the stadium lights could have easily failed (in which case his attempted catch would have gone awry).⁸

Elsewhere I have expressed this general point in terms of the contrast between *intervening* and *environmental* luck and the different ways in which they impact on achievements. Intervening luck is when something gets between the subject's manifestation of relevant abilities and her success to ensure that while both success and relevant ability are on display, the success is lucky (and thus unsafe) because the former is not attributable to the latter. Think, for example, of a skilled archer firing at a target but where a dog grabs the arrow mid-flight and places it in the bulls-eye. Luck of this kind undermines achievement, as the success is not now attributable to the subject's ability (but rather, in this case, to the 'helpful' dog). Relatedly, the corresponding kind of epistemic luck, *intervening epistemic luck*, undermines cognitive achievement. This is just what we find in the standard Gettier-style cases that characteristically involve *intervening epistemic luck*.⁹

In cases of *environmental luck*, in contrast, nothing intervenes between the subject's manifestation of relevant ability and her success, but the success continues to be lucky (and thus unsafe) due to features of the environment. To return to our archer case, imagine this time that nothing gets in between the archer skillfully firing the arrow and hitting the target, but suppose that something could very easily have intervened (but didn't). Let's say that unbeknownst to our archer there were

⁷ Here is Sosa:

"We must now reject familiar and general principles of deductive closure. Default hinges are assumptions that stand fast by being faultlessly assumed. By standing fast, such competent assumptions can buttress knowledge even if they are not themselves known to be true or even safe. But we can't possibly come to know that *p* by reasoning based essentially on the sheer prior assumption that *p*, not even if the assumption is true, nor even if it is true and competently assumed to be true. You cannot get to such knowledge through any deduction that relies essentially on a sheer assumption not itself already known to be true, not when the assumption has the very same content as the conclusion to be drawn." (Sosa 2020, 136).

⁸ Actually, background epistemic luck is not quite the same as environmental epistemic luck (a notion that I will introduce presently), but the distinction doesn't matter for our present purposes. For further discussion of this distinction and its importance for epistemological debates, see Pritchard (*forthcomingc*, ch. 8).

⁹ Think, for example, of Chisholm's (1977, 105) famous 'sheep' case.

wind machines set up to knock the arrow off-course, but at the last moment they malfunctioned and hence didn't affect the shot. What is interesting about this case is that although the success is still lucky (it could very easily have been a failure), the purely environmental luck in play doesn't undermine the subject's achievement (as the success is still attributable to her manifestation of ability). Genuine achievement can thus be lucky in this specific sense, and hence unsafe, without it being any less of an achievement. And what goes for achievement also applies to cognitive achievement and the corresponding form of epistemic luck—i.e., environmental epistemic luck is compatible with cognitive achievement.¹⁰

Although Sosa doesn't express his point in terms of achievements (he prefers the language of performances), this is essentially what his view amounts to, and I think he is correct on this front. In terms of his proposal about unsafe background assumptions, Sosa's point is that this in itself would only make one's target belief lucky in the purely environmental sense just noted, and that's compatible with cognitive achievement, and thus apt belief. The fielder, for example, still manifests an achievement when he legitimately assumes that the lights are working properly, even if those lights could very easily have failed (but didn't). What goes for the fielder's achievement in the baseball game also goes for his corresponding cognitive achievement regarding any true belief that he properly forms while this background assumption is operative.

So far so good. I would, however, diverge from Sosa on three points. The first is that I don't think that apt belief, and thus cognitive achievement, amounts to knowledge. In particular, this equivalence cannot hold because knowledge, unlike cognitive achievement (/apt belief), is not compatible with lucky (unsafe) belief. That is, if one's cognitive success (true belief) could have easily been a failure, then it isn't knowledge. Knowledge excludes luck (fragility) of this kind because it excludes high levels of epistemic risk (i.e., the epistemic risk that one's belief is false).¹¹ That why, when one knows, one's basis for belief is such that one couldn't easily be wrong, which is what safety demands. In contrast, Sosa's proposed alternative picture would commit us to allowing that knowledge can coexist with high levels of epistemic risk, which I think is much too high a price to pay.¹²

Consider a variant of Fred Dretske's (1970) cleverly disguised mule case whereby the epistemic luck in play is instead purely environmental. Our agent is in fact looking at a genuine zebra in the zebra enclosure at the zoo (as opposed to Dretske's case, in which our agent is looking at a cleverly disguised mule). We can further stipulate that our agent is also non-negligently assuming, even though she has no

¹⁰ The classic example of environmental epistemic luck is, of course, Goldman's (1976) 'barn façade' case (which he in turn attributes to Carl Ginet). Note that both environmental and intervening epistemic luck are sub-varieties of veritic epistemic luck—i.e., the luck that one's belief is true, given how it was formed. I discuss the distinction between intervening and environmental epistemic luck and its epistemological significance in a number of places, but see especially (2009, 2012, 2016, 2020) and Pritchard et al. (2010, chs. 2–4).

¹¹ Clearly Sosa disagrees with this intuition, and to be fair he is not alone on this front—see, for example, Kelp (2016)—though it is certainly a minority view among contemporary epistemologists.

