

## REPLIES TO MY CRITICS

ERNEST SOSA  
Department of Philosophy  
Rutgers University  
ernsosa@rci.rutgers.edu

**SUMMARY:** This paper is a response to the four critics of *A Virtue Epistemology* (2007). It responds to Claudia Lorena García, Miguel Ángel Fernández, Jonathan Kvanvig, and Ram Neta, in that order.

**KEY WORDS:** apt belief, epistemic competence, knowledge, skepticism, virtue epistemology

**RESUMEN:** Este artículo es una respuesta a los cuatro críticos de *A Virtue Epistemology* (2007). Ofrece respuestas a Claudia Lorena García, Miguel Ángel Fernández, Jonathan Kvanvig, y Ram Neta, en ese orden.

**PALABRAS CLAVE:** competencia epistémica, conocimiento, creencia apta, epistemología de las virtudes, escepticismo

My grateful thanks to these critics and friends for their attention to my work, and for the obvious care taken over their comments. These have prompted me to develop and improve the account in the book.

My responses below are designed to be self-standing. I have reproduced the questions and objections that I will take up, so that they and my responses can be understood from what I have to say.

### 1. *Reply to Claudia Lorena García*

Claudia Lorena García makes many good points. Her main critique, however, concerning skepticism and dreams, is off target, since the varieties of skepticism in her paper are absent from my book. If a skeptic challenges us to prove ex nihilo that we see a hand, or a fire, we should decline. Nor should we take the bait when he offers us as premises just obvious a priori truths and facts about our own current subjective states. Even given such premises, pure reason (deductive or inductive) will not yield nearly enough of what we believe ordinarily. Our knowledge of hands and fires is not explicable in terms of such reasoning from the given. Far from accepting such skeptical challenges, we should decline the presuppositions from which they are launched. We should reject the sheer assumption that only reasoning from the foundational given will provide or explain whatever knowledge we may enjoy.

I take more seriously a more serious skeptical challenge. Here's how I would distinguish it from those of lesser concern.

### 1.1. Varieties of Skepticism

Skepticism comes in various forms. Pyrrhonian skepticism is a determined suspension of judgment, through tropes that counter whatever reasons may be offered in favor of belief. The Academic skeptic, by contrast, is more assertive, if only by claiming that we know nothing, either in general, or in some large department of our supposed knowledge: the external world, for example, or other minds, or morality. Any attempt to refute the fully global claim therefore begs the question.

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

Such varieties of skepticism are less problematic than the variety of main interest in the book. My skeptic spots a commitment ostensibly at the heart of common sense, and shows how it entails that we know very little, either in general or in some main department.

### 1.2. Skepticism, Sensitivity, and Safety

Consider, for example, the idea that a belief constitutes knowledge only if *sensitive*, only if it satisfies the following condition: *had its propositional content been false, it would not have been held by the believer*. In order for you to know that you see a hand, your belief must be *sensitive*, in that *had you not seen a hand, you would not have believed that you saw one*. If this sensitivity requirement is indeed among our core commitments, the skeptic is in a good position. Belief that one is not radically misled cannot be sensitive. Indeed, skeptical scenarios are framed to secure precisely this result. If you were now a brain in a vat being fed experiences as if you enjoyed normal perception, that would not stop you from believing that you were *not* radically deceived.

Various responses to such skepticism have been developed over many decades, including closure-denying tracking approaches, and contextualist semantic ascent. A rich and subtle dialectic can be

found in an extensive and still active literature. No such sensitivity-influenced approach is fully satisfactory, though of course that is not something I can argue in the space available here.

There is anyhow, in my view, a better response to the skeptic. It begins by noting that subjunctive conditionals do not validly contrapose,<sup>1</sup> which suggests a requirement of safety rather than sensitivity, as follows.

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

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

### 1.3. Why the Dream Scenario Is Special

Our line of reasoning is effective against radical scenarios, such as the brain in a vat, the evil demon, the Matrix, and so on. Only the dream scenario stands apart. Dreams being so common, the possibility that one dreams is not outlandish. Therefore we cannot defend the safety of our belief that we are awake by adducing how remote is the possibility that we go wrong in so believing. Too easily for comfort might we have been not awake but only dreaming.

García quite reasonably asks: “*Which* dreaming scenario?” The scenario that life is but a dream —*all of life*— is a dream scenario, but it is hardly less outlandish than BIV or evil demon scenarios.

<sup>1</sup> For example, if water flowed out of my kitchen faucet, it would be false that water flowed while the main house valve was closed. But consider the contrapositive of this conditional: namely, that if water flowed out the faucet while the main valve was closed, then water would not flow. This contrapositive would be obviously false, showing contraposition to be invalid for such conditionals.

That is, accordingly, *not* the dream scenario of special interest to us. The more relevant dream scenario is the one posed by Descartes when he wonders, as he sits before a fire, whether he is *then* just dreaming. Our dream scenario is one that arises for any arbitrary case where we consider whether our ostensible perceptual knowledge—of a fire, say, or a hand—is real *on that occasion*.

Such a dream scenario has a distinctive importance by comparison with the familiar radical scenarios. Unlike the others, it is not outlandish. And, for another thing, it threatens our perceptual beliefs directly, not by way of closure.<sup>2</sup>

The threat involves *danger* or *risk*, the danger or risk of believing falsely. This does not depend on awareness by the subject in danger that or how he may be at risk, nor on his evidence concerning such risk, nor even on whatever evidence may be *available* to him. So I would reject that line of reasoning in García's paper. As one strides across a minefield one can be in great danger, even with no inkling of that fact, nor any available evidence for so believing.

The danger to one's attaining ordinary knowledge does require a certain orthodox conception of dreams, according to which beliefs and experiences in our dreams are hosted not only in the dream but also in actuality, while we dream. Only thus would our ordinary perceptual beliefs be threatened by the possibility that in a realistic dream we might believe the same on the same experiential basis.

Is that really how we should conceive of our dreams? Are dreams made up of conscious states just like those of waking life except for how they fit their surroundings? The orthodox answer is in the affirmative. Dream states and waking states are thought intrinsically alike, though different in their causes and effects. The orthodox view is deeply flawed, however, or so I argue in the book, before suggesting a better view. To dream is to imagine, not to hallucinate. My first chapter defends this in detail.

Ordinary perceptual beliefs can hence retain their status as apt, animal knowledge, despite the nearby possibility that one is asleep and dreaming. Ordinary perceptual beliefs can still attain success through perceptual competence, despite the fragility of that competence and of its required conditions. However fragile a performer's competence may be, and however fragile may be the conditions appropriate for

<sup>2</sup> Alternatively, one might argue that just as the modal proximity of enough possible fake barn encounters creates a problem for the belief that one perceives a barn, so the modal proximity of enough possible dreams creates a problem for the belief that one perceives an external reality.

its exercise, if a performance does succeed through the exercise of that competence in its proper conditions, it is then an apt performance, one creditable to the performer. Knowledge is just a special case of such creditable, apt performance. Perceptual knowledge is unaffected by any fragility either in the knower's competence or in the conditions appropriate for its exercise. The knower's belief can thus remain apt even if unsafe through the proximity of the dream possibility.

## 2. *Replies to Miguel Ángel Fernández and Jonathan L. Kvanvig*

My book thus develops an account of ordinary perceptual knowledge, one that comes under further scrutiny by Miguel Ángel Fernández and by Jonathan Kvanvig. I turn next to their discussions.

In the book a kaleidoscope perceiver forms beliefs about a surface that he sees. A jokester controls both its color and the quality of the ambient light, and is liable to arrange either of the following combinations with equal ease: first, white-light + red-surface; second, red-light + white-surface. These outcomes are moreover, by hypothesis, indistinguishable to our subject. Given that the ambient light could as easily be bad as good, then, even if the subject sees the surface in its true colors, and knows accordingly, when the light is good, he does *not* know that he knows any such thing. Falling short of such meta-knowledge, he thereby falls short of reflective knowledge. If we wish to uphold our ordinary perceptual knowledge even when we might too easily be dreaming, therefore, we need to find some difference between the kaleidoscope perceiver in danger of being misled by bad light, and the ordinary perceiver in danger of being misled by a dream.