¹² I've engaged with this general issue with Sosa regarding the possibility of unsafe knowledge before—see, e.g., Pritchard (2009, 2012, 2016).

specific rational basis for this assumption, that there is no deception taking place (i.e., there's no attempt to pass off an animal as something different to what it is). In normal conditions, such an assumption would not only be true but also safe, in that it would not easily be the case that deception of this kind would be occurring. It is thus plausible that in normal conditions someone who believes on an ordinary perceptual basis that the creature in the zebra enclosure is a zebra, with the target 'no deception' assumption operating in the background, thereby has knowledge of what she believes. A belief so formed would be both true and safe, and this cognitive success would in addition be attributable to her cognitive abilities.

But consider now that abnormal conditions are in play so that the operative assumption, while actually true, is false in close possible worlds, such that the zebra before our agent could very easily have been an indistinguishable cleverly disguised mule (e.g., at the last minute the zoo owners elected, on a whim, to not go through with their deception). While I grant that this would not undermine the subject's cognitive achievement in forming a true belief about the zebra—this cognitive success would still be attributable to her cognitive abilities—it doesn't seem to me plausible to ascribe knowledge to the agent. The natural explanation as to why is that the luckiness of the truth of the belief, so formed, reveals that there is a high degree of epistemic risk in play. Given the fragility of the truth of the assumption at issue, it follows that forming a belief on the target basis could very easily have led to error. Think about it this way: if you were the agent in this case and came to discover how lucky it was that you formed a true belief, then would you naturally ascribe knowledge to your past self that there was a zebra in the enclosure? I don't think so. The moral to be extracted is that one's operative assumptions need to not only be true but also safe, for otherwise their lack of safety can infect one's beliefs and thereby lead to knowledge-precluding levels of epistemic risk.

As we noted above, Sosa claims that allowing that one's non-negligent background assumptions can be unknown entails the rejection of closure. In terms of our zebra example, one who knows that the creature in the enclosure is a zebra is able to competently deduce that she is not the victim of a deception regarding the animals, such as being presented with a cleverly disguised mule. Accordingly, if the latter proposition, which was merely being assumed, isn't known (and can't become known via this inferential route, as Sosa claims), then closure has to go. I think rejecting the closure principle is a mistake, and this brings to me to my second critical point about Sosa's proposal. Sosa is certainly right that one can have knowledge without knowing the background assumptions that are operative in the acquisition of that knowledge. But it doesn't follow from this point alone that closure should go.

Recall that we noted that closure in its most compelling form is a competent deduction principle, whereby one's belief in the entailed proposition is based on the competent deduction from one's knowledge. Indeed, we need to formulate closure this way for otherwise it is subject to some fairly obvious counterexamples, such as when one's belief in the entailed proposition doesn't amount to knowledge simply because it isn't based on the competent deduction but rather on a different basis entirely that is independently problematic. With closure so construed, however, one doesn't need to deny this principle in cases of local background assumptions of the kind that Sosa envisages.

Consider again our zebra case. It is plausible that one might know that there is a zebra in the enclosure without paying any thought to one's operative assumption that there is no deception taking place regarding the animals. In this sense, one can know the former while failing to know the latter. Note, however, that in normal conditions it would be odd to regard the latter as a completely groundless assumption. Sure, one might not have *explicitly* formulated any supporting reasons for it, as one needn't have considered it at all, but nonetheless if called upon to do so one surely could formulate reasons in this regard. Doesn't one ordinarily have all sorts of reasons at one's disposal for thinking that this assumption holds? Why would a reputable zoo do such a thing? What would the incentives possibly be? How would they avoid detection? Wouldn't there be penalties were they to be found out that would outweigh any possible benefits? Wouldn't such a deception, even if expertly done, soon be found out (even if the regular zoo patron wouldn't spot it, surely someone would)? And so on. Since our agent has, by hypothesis, never explicitly considered this assumption we can think of these grounds as potential rather than actual, in that a regular subject would formulate them if called upon to do.

Notice, however, that competent deduction closure is specifically cast as a diachronic principle, one that generates a belief in the entailed proposition on this basis. Accordingly, where the entailed proposition is merely an assumption, one would expect the agent to reflect on what rational basis there might be for now actively endorsing this claim. With this in mind, surely in the normal case the subject would be prompted to formulate the relevant rational support that we have just noted would be available to them on reflection. Accordingly, this would be a case where the subject *does* come to know what was previously only assumed, not because she is acquiring this knowledge simply through the competent deduction (which would invite a charge of bootstrapping), but rather because the competent deduction is prompting her to reflect on reasons that are available to her which she hasn't previously made explicit. In the normal case, then, there is no need to deny closure for cases of reasonable background assumptions, for while they might not be known prior to undertaking the competent deduction, we would expect them to become known once the competent deduction prompts one to reflect on the rational basis for this assumption.

Moreover, in the abnormal case where the subject doesn't come to know what they previously assumed—where, for example, she is oblivious to there being these sorts of rational bases available—then while it will remain true that the agent fails to know the entailed assumption, it will also now no longer be plausible to regard the agent as knowing the entailing proposition. For example, if someone really is oblivious that there might be good reasons to assume that there is no deception taking place at the zoo, even when prompted to reflect on the matter, then why would one still be inclined to attribute to them knowledge that what they are looking at is a zebra? Either way, therefore, there is no need to deny the closure principle.¹³

¹³ For further discussion of this sort of line about why the closure principle is compatible with unknown background assumptions, see Pritchard (2010, 2022b).