Fernández objects as follows.

[According] to Sosa, both the kaleidoscope perceiver and the ordinary perceiver have [animal knowledge, apt belief] that  $p$ , because the object-level beliefs of both fulfill the two parts of the [following condition]:

For any correct belief that  $p$ , the correctness of that belief is attributable to a competence only if

- $C_i$ : it derives from the exercise of that competence in appropriate conditions for its exercise, and
- $C_{ii}$ : that exercise in those conditions would not then too easily have issued a false belief.

[...]

Sosa argues that the ordinary perceiver does, whereas the kaleidoscope perceiver doesn't, have an apt belief that he has an apt belief that  $p$ , i.e. R-knowledge [reflective knowledge] that  $p$ . Following the same pattern of explanation as before, this would have to be because the second-level belief of the ordinary perceiver does fulfill conditions  $C_i$  and  $C_{ii}$ , whereas the second-level belief of the kaleidoscope perceiver doesn't fulfill at least one of these conditions.

[...]

It is unclear why we should think that the ordinary perceiver's meta-belief does fulfill condition  $C_{ii}$ , and so the question whether the epistemic situations of the ordinary perceiver and the kaleidoscope perceiver are disanalogous up to the point that only the former has R-knowledge [reflective knowledge], remains open. (pp. 37–39)

This highlights the contrast of interest. *On one side* is the kaleidoscope perceiver whose circumstances might too easily go bad, which denies him knowledge that  $p$  in its reflective mode while allowing it to him in its animal mode. That is to say, the subject still believes aptly that  $p$ , but no longer aptly discerns that he does so. *On the other side* is the ordinary perceiver who knows that  $p$  in both animal and reflective modes, despite how easily he might be dreaming. Fernández suspects that this contrast has not been properly defended. He suspects in particular that the ordinary perceiver too fails the test whose failure denies reflective knowledge to the kaleidoscope perceiver. He worries that the ordinary perceiver too is deprived of such knowledge by his equal violation of a certain necessary condition.

The condition of interest is  $C_{ii}$ , that the competence to tell whether one has apt belief *not* be too easily misleading. It is not clear that the ordinary perceiver can satisfy this condition for his competence to tell whether he is or is not dreaming. It is not clear that the ordinary perceiver is any more able to ward against being so misled than is the kaleidoscope perceiver able to ward correspondingly (once the hidden jokester controls the light).

That worry worries me persistently in the book. Fernández and Kvanvig are right to press at just this point. Here I will try to develop my approach more fully, correcting it along the way, so that the core of my proposed solution will stand out more clearly and attractively. I begin by responding directly to the questions and objections raised.

Kvanvig has two main problems for my account. He argues, first, that we must attribute reflective knowledge to the kaleidoscope perceiver, not only animal knowledge. He considers the account of reflective knowledge as  $K+$ , i.e., as  $KK$ , apt belief aptly noted, or animal knowledge on top of animal knowledge. And he argues that this unified theory of animal and reflective knowledge cannot stop the kaleidoscope perceiver from enjoying reflective knowledge. Only an account of reflective knowledge in terms of reflective *coherence* will do the job.

Kvanvig's second problem for my account derives from a particular sort of Gettier case, one in which misleading but highly credible testimony, beyond the subject's ken, spoils his knowledge that  $p$ . The subject's belief that Tom Grabit stole a library book (based on superb evidence) is undone by Grabit's mother when she says that the thief was really Tom's twin, not Tom. According to Kvanvig, I am unable to explain why it is that here the subject falls short of knowledge, even when we suppose that the mother is lying.

I respond to these two objections in order.