Note too that this story about potential rational bases for one's background assumptions not only allows one to retain the closure principle but is also very amenable to a virtue-theoretic construal. That is, one could argue that agents can have competences that enable them to virtuously identify background assumptions that are plausible candidates for having potential rational bases of the relevant kind. This would thus ensure that such virtuous subjects tend to embrace only the reasonable background assumptions. Moreover, there would be an obvious benefit to having such a competence within one's cognitive economy given the intellectual resources involved in formulating rational bases for background assumptions when this isn't required for one's wider cognitive functioning. Indeed, this virtue-theoretic way of thinking about reasonable background assumptions would dovetail with much of Sosa's thinking in this regard—for example, the idea that there is a kind of intellectual negligence involved in reflecting on the rational basis of one's background assumptions when there is no specific reason to do so.

In any case, the key point is that there isn't the rationale for denying closure on the basis of local background assumptions that Sosa imagines. I think part of the reason Sosa doesn't recognize this point is that he doesn't distinguish between the local background assumptions and the kind of global background assumptions involved in the denials of radical sceptical hypotheses. The latter are very different from the former as there simply isn't the same kind of potential rational support available for them. There is a good reason for this, as one core feature of such hypotheses is that by calling one's beliefs *en masse* into question they thereby also bring one's rational support as a whole into question too. Since the sceptical hypothesis entails global error, so the relevant background assumption involves the denial of such global error.

The same is not true in the local case. The relevant error-possibility—that there are cleverly disguised mules in the zebra enclosure, say—is a local error-possibility in the target sense, as it doesn't call one's beliefs in general into question. Similarly, the relevant background assumption involves only the denial of such localized error. This is why the possibility that there might be deception taking place at this zoo is entirely compatible with the sort of reasons we cited above as potential support for dismissing this possibility, such as that reputable zoos don't engage in such activities. In contrast, the possibility that one is, say, a brain-in-a-vat being systematically 'fed' deceptive experiences would undermine the corresponding reasons one might cite for rejecting this scenario. (For example, that the current state of technology is not sophisticated enough for there to be such envatted brains is undermined by the sceptical scenario itself, as of course the victim of such a scenario might well be engineered to have false beliefs about the degree of advancement of contemporary technology).

The crux of the matter is that it is global, rather than local, background assumptions that raise a problem for closure, since they seem to preclude the possibility of one's belief in the entailed proposition rising to the status of knowledge in the context of a relevant competent deduction. In particular, in a case where one's background assumption is the denial of a radical sceptical scenario it seems as if the only closure-retaining option is to suppose that one can come to know this claim via the competent deduction itself, which is bootstrapping.

This brings me to my third critical point about Sosa's proposal, which is the implausibility of treating our commitment to the denials of radical sceptical hypotheses as being unknown, but epistemically legitimate nonetheless, background assumptions. The rationale for allowing that one's local background assumptions can remain unknown is that normally we simply don't need to know them, and thus there is a kind of intellectual negligence involved in putting in the cognitive work to know them when this isn't required. Indeed, like the fielder who assumes that the lights are functioning okay, it's not just that putting in work that one doesn't need to do wastes one's cognitive resources, but it also might undermine one's overall intellectual performance, such as by distracting one from something more important.

Sosa's common treatment of both local and global (anti-sceptical) background conditions takes it for granted that this sort of negligence-avoiding line about the propriety of local background conditions also applies to the global background conditions. Given that there is an in principle barrier to gaining rational support for the latter that doesn't apply to the former, however, I think it is clear that this supposition does not hold water. In particular, in assuming that a radical sceptical scenario doesn't obtain we aren't merely setting the issue of rational support to one side to ensure that our proper focus is on other, more important, cognitive matters as we might when it comes to local background assumptions. On the contrary, we already know that no such rational support is going to be available.

Although Sosa's discussion offers a common treatment of both global and local background assumptions, it is significant that when he turns his attention to the latter he brings in considerations that simply don't feature in his discussion of the former. In particular, the point about avoiding intellectual negligence morphs into a subtly—but, I think, importantly—different claim about the necessity of our anti-sceptical assumptions. Consider this passage:

“Because we are essentially rational animals, we have *no real option* on how to proceed cognitively over the enormous span of the animal knowledge we rely on in any ordinary day. [...] Any particular human, no matter how inept or adept, will be bound to proceed within some parameters of perceptual belief acquisition. But no one is likely to *entirely* avoid illusion and every other perceptual error.” (Sosa, 2020, 138, *italics in the original*)

I think we can disentangle two thoughts in play here. The first is that we can't help but be committed to the denials of radical sceptical hypotheses if we are to be rational creatures who undertake inquiries and so forth, given that the alternative would undermine our rational practices altogether. While this claim is plausible, it doesn't have the anti-sceptical import that Sosa clearly thinks that it does. After all, even the radical sceptic would surely grant that this is the case. Indeed, this is precisely why our epistemic situation is held to be so paradoxical, in that we are both obliged to have anti-sceptical commitments and yet we also recognize that we have no rational basis for these commitments. Moreover, notice that the non-optionality of our anti-sceptical commitments in itself does not provide us with any rational basis for having them. This is why these commitments are just assumptions.

The second thought in play is the idea that infallibility in our epistemic practices is not a realistic goal and hence it is reasonable to countenance the possibility of some error in our beliefs. This claim also seems plausible, but again it doesn't have the anti-sceptical import that Sosa imagines. The radical sceptic is not appealing to infallibilist intuitions to motivate their radical scepticism. Their contention is not that our beliefs fail to meet an unrealistically high epistemic standard but rather that they fail to meet any epistemic standard at all, even a low one. That's just the point that's raised by our inability to exclude radical sceptical scenarios, as the idea is that absent such an exclusion, we have no rational basis for our everyday beliefs at all. (For example, one's apparent perception that such-and-such is the case is no rational basis for the target belief if one can't exclude the possibility that one is a brain-in-a-vat who is being programmed to have indiscriminable, but deceptive, experiences). Making this point about infallibility in the context of the radical sceptical problem thus diverts our attention from what is really at issue.

The upshot is that Sosa doesn't really have anything to say about why our background anti-sceptical assumptions are epistemically legitimate even though unknown. While there is a story to be told in this regard that applies to the local background assumptions, it simply doesn't carry over to the global anti-sceptical background assumptions. The problem, however, is that without this account then Sosa's response to radical scepticism fails to offer us an intellectually satisfying response to the problem.

3 Sosa on hinge commitments

So where did Sosa go wrong in his thinking on radical scepticism? I think that the answer to this may lie in his understanding of the Wittgensteinian notion of our 'hinge' certainties. As we will see, Sosa (2020) relates his discussion of background assumptions directly to this notion, thereby indicating that he regards at least one core type of hinge commitment as a kind of groundless, but nonetheless epistemically legitimate, background assumption. Sosa (2020) concludes his latest monograph with an extended discussion of G. E. Moore's (e.g., 1925; 1939) treatment of our everyday certainties and the radically different way that Wittgenstein (1969) conceives of them, so let's start there.

As is now familiar, Moore was intrigued by the special status of our everyday certainties: the kinds of claims that we take to be obviously true. Since Moore prized commonsense over any philosophy that might challenge it, he therefore thought that they were knowledge, even though he granted that it was not clear what epistemic basis grounded this optimal certainty that we have in these propositions.¹⁴

¹⁴ Moore is often taken to have been the first to identify this particular issue regarding our everyday certainties. As I've argued elsewhere, however, he is in fact responding to a common philosophical issue of his time, one that goes back at least as far as Newman (1979 [1870]). As Newman makes clear, the underlying issue here arises out Locke's epistemology, which demands, quite plausibly, both that our epistemic support should be proportionate to our level of commitment (so optimal certainty demands optimal levels of epistemic support), and that we should be able to make explicit what this epistemic support is and how it performs this role. For discussion, see Pritchard (*forthcomingb*).

Wittgenstein, famously, took a very different line. He argued that our everyday certainties play a ‘hinge’ role in our rational practices by providing the framework relative to which we undertake rational evaluations. As a result, these hinge commitments are not themselves subject to rational evaluation, and hence they are unknown.¹⁵

Sosa finds both proposals problematic and offers a nuanced discussion of the various issues in play as he explores the options. In particular, he rejects Wittgenstein’s own account of hinge commitments, which is important, given that I will be defending Wittgenstein in this regard. What is salient for our purposes, however, is the conclusion that he draws, which is twofold. The first part is to maintain that we can know at least some of these hinge certainties via a special kind of epistemic competence that enables us to make the relevant discriminations in the target domain. Sosa concedes that we have little understanding of how such a competence might work, but he claims that since we don’t have much understanding of how any competences work this is not a substantial concern. As always, Sosa’s unusual candor when it comes to evaluating the merits and demerits of features of his position is refreshing.

But there is a further twist, as Sosa further argues that this approach won’t work for all our hinge certainties, especially the ones most relevant to the sceptical problematic. His line here, however, references the previous discussion of our background commitments, as he claims that we need to think of these particular hinge commitments not as beliefs at all but as assumptions. Crucially, however, he also points out that the virtue epistemology that he offers provides the means to make sense of how such assumptions can be epistemically legitimate, thereby blocking radical sceptical worries. Here is how his discussion concludes:

“We must go beyond beliefs for the desired account [*of our hinge commitments*]. [...] the desired explanation requires not background beliefs but background assumptions, and that is why we should welcome the route provided by virtue epistemology to the plausibility of such assumptions. Our normal human epistemic flourishing requires not just fast-standing human belief systems acquired *sans* ratiocination. It requires also fast-standing assumptions that incur no relevant epistemic negligence or recklessness, as they block inappropriate skeptical concerns.” (Sosa, 2020, 220–21)

Indeed, this particular line on hinge commitments was prefigured in an earlier discussion:

“That opens up a further interpretative option concerning Wittgenstein’s hinges. Some at least of these, some important ones pertaining to radical skep-

¹⁵ There is now a wealth of literature examining Wittgenstein’s *On Certainty*. For some of the main texts in this regard, see Strawson (1985), McGinn (1989), Williams (1991), Stroll (1994), Moyal-Sharrock (2004), Wright (2004), Coliva (2015; 2022), and Schönbaumsfeld (2016). Naturally, in what follows I will be putting forward my own favoured interpretation, as developed, for example, in Pritchard (2015, *passim*). For a recent survey of this literature, see Pritchard (2016). For a comparative treatment of Moore and Wittgenstein on our everyday certainties, see Coliva (2010) and Pritchard (2021a, 2021b, 2021c).