*Contra* Kvanvig's first objection, the kaleidoscope perceiver fails to satisfy my requirements for iterated, reflective knowledge. Why so? Because he fails to satisfy a condition that my account lays down as necessary for knowledge. Kvanvig himself quotes that condition as follows.

Sosa formulates the following necessary condition for basing on indications of truth:

(I') S has animal knowledge that  $p$  based on indication  $I(p)$  only if either (a)  $I(p)$  indicates the truth outright and S accepts that indication as such outright, or (b) for some condition C,  $I(p)$  indicates the truth dependently on C and either (i) S accepts that indication as such not outright but guided by C (so that S accepts the indication as such on the basis of C), or else (ii) C is constitutive of the appropriate normalcy of the conditions for the competence exercised by S in accepting  $I(p)$ . (p. 55)

Why does the kaleidoscope perceiver fail to satisfy this condition in believing *that he knows the surface to be red*? Because he has no indication of this fact (the one in italics), none that indicates it either outright or relative to some condition C that satisfies either of b(i) or b(ii). As the example is set up (at least implicitly), the kaleidoscope perceiver cannot tell whether the jokester is spoiling

the light, nor even that there is a jokester in control. Accordingly, the perceiver would take himself to know the color of the surface regardless of whether the jokester spoiled the light. The perceiver has no indication that he knows, no tell-tale indicator that would be present outright only if he did know, nor even any that he can trust based on some further condition. No condition that he can discern is one relative to which he can trust some indication to the effect that he knows. The kaleidoscope perceiver hence violates I' and lacks animal knowledge that he has animal knowledge that the surface is red.

As for the Grabit case, the subject's epistemic belief-forming meta-disposition might too easily deliver an incorrect answer to the question whether he knows that Grabit is the thief. It might too easily take into account the easily available, highly credible (though lying) testimony of the mother, and would then deliver that Grabit is innocent.<sup>3</sup>

Stepping back, here's how the case looks to me. The subject sees Grabit run from the library holding the book, with the guard in hot pursuit. On that basis he forms the belief that Grabit is stealing a book, and let's suppose that to be a good basis (adding also loud shouts of "Thief!" from the guard, from the librarian at the door, etc.). Based on such reliable and truth-indicating evidence, let's say he does know at the animal level that Grabit is stealing a book. This is analogous to the barn-perceiver's knowing at the animal level that he sees a barn. However, both the Grabit subject and the barns subject fall short of reflective knowledge. In neither case does the subject know that he knows. In each case, too easily might the subject have received misleading deliverances that he would not have spotted as misleading.

We still need to specify more clearly the relevant contrast, however, between first-order and second-order beliefs, such as those below.

*First-order:* the belief of the kaleidoscope-perceiver as to whether the surface he sees is red; the belief of the barn-perceiver as to whether he faces a barn; the belief of the Grabit subject as to whether Grabit stole the book.

*Second-order:* the belief of the kaleidoscope perceiver as to whether he knows that the surface he sees is red; the belief of the barn-perceiver as to whether he knows he faces a barn;

<sup>3</sup> In line with my suggestion, our intuition that the subject fails to know grows stronger in direct proportion to the modal proximity of the possibility that we avail ourselves of the lying testimony.



the belief of the Grabit subject as to whether he knows that Grabit stole the book.

The correctness of each first-order belief manifests a competence exercised in its normal conditions. By contrast, each second-order belief derives from a disposition to judge when one knows, where too easily might *that same disposition* have been exercised in relevantly similar conditions while yet the belief so formed was false.

Consider the kaleidoscope subject, the barns subject, and the Grabit subject. Each forms his first-order belief through a disposition exercised in conditions wherein it constitutes a competence dependently on whether the distinguished manifestations are and would be triggered. The relevant first-order epistemic competences here are respectively: color-discerning, barn-discerning, and thief-discerning. Each of our three subjects has the relevantly discerning competence, in virtue of having in that circumstance a disposition to issue the distinguished manifestations under the relevant triggers.