ticism, can now be viewed as default assumptions properly shared by humanity in our joint life and action. We are fellow pilots leading lives that are intricately interrelated. Our success can then be individually and collectively creditable even in the absence of security. So, it can be based just on default assumptions that need only be true, that need not be known to be true, nor even safe.” (Sosa, 2020, 183–84)¹⁶

We thus get the idea that our hinge commitments, at least insofar as they are relevant to the problem of radical scepticism, should be understood as non-negligent default assumptions. So construed, the fact that they are unknown needn’t entail any standing invitation to radical scepticism. We can’t know that we are not the victim of radical sceptical scenarios, but we recognize that we are committed to regarding them as false in all our intellectual endeavors. We are thus committed to assuming their falsity. Nonetheless, just as background assumptions in general can be epistemically legitimate even when groundless provided that they are undertaken in a non-negligent fashion, so our hinge assumptions that we are not victims of sceptical scenarios can also be epistemically legitimate.

I want to suggest that this is not the right way to think about hinge commitments. In particular, in understanding them in this fashion Sosa fails to recognize how they in fact undermine the radical sceptical challenge. To conceive of our hinge commitments as groundless assumptions is already to concede too much to the sceptic, for it implies that the radical sceptic is at least right to the extent that they have identified an important cognitive lacuna, in that there is a crucial class of fundamental commitments that are in the market for knowledge but which we simply cannot know.¹⁷ What Wittgenstein is showing us, however, is that it is the product of a faulty philosophical picture that we would even expect our hinge commitments to be known, much less to suppose that our failure to know them indicates a cognitive lack on our parts. As we will see, this is also related to why Wittgenstein is quite explicit that our hinge commitments are not assumptions or, indeed, any kind of propositional attitude that could be properly assessed for being reasonably or unreasonably held.

In order to unpack these points, we need to look again at Wittgenstein’s account. Wittgenstein is primarily interested in a particular kind of certainty that is associated with one’s worldview as a whole. What Wittgenstein is trying to get us to recognize is that in acquiring a worldview at all it is necessary that certainty permeates one’s commitment to the worldview. Wittgenstein repeatedly emphasizes the

¹⁶ See also this passage:

“The kinship between some at least of Wittgenstein’s hinges and our own background conditions comes again to the fore. Both are assumed by default to buttress our coherent commonsense perspective as well as the particular judgments that we make from that perspective. Hinges can thus enable our body of knowledge without our knowing them to be true, and without our reaching them as conclusions (through “ratiocination”). As with hinges, background conditions can be assumed to obtain, with assumptions that stand fast for us, and can properly do their enabling even when they are just assumed to be true, and even when that assumption is both unsafe and in ignorance.” (Sosa 2020, 130, n9).

¹⁷ Sosa is not alone on reading Wittgenstein in *On Certainty* as exposing an important cognitive limitation on our parts. As I discuss in Pritchard (2023), this is a common misinterpretation, one found, for example, in the work of Wright (2004).

holistic nature of this commitment (e.g., OC, §140–44). I've elsewhere termed this the *über hinge commitment*, and it is essentially the certainty that one's worldview is broadly correct—i.e., that it is not radically or fundamentally in error.¹⁸ One's *über hinge* certainty cannot be itself rationally evaluable, since it is rather a prerequisite of being in a position to undertake rational evaluations at all. Accordingly, it is a certainty that is visceral, 'animal', 'primitive'. (OC, §§475 & 359) It is not (and could not be) reasonable or unreasonable, but is simply there, 'like our life'. (OC, §559) Rather than being grounded in our reasons it is instead manifest in our actions which reveal our complete conviction in our worldview. (OC, §204).

This is why we can make no sense of the universal doubt envisaged by the radical sceptic. As Wittgenstein puts it: "A doubt that doubted everything would not be a doubt." (OC, §450) The point is that in order to undertake the kind of rational evaluation involved in doubting already requires the certainty in the worldview that precludes radical doubt. As a result, putative presentations of radical sceptical doubt will either be interpreted as merely performative (as when the philosopher says that she is embracing sceptical doubts but her actions reveal her continued certainty in the worldview) or else will be interpreted not as a rational undertaking at all but as an indication of mental illness. (OC, §§71–75).

Once one grants that the *über hinge* commitment is required for all rational evaluation, then it follows that this overarching certainty will apply to the denial of any claim that calls one's worldview into question as a whole. This is the sense in which we are all hinge committed to the denials of radical sceptical hypotheses, even before we encounter the details of any particular radical sceptical hypothesis. A key innovation in Wittgenstein's discussion is to recognise that one's visceral *über hinge* certainty in one's worldview can also be represented in one's certainty in specific commonsense propositions, such as that (in normal circumstances) one has hands (e.g., OC, §1), what one's name is (OC, §425), and the language that one is speaking (OC, §158). Wittgenstein suggests that these everyday certainties are quite different from other kinds of commitments that we have in that, as manifestations of the *über hinge* commitment, they are not grounded in reasons. The reason why they manifest one's *über hinge* commitment is that they are the kind of claim where doubt would call one's worldview as a whole into question. As Wittgenstein puts it, a doubt here would "drag everything with it and plunge it into chaos." (OC, §613).