Things change when we move to the second order. Here again the relevant manifestations would fail to be triggered reliably by the relevant triggers. Too easily might the subject have believed he knew when he did not know. The color appearance might now too easily mislead, as might the barn-façade appearance, as might the thievery appearance (the *total* thievery-relevant appearance, including the misleading testimony). What makes the difference for these second-order cases? Why are they importantly different from the first-order cases? Answer: there is nothing here like the good light, or the backsides normally attached to the façades, or testimony that is truthful (as testimony normally is). These are conditions respectively required, on the first order, for determining whether the subject has the relevant *complete* color-sight competence, or sortal-identification competence, or ability to attain truth through credible testimony. And they are conditions that are *present* for the subject's target belief (that he sees a red surface, for example, or a barn, or a thief) but *absent* in the ostensibly problematic modally proximate alternatives. When such conditions appropriate for the determination of an inner competence are present, finally, and the inner competence remains, it follows that the relevant *complete* competence is also present.

What's different on the second order? Since nothing there indicates abnormality, the subject eases into default trust in his color-detection, sortal-detection, and reliable testimonial mechanisms. This disposition of default trust in our faculties and their deliverances is exercised uniformly, absent some special indication to the contrary,

as we move from the good light to the bad light, or from the real barn to a fake barn, or from the true testimony to the lying testimony, or more generally from a situation wherein our epistemic competence can be trusted to one where it cannot be trusted.

Unpersuaded, our critic presses his doubt:

What, exactly, is the difference from what happens on the first order? The subject cannot tell on the second order when the conditions for exercising his first order color sight are normal and when they are not. He is not privy to the control of the jokester over the quality of the light. So, why does that not remove his first-order competence, in such a way that even when the jokester allows the light to remain good, the kaleidoscope perceiver still does not discern the redness of the surface he sees? *If* the first-order competence *is* thus removed, then the kaleidoscope perceiver does *not* after all attain apt belief, and animal knowledge, that the surface is red.

This goes wrong in supposing that if the kaleidoscope perceiver cannot tell when the light is good, then he lacks the relevant faculty, i.e., the inner color sight competence. This supposition is false because it ignores a crucial factor: the “appropriate conditions” that need to be in place if the presence or absence of that visual faculty is to be determined by whether or not the subject would issue the successful manifestations under the relevant triggers. Again, the fact that you cannot discern colors and shapes under bad light takes away your *complete*, situational competence, but it does not impugn your faculty, your *inner* competence, your good color eyesight. You retain that inner competence even when the light is bad, and you make lots of mistakes. To manifest an inner competence you must issue a relevant manifestation under a relevant trigger *in appropriate conditions*. (It follows of course that the exercise of such a faculty in its appropriate conditions *cannot* lead us astray, at least not too often.) The kaleidoscope perceiver does issue a relevant manifestation (when he takes the surface seen to be red) under a relevant trigger (looking at the surface and posing the question as to its color) in appropriate conditions (conditions of good light, those *required* for the quality of his relevant conduct to determine the presence or absence of the relevant faculty, or inner competence). The kaleidoscope perceiver therefore retains his inner visual faculty and he exercises that inner competence in its appropriate outer conditions —the good light, etc.— enjoying thus a relevant complete competence (inner

and outer). And he manifests this complete competence in correctly taking the surface seen to be red.

Nothing analogous appears on the second order. Take the second-order disposition to tell whether your first-order perceptual belief is apt. This, I suggest, is a default disposition by which one just assumes that one's relevant first-order faculty and conditions are appropriate unless there is some sign to the contrary. That is how one tells that the light is good, or that one is awake, not just dreaming. These are just things one assumes by default, absent tell-tale signs to the contrary. What further condition is now appropriate for determining the possession or lack of the relevant second-order faculty?

One trusts one's cognitive dispositions by default in the absence of contrary signs. Consider now the following condition: *that there would be tell-tale signs if one's relevant first-order disposition were then unreliable*. This condition is normally satisfied when we exercise our perceptual dispositions. But it fails to be satisfied when, hidden from view, the jokester controls the light. In this second case we lack the complete competence (inner plus outer) for telling whether we know by sight the color of the seen surface.