Note that on this account of hinge commitments they are not just any background commitment we might have. Rather, they need to be the kind of commitment that bears the appropriate relationship to the *über hinge* commitment, such that doubt of the former would entail doubt of the latter. This means that they are almost exclusively concerned with mundane everyday items of commonsense, such as concerning whether one has hands or what one's name is. Indeed, it is odd to think of most of our hinge commitments as part of the background at all, as they seem to rather be very much in the foreground. Wittgenstein's point is not that these are tacit commitments we have, but rather that they are so familiar and mundane that we don't fully appreciate their odd epistemic status. In particular, we don't notice that they are very different from other sorts of everyday claims in that they are not themselves subject

¹⁸ See, especially, Pritchard (2015, part 2).

to rational evaluation. As Wittgenstein puts it, they “lie apart from the route travelled by inquiry.” (OC, §88) That’s not the same as being tacit background claims.

It follows that our hinge commitments are not theoretical claims, much less specifically philosophical in nature. Propositions of this kind simply don’t bear the relevant relationship to the über hinge commitment nor, relatedly, do they inspire the visceral certainty that is distinctive of our hinge commitments. This is salient to our discussion of Sosa because at least some of the putative hinge claims that he has in mind (and which he claims can be known) look very theoretical in nature. This includes the claim that there is an external world, in the sense of a world that is metaphysically independent of all human minds. This is not a hinge commitment on the Wittgensteinian picture, however, not least because it is not a core element of commonsense that manifests our über hinge commitment.

Indeed, matters are even worse on this front as Wittgenstein is quite clear that he regards a proposition such as ‘There is an external world’ as being simply meaningless, as constituting a philosophical misuse of ordinary language.¹⁹ (OC, §35–37) Given that our hinge commitments are contentful, it follows that it cannot be a hinge commitment. Even if one resists this particular contention, however, the point just made that such philosophical claims cannot be hinge commitments still applies. Either way, then, at least some of the putative hinge commitments that Sosa has in mind are not plausible candidates for this description once we understand this notion correctly.²⁰

An important consequence of our account of hinge commitments is that they are not to be understood as assumptions, or indeed as any kind of propositional attitude that can be evaluated (as assumptions clearly can) as reasonable or unreasonable. Indeed, Wittgenstein is quite explicit that he doesn’t think of our hinge commitments as assumptions, as he is at pains to contrast the former with the latter (e.g., OC, §343). The crux of the matter is that to treat them as assumptions is to completely misunderstand the kind of propositional attitude in play. Assumptions can be assessed as reasonable or unreasonable because they reflect our epistemic agency: we can be held to account for the assumptions that we elect to make. Our hinge commitments are not products of our epistemic agency, however, but are rather completely brute in nature. This is why they can’t be assessed for whether they are reasonable or unreasonable. Relatedly, our hinge commitments cannot be discharged in the way that assumptions can, nor is it possible to be agnostic about the truth of one’s hinge commitments in the way that one can be about one’s assumptions, hypotheses, presuppositions and so forth. Indeed, as our actions reveal, one remains completely convinced in one’s hinge commitments even when one comes to realise that they are not grounded in reasons.

¹⁹ For further discussion of this exegetical point, see Williams (2004, 2018) and Pritchard (2015, part 2).

²⁰ Although it doesn’t bear directly on our discussion here, it is worth noting that there is a second class of proposition that is often falsely classed as a hinge commitment, over and above the kind of philosophical claim just mentioned. This is the kind of tacit everyday commitment that we might call (following Greco 2016) ‘common knowledge’. As I discuss in Pritchard (2022a), while common knowledge does share some properties with our hinge commitments, they nonetheless concern a distinct kind of everyday certainty, one that doesn’t bear the required relationship to the über hinge commitment.

Recognising the distinctive nature of the propositional attitude in play in our hinge commitments is crucial to understanding that it simply cannot be an assumption, or indeed any kind of propositional attitude that can be assessed for rationality. This is why I've elsewhere argued that our hinge commitments cannot be beliefs in the specific sense of belief that is relevant to epistemology—i.e., that propositional attitude that is a constituent part of knowledge (*K-apt belief*). Unlike the more permissive folk notion of belief (which covers a range of propositional attitudes, not all of them reasons-responsive), *K-apt belief* bears some basic conceptual connections to reasons and truth. In particular, one cannot *K-apt* believe that *p* while recognising that one has no rational basis for the truth of *p*. And yet we've noted that one does remain hinge committed even after recognising the groundless nature of one's hinge commitments.²¹

One interesting consequence of the fact that hinge commitments are not *K-apt* beliefs is that even though they are essentially unknown they are also compatible with the closure principle. Recall that we noted that this principle needs to be formulated such that one acquires a belief on the basis of the competent deduction in question. The notion of belief that is in play here, however, is clearly *K-apt belief*, given that the subject is meant to be acquiring knowledge from this competent deduction. In cases where the entailed proposition is a hinge commitment, however, one simply cannot *K-aptly* believe the target proposition, much less come to *K-apt* believe it on such a rational basis as a competent deduction. The putative failure of closure when it comes to our inability to know the denials of radical sceptical hypotheses that we noted above thus doesn't materialise, at least so long as we follow Wittgenstein in thinking of our hinge commitments as involving a distinctive kind of propositional attitude.