As soon as this disposition of default trust would too easily mislead in our circumstances, because there would be no warning signs, we lose the relevant complete competence, which deprives us of the ability to tell aptly that our relevant first-order belief is apt. Having *lost* our relevant complete second-order *competence* there's no way for us now to manifest it by correctly believing that our relevant first-order belief is apt. What removes our complete competence, our ability to tell whether our first-order competence is trustworthy, is just this: that now there would be no tell-tale signs that would tip us off, in too many cases where things go wrong in modal proximity to the actual case.<sup>4</sup>

Consider again the ordinary perceiver (especially an alert believer running outside at noon, not a drowsy subject in bed). Some might *still* doubt that he enjoys reflective knowledge. To anyone racked with such doubt, we can offer a fallback position. We can deny that while dreaming we really believe the things that we believe *in our dream*. Among the things we may believe *in* a dream without really believing it, *while* we dream, is the proposition that we are then awake and perceiving things. Once it is granted that we form no such beliefs while dreaming, we can deny in this further way

<sup>4</sup>The position sketched here is developed more fully in my "Perceptual vs. Reflective Orders of Epistemic Competence" (forthcoming).

that the dream possibility poses any real epistemic risk to ordinary perceptual beliefs. (This is the fallback defended in the first chapter of my book.)

### 3. *Reply to Ram Neta*

One main epistemological project aims to explain what in general might be the sources of our knowledge. I don't mean the merely historical or causal sources. A scholarship stipend might help bring about one's knowledge of physics, by enabling one to earn a degree in that field. Turning one's head at just the right moment might be a source of one's knowledge that a bird is flying by. These are not the sources of special interest in epistemology. Among the things with distinctive epistemological interest are rather the following: What constitutes knowledge? What constitutively or necessarily gives to a belief a normative status that it must have if it is to constitute knowledge? If knowledge is constituted by a certain sort of true belief, we then want to know the conditions that a true belief must satisfy, in addition to being a belief and being true, in order to constitute knowledge. By this I mean the conditions *in virtue of satisfying which* a true belief amounts to knowledge.

In my view a belief needs to be *competent* at least, if it aspires to be knowledge. It needs also to be *true*, of course. But each of these is a status that a belief can have without having the other. What is more, a belief can have *both* of these while still falling short. In order to qualify as knowledge, a belief must be correct in a way that manifests the believer's epistemic competence.

*Epistemic* rational justification is best viewed as either tantamount to or importantly involved in epistemic competence. Externalists stretch the term to cover blindsight and chicken-sexer believers, even independently of whatever track-record inductive justification these may eventually amass. Internalists will want to put aside whatever positive status such believers may gain through their subpersonal competences, distinguishing such status from true rational justification. In their view, such true justification requires the sort of access we have to our own present conscious mental states, and to the conscious reasoning that we do on their basis.

Blindsighters and chicken-sexers plausibly know sans reflective access to the truth or aptness of their beliefs. We might accordingly do well to surrender the term "rational justification" to internalists. A belief might be *epistemically competent* without being *rational*

*justified*, after all, and rationally justified without being fully epistemically competent. It is competence that matters for knowledge. Good reasoning, which results in rational justification, is only one way to attain competent belief.

However that may turn out, let us here focus on justification in any case. And let us restrict our explanatory project to the question of what in general might render a belief epistemically justified. Sometimes a belief is justified because it is based on good reasons in the believer's possession. But what is it for a good reason to be thus in someone's possession? Let us here think of reasons as ostensible facts, for the sake of argument (though alternatively we could think of reasons as mental states of the believer). On that conception, a reason is in the relevant sense *had* by someone only if they believe that ostensible fact. With this clearly in mind, we need to reconsider the explanation of a belief's justification by appeal to "reasons in the believer's possession". Plausibly, those reasons are *owned* by the believer only provided he believes them in turn.<sup>5</sup> The explanation of the belief's justification will then be satisfactory only if these further beliefs are *themselves* also justified. (By a *belief's justification* here I mean the epistemic justification of that belief state—the believer's so believing *justifiedly*—which is presumably required for that belief to constitute knowledge.) And this ushers in the regress/circle/foundations problematic. Not *all* the beliefs in a body of beliefs could possibly be justified just in virtue of being based on *other* beliefs in that same body of beliefs. Epistemic justification demands some other source of justification that will somehow relate that body of beliefs to the surrounding world. This is the foundationalist position.

Foundationalism tries to stop the regress of epistemic justification by appeal to the given. Experiences hosted passively lie beyond justification and unjustification. Yet they do provide justification for related beliefs. A migraine will make one justified in believing that one's head aches even though it makes no sense to wonder about its own justification.

Such appeal to the given may or may not provide the basis for an account of empirical epistemic justification. It remains to be seen

<sup>5</sup> Here one should distinguish the sense in which Watson may perhaps "have" reasons to suspect what Holmes suspects, even if only Holmes has made those reasons his own in a way that helps explain why he, unlike Watson, is epistemically justified in suspecting someone in particular. Watson *may*, I say, "have" those reasons if they are right before his nose. More properly, however, one should say that those reasons are *available* to him, however easily, and not yet strictly in his possession.

whether that initially promising suggestion can be developed into a fuller account. However that may turn out for *empirical* justification, it is doubtful that it could be extended to cover a priori as well as a posteriori justification. It is hard to see what could possibly play the role of sensory experience for our beliefs of simple math and logic.

It may be thought that the intuitions constitutive of a paradox, such as the sorites, *are* a priori correlates of sensory experience. Even when our solution to the paradox involves rejecting one of those intuitions as false, the force of the intuition may yet remain. Paradox-enmeshed intuitions seem thus similar to perceptual appearances involved in a perceptual illusion, such as the Müller-Lyer. The lines still appear incongruent even after measurement reveals them to be congruent. Similarly, the rejected proposition in the paradoxical cluster may still seem true even after one believes it to be false, in line with one's solution.

Clearly there are such pre-belief a priori seemings. What is not so clear is whether they are relevantly like sensory experiences. After all, one's visual experience as of a chessboard may or may not be accompanied by a corresponding seeming. In an expert that experience will prompt the seeming as to the number of squares immediately upon his asking the question. The visual field of a tyro will have the same number of squares, presumably, although that will prompt no immediate seeming for him. He will need to count in order to so much as hazard a good guess. The Müller-Lyer subject will also have such a seeming, one derived from his visual experience of the lines. But the contrary seeming derived from measurement-cum-memory will normally prevail.

Seemings are hence to be distinguished from experiences. Undeniably, there are a priori seemings, such as those operative in a paradox. And these seemings can be overridden as part of the solution for the paradox, as happens also in the Müller-Lyer case. What is missing on the a priori side is any *further* mental state that can serve as a counterpart of the empirical sensory experience. The chess expert and the novice share a visual experience involving 64 squares. Only to the expert does it seem intellectually, however, that the squares number 64. This is beyond the novice.

Compare this: when it seems to us that a set could not possibly be equinumerous with a proper subset, this seeming has no apparent correlated experience distinct from it. Nothing beyond consciously considering the set-theoretic proposition prompts our assent to it. This is quite unlike what happens in the case of the chessboard.

There the expert also considers consciously a proposition, as to the number of squares before him. But that is not by itself sufficient to prompt the attraction to assent. The experience also plays a role, as does the expert's background knowledge about chessboards. His visual experience yields the chess expert's seeming, a role unmatched by any correlate for such experience when it seems to us that a set cannot be equinumerous with a proper subset.

Intuitive seemings are hence distinctive in deriving from the sheer understanding of the propositional content involved. One understands the question and straightaway one is attracted to an answer, with no benefit of ulterior reason. The attraction derives from nothing beyond sheer understanding of the content. The chess expert can entertain the proposition whether he faces a 64-square array, can do so with full understanding, and yet remain unmoved, if his eyes are closed.