That the propositional attitude involved in our hinge commitments excludes the possibility that they are assumptions is obviously in conflict with Sosa's account of them as legitimate background assumptions. Even so, one might think that Sosa's proposal is more attractive than what Wittgenstein has in mind. Wouldn't it be better to think of our anti-sceptical commitments as assumptions that can, accordingly, be potentially reasonably held, as opposed to merely being the product of a visceral overarching certainty that is unresponsive to rational considerations, and which thus cannot be evaluated as reasonable or unreasonable?

In order to appreciate why Wittgenstein's approach is preferable, we need to understand the important diagnostic story that he is offering regarding the problem of radical scepticism. The point that Wittgenstein is trying to convey with his account of hinge commitments is meant to be a general one about the nature of rational evaluation rather than a contingent claim about our specific rational practices. (As he emphasises in a number of places, the point he is making is one of 'logic' (e.g., OC, §342)). In particular, he is trying to get us to see that the very idea of fully general rational evaluations is simply incoherent, as all rational evaluation presupposes arational hinge commitments. It is thus not an incidental lack on our parts that we fail to know our hinge commitments, as if this were a cognitive limitation that unfortunately we cannot overcome. Rather it is in the very nature of hinge

²¹ For further development of these points, see Pritchard (2015, part 2; 2018, *forthcominga*).

commitments that they are unknowable. Accordingly, while one does not know one's hinge commitments, it is not as if one is *ignorant* of them, as they are not even in the market for knowledge.²² It follows that our inability to know our hinge commitment is no more reflective of a cognitive limitation on our parts than our inability to conceive of a circle-square reflects an imaginative limitation.

Moreover, Wittgenstein is also showing us how this faulty conception of our rational practices is not the product of our ordinary ways of thinking in this regard but rather the result of dubious philosophy masquerading as commonsense. This entails that there is nothing paradoxical about our epistemic situation as the radical sceptic claims, as the putative tension in our thinking in this regard is not arising out of our ordinary practices but rather out of faulty philosophical presentations of them. In particular, our ordinary rational practices, with their essentially local rational evaluations, are perfectly in order. That we are unable to rationally ground our worldview as a whole does not reveal a limitation of those practices, much less a failing, but is rather built into the nature of what a rational practice must be.

Sosa's conception of the sceptical problematic is one on which the radical sceptic exposes a cognitive lacuna that is to be plugged by means of a groundless global background assumption, albeit one that is made epistemically legitimate by the fact that it is a necessary for living a productive intellectual life. If Wittgenstein is right, in contrast, there is no lacuna in the first place, as the sceptical problematic is rather arising from a faulty philosophical picture that should be rejected out of hand. This is the sense in which Wittgenstein is offering a diagnostic story that dissolves the putative radical sceptical paradox.²³ In failing to appreciate this radical dimension to Wittgenstein's proposal, Sosa adopts a conception of hinge commitments that ends up making important concessions to radical scepticism.²⁴

Funding There is no funding associated with this paper.

Declarations

Conflict of interest The author confirms that they have no relevant financial or non-financial interest to disclose.

²² That one doesn't count as ignorant of a proposition that is not in the market for knowledge reveals that there is more to ignorance than merely a lack of knowledge. For further discussion of this point, including other cases where ignorance comes apart from lack of knowledge, see Pritchard (2021a, 2021b).

²³ At any rate, it dissolves that form of radical scepticism which trades on universal rational evaluations. See Pritchard (2015, *passim*) for discussion of a distinct variety of radical scepticism that doesn't trade on this claim.

²⁴ I am, as ever, grateful to Ernie Sosa for discussion of these issues. Thanks also to the very helpful report of an anonymous referee from *Philosophical Studies*. This paper was written while a Senior Research Associate of the *African Centre for Epistemology and Philosophy of Science* at the University of Johannesburg.

References

- Chisholm, R. (1977). *Theory of knowledge* (2nd ed.). Prentice-Hall.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13, 57–89.
- Coliva, A. (2010). *Moore and Wittgenstein: Scepticism, certainty and common sense*. London: Palgrave.
- Coliva, A. (2015). *Extended rationality: A hinge epistemology*. Palgrave Macmillan.
- Coliva, A. (2022). *Wittgenstein rehinged*. Anthem.
- DeRose, K. (1995). Solving the Skeptical problem. *Philosophical Review*, 104, 1–52.
- Dretske, F. (1970). Epistemic operators. *Journal of Philosophy*, 67, 1007–1023.
- Dretske, F. (1971). Conclusive reasons. *Australasian Journal of Philosophy*, 49, 1–22.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73, 771–791.
- Greco, J. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge University Press.
- Greco, J. (2016). Common knowledge. *International Journal for the Study of Skepticism*, 6, 309–325.
- Hawthorne, J. (2005). The case for closure. In E. Sosa & M. Steup (Eds.), *Contemporary debates in epistemology* (pp. 26–43). Oxford: Blackwell.
- Kelp, C. (2016). Epistemic Frankfurt cases revisited. *American Philosophical Quarterly*, 53, 27–37.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74, 549–567.
- McGinn, M. (1989). *Sense and certainty: A dissolution of scepticism*. Blackwell.
- Moore, G. E. (1925). A defence of common sense. In H. Muirhead (Ed.) *Contemporary British Philosophy* (2nd series) Allen & Unwin.
- Moore, G. E. (1939). Proof of an external world. *Proceedings of the British Academy*, 25, 273–300.
- Moyal-Sharrock, D. (2004). *Understanding Wittgenstein's on certainty*. Palgrave Macmillan.
- Newman, J. H. (1979 [1870]). *An Essay in Aid of a Grammar of Assent*, Notre Dame, IN: University of Notre Dame Press.
- Nozick, R. (1981). *Philosophical explanations*. Oxford University Press.
- Pritchard, D. H. (Forthcomingc). *Tempting fate: Risk, luck and the meaning of life*. Princeton, NJ: Princeton University Press.
- Pritchard, D. H. (Forthcominga). Beliefs, delusions, hinge commitments. In *Synthese*.
- Pritchard, D. H. (Forthcomingb). Quasi-fideism and virtuous anti-evidentialism: Wittgenstein and Newman on knowledge and certainty. In J. Milburn & F. Aquino (eds.), *Newman and Contemporary Philosophy* Routledge.
- Pritchard, D. H. (2005). *Epistemic luck*. Oxford University Press.
- Pritchard, D. H. (2007). Anti-luck epistemology. *Synthese*, 158, 277–297.
- Pritchard, D. H. (2009). Apt performance and epistemic value. *Philosophical Studies*, 143, 407–416.
- Pritchard, D. H. (2010). Relevant alternatives, perceptual knowledge, and discrimination. *Noûs*, 44, 245–268.
- Pritchard, D. H. (2012). Anti-luck virtue epistemology. *Journal of Philosophy*, 109, 247–279.
- Pritchard, D. H. (2015). *Epistemic angst: Radical skepticism and the groundlessness of our believing*. Princeton University Press.
- Pritchard, D. H. (2016). Epistemic risk. *Journal of Philosophy*, 113, 550–571.
- Pritchard, D. H. (2018). Disagreement, of belief and otherwise. In C. Johnson (Ed.), *Voicing dissent: The ethics and epistemology of making disagreement public* (pp. 22–39). Routledge.
- Pritchard, D. H. (2020). Anti-risk virtue epistemology. In J. Greco & C. Kelp (Eds.), *Virtue-theoretic epistemology: New methods and approaches* (pp. 203–24). Cambridge University Press.
- Pritchard, Duncan. (2020). Scepticism and certainty: Moore and Wittgenstein on common sense and philosophy. In R. Peels & R. van Woudenberg (Eds.), *The Cambridge companion to common-sense philosophy* (pp. 247–264). Cambridge University Press. <https://doi.org/10.1017/9781108598163.012>
- Pritchard, D. H. (2021a). Ignorance and inquiry. *American Philosophical Quarterly*, 58, 111–123.
- Pritchard, D. H. (2021b). Ignorance and normativity. *Philosophical Topics*, 49, 225–243.
- Pritchard, D. H. (2022a). Hinge commitments and common knowledge. *Synthese*. <https://doi.org/10.1007/s11229-022-03647-5>
- Pritchard, D. (2022b). In defence of closure. In M. Jope & D. Pritchard (Eds.), *New perspectives on epistemic closure* (pp. 113–130). Routledge. <https://doi.org/10.4324/9781003104766-8>

- Pritchard, D. H. (2023). Hinge commitments and trust. *Synthese*. <https://doi.org/10.1007/s11229-023-04378-x>
- Pritchard, D. H., Millar, A., & Haddock, A. (2010). *The nature and value of knowledge: Three investigations*. Oxford University Press.
- Sainsbury, R. M. (1997). Easy possibilities. *Philosophy and Phenomenological Research*, 57, 907–919.
- Schönbaumsfeld, G. (2016). *The illusion of doubt*. Oxford University Press.
- Sosa, E. (1991). *Knowledge in perspective: Selected essays in epistemology*. Cambridge University Press.
- Sosa, E. (1999). How to defeat opposition to moore. *Philosophical Perspectives*, 13, 141–154.
- Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. 1). Oxford University Press.
- Sosa, E. (2009). *Reflective knowledge: Apt belief and reflective knowledge* (Vol. 2). Oxford University Press.
- Sosa, E. (2015). *Judgment and agency*. Oxford University Press.
- Sosa, E. (2020). *Epistemic explanations: A theory of telic normativity, and what it explains*. Oxford University Press.
- Stine, G. (1976). Skepticism, relevant alternatives and deductive closure. *Philosophical Studies*, 29, 241–261.
- Strawson, P. F. (1985). *Skepticism and naturalism: Some varieties*. Columbia University Press.
- Stroll, A. (1994). *Moore and Wittgenstein on certainty*. Oxford University Press.
- Williams, M. (1991). *Unnatural doubts: Epistemological realism and the basis of scepticism*. Blackwell.
- Williams, M. (2004). Wittgenstein's refutation of idealism. In D. McManus (Ed.), *Wittgenstein and scepticism* (pp. 76–96). London: Routledge.
- Williams, M. (2018). Illusions of doubt: Wittgenstein on knowledge and certainty. In D. Machuca & B. Reed (Eds.), *Skepticism: from antiquity to the present* (pp. 481–505). London: Bloomsbury.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.
- Wittgenstein, L. (1969). G. E. M. Anscombe, G. H. von Wright, (tr.) D. Paul, G. E. M. Anscombe (eds.), *On Certainty* Blackwell.
- Wright, C. (2004). Warrant for nothing (and foundations for free)? *Proceedings of the Aristotelian Society*, 78, 167–212.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.