Ram Neta presents a view that he takes to represent a main tradition in epistemology, and supposes my own preferred virtue epistemology to oppose that traditional view. Here's how he presents that view.

According to a historically influential Cartesian picture of knowledge, all of the knowledge that normal human adults possess is knowledge that is, in one or another way, *built out of* the epistemic deliverances of critical reasoning —i.e., it is built out of our knowledge of just those facts that determine what we rationally ought to believe or intend. According to an especially prominent foundationalist version of this picture, our empirical beliefs about the world around us are justified by our knowledge of our own sensory experiences and apparent memories: critical reasoning tells us what experiences and apparent memories we're having, and it also tells us what justificatory relations obtain between our knowledge of those experiences and memories, on the one hand, and our beliefs about the world around us, on the other. (p. 64)

Neta then attributes to me a critique of that view of knowledge, on both its a priori and its a posteriori sides. He acknowledges my claim that there is nothing relevantly like sensory experience to stop the regress through either rational intuition or perception or introspection. Nothing on its own provides justification for our intellectual seemings or beliefs, either a priori or a posteriori, while itself remaining beyond justification and unjustification. And he adverts to my account of foundational justification in terms of seemings or beliefs that manifest the subject's epistemic competence. But that is then followed by this puzzling objection:

So what reason is there to think that, when we are properly exercising our powers of reflection in reflecting upon some fact, the very fact that we are doing so is not itself accessible to our own powers of reflection? In order for Sosa to make his case against the internalist foundationalist picture of a priori knowledge, and to show that some appeal to virtue is necessary for an adequate account of such knowledge, he must provide an answer to this question. So far as I can see, he does not do so anywhere in his 2007 book. (p. 68)

Consider our belief that  $0 < 1$ . Suppose it is granted (if only for the sake of argument) that such a belief is foundationally justified just because we form it by properly exercising our powers of reflection. Is that not already to grant that some appeal to virtue is necessary for an adequate account of such knowledge (or at least contributes to such an account)? According to Neta, the fact *that we there properly exercise our powers* may well itself be accessible to our powers of reflection. That very fact (in italics) —call it F— then suffices to justify us immediately in believing that we so do.

I have two comments about that suggestion. First of all, I am unable to see why the fact that F justifies belief in F should be any impediment to the role that F is said to play in accounting for our justification for believing that  $0 < 1$ . Given that role, however, our competence (epistemic virtue or power) does help account for the epistemic justification of our belief. This is so even if we do have immediate access to the fact that such a foundational a priori belief is formed, thus, through a proper exercise of our intellectual virtue.

In the second place, just the fact that I am exercising my competence properly seems insufficient to justify me, all by itself, in believing accordingly. Indeed, the view that it is sufficient is subject to my objection against the traditional myth of the given: if someone believes that he is exercising his competence properly, but is very poor at judging whether he is doing so, then his belief that he is exercising his competence properly may not itself be an epistemically justified belief. Here again I would argue that the sheer truth of that fact is insufficient on its own to account for the believer's justification in holding the relevant belief. The believer's ability to discern such facts reliably is also required.

Consider, in any case, the master thesis of virtue epistemology: *that knowledge is apt belief*, which imports a central appeal to competence, since what makes any apt belief apt is that its correctness manifests the believer's epistemic competence. These core components of virtue epistemology *commit* one to no particular view on



whether the presence of our competence must itself be reflectively accessible. That it must *always* be accessible to *pure armchair reflection* seems quite implausible, but this is a separable, externalist, component of virtue epistemology. Descartes, for example, was a virtue epistemologist committed to a view of certainty as superlatively apt belief.<sup>6</sup> But he was also an *internalist* virtue epistemologist: the infallible reliability of our relevant epistemic competences was for him itself accessible to a priori (theological) reflection.

*Received: April 26, 2010; accepted: May 5, 2010.*

<sup>6</sup> As I argue in the first of my Carus Lectures (delivered at the American Philosophical Association meetings of February, 2010), forthcoming.