

Doncaster pandas and Caesar's armadillo

Scepticism and *via negativa* knowledge

Levi Spectre & John Hawthorne

Forthcoming in *Philosophy and Phenomenological Research*

The external world sceptic tells some familiar narratives involving massive deception. Perhaps we are brains in vats. Perhaps we are the victim of a deceitful demon. You know the drill. The sceptic proceeds by observing first that victims of such deceptions know nothing about their external environment and that second, since we cannot rule out being a victim of such deceptions ourselves, our own external world beliefs fail to attain the status of knowledge. Discussions of global external world scepticism tend to focus on the second step, where a number of well-known lines of resistance have been offered.¹ But there has been little attention to the first, seemingly innocuous step. That will be the focus of this paper. Part one – sections 1, 2, and 3 – will explain why these standard narratives are not convincing examples of cases where there is no knowledge of the external world. In part two – section 4 – we shall undertake a useful case study. David Lewis's 'Elusive Knowledge' is often thought of as presenting an epistemological vision that is somewhat friendly to external world scepticism: as Lewis himself presents things, there are contexts where external world knowledge ascriptions are uniformly false, and where true knowledge ascriptions are limited to either axiomatic truths or truths about our inner life. We examine his discussion in the light of the preceding reflections and show that the framework he presents is not so concessionary to global external world scepticism after all.

1 | PRELIMINARIES

Let us call a proposition *strongly external* if its truth is never settled by the truths about someone's inner life. Take any proposition p that characterizes a possible being's inner life. A strongly external proposition is such that it is not entailed by any such p . The global external world sceptic maintains that *no strongly external proposition is known*. (There are propositions that are not strongly external but are in some looser sense 'about' the external world. Consider, for example: 'I am either having a headache or eating a Boston creme doughnut' and

¹ These include but are not limited to: (i) A denial that our evidence in the good case is limited to phenomenal evidence; (ii) Arguments to the effect that our mundane beliefs form part of a better explanation of the experiential data than various sceptical alternatives; (iii) An insistence that knowledge only requires freedom from error at *close worlds*; (iv) An appeal to normality considerations, whereby the relative abnormality of sceptical scenarios is intimately linked to our knowledge that they don't obtain; (v) An appeal to the reliability of the relevant empirical belief-forming mechanisms; (vi) A Bayesian framework that advocates for ur-priors that assign low priors to various sceptical narratives; and so on.

'It's not the case that I have a headache from listening to an Iron Maiden concert'.² It ought to be no part of external world scepticism that propositions expressed by sentences like these cannot be known.)

The kind of global external world sceptic that interests us here is one that tries to bring us over to their gloomy outlook by telling one of the rich and familiar narratives involving rampant deception. This is to be sharply contrasted with the kind of sceptic that, at the outset, puts forth a tendentious account of knowledge according to which a proposition can be known only if it is entailed by one's phenomenal evidence. As we have set things up, it is trivial to show that external world propositions are not entailed by our phenomenal evidence – all that is required is to notice that phenomenal facts are truths about our inner life. When it comes to this last kind of sceptic, pretty much all the action concerns the tendentious theory of knowledge. Aside from that, the only issue when it comes to applying scepticism to particular propositions will be to show that the proposition in question is strongly external – and one typically doesn't need a rich narrative to convince oneself of that. The sceptic that interests us here – and which is most familiar – doesn't begin with a tendentious theory of knowledge. Rather they tell stories about wholesale deception and trust that it will be intuitively compelling that for all we know we are the victims of such deception. If a thesis connecting phenomenal evidence and knowledge is suggested at all, it will be later in the dialectical game, once the force of the narratives has been conceded and appreciated.

2 | THE EVIL GENIUS HYPOTHESIS

In the first of Descartes' (1641) meditations, he imagined an evil genius 'not less powerful than deceitful' that had 'employed all his energies to deceive me'. When confronted with this narrative, there is a strong temptation to concede that, for all we know, we are the victim of such deception. But how does this concession relate to global external world scepticism? However exactly the journey to scepticism gets fleshed out, the usual starting point, rarely remarked on at length, is that if we are, in fact, in the 'bad case' of evil genius deception, then we know nothing about the external world:

IGNORANCE: If we are the victim of an evil genius, we don't know any (strongly) external world propositions.³

² Strongly external propositions should be distinguished from what may be called *super strongly external* propositions, propositions such that both they *and their negations* are strongly external. To see the difference, consider 'I have a headache from listening to an Iron Maiden concert', a proposition that is strongly but not super strongly external: No inner life fact can conclusively establish it, but certain (headache-free) inner lives can establish its negation.

³ Despite its largely anti-sceptical leanings and IGNORANCE's particularly menacing sceptical implication, contemporary epistemologists often accept it and hardly ever explicitly reject it. For instance, 'You can't know anything about the external world on the basis of your perceptual experiences', if 'You can't know you're not being deceived by an evil demon' (Pryor 2000, 528). And since you can't know you're not if you are a deceiving demon's victim, that means IGNORANCE is true according to this view. The major step here – why knowing nothing about the external world follows from never knowing 'you're not being deceived by an evil demon' – is 'some sort of Closure Principle' (Pryor 2000, 522). However, no closure principle on its own has – as far as we can see (but perhaps counter-closure does, see below) – this implication. (Similarly, Markosian 2014, speaking of the proposition that

This is the thesis that we will subject to careful scrutiny. How is IGNORANCE to be justified? The most obvious way to try to justify it is by appealing to the factivity of knowledge. Here is the line of argument:

IGNORANCE VIA FACTIVITY: For any p , if we know p , then p . But every strongly external proposition p that we believe is false in the evil demon scenario. So if the evil demon scenario obtains, we don't know any external world or strongly external propositions.⁴

But that line of argument is no good. It is not hard to find strongly external propositions that we believe and that are true in the evil demon scenario. A plentiful supply of them can be generated by reminding ourselves of negated propositions that we believe.⁵ Each of us (i.e., the authors) believes that we are not having lunch with Barack Obama. Each of us believes that

we are not brains in vats, states: 'And if you can't know that, then how can you know anything about the external world'. And see the quotes below in footnote 4 from *The Routledge Companion to Epistemology* or from *The Internet Encyclopedia of Philosophy* that relay the same idea with different intermediary steps.) Note that these authors are explicitly speaking in favour of a view according to which a failure to know we are not brains in vats implies a failure to know anything else about the external world. But if, contra IGNORANCE, actually being brains in vats is compatible with such knowledge, then *a fortiori*, a failure to know we are not brains in vats, is so compatible. We should also mention that even the best-known closure deniers – i.e., Dretske 1970 and Nozick 1981 – raise no special concerns about IGNORANCE. (Famously, they argue that knowing, say, one has hands is compatible with not knowing one is not the victim of a demon or vat scientist. But this idea is neutral about IGNORANCE itself.) Magidor (2018) comes closer to rejecting IGNORANCE: Relying on externalist considerations, she argues that there are possible brains in vats who know they are brains in vats. But the knowledgeable brains in vats she entertains are not like us – they believe they are brains in vats. Her remarks are silent on whether, if we are in fact envatted, we in fact still have some external world knowledge.

⁴ That IGNORANCE VIA FACTIVITY is the operative underlying argument is very much suggested by the opening words of Hickey's (n.d.) *Internet Encyclopedia of Philosophy* 'The Brain in a Vat Argument' entry: 'If you cannot now be sure that you are not a brain in a vat, then you cannot rule out the possibility that all of your beliefs about the external world are false'. Others have suggested, maybe less directly, that a supreme deceiver – that controls our sensory input, i.e., a Cartesian demon or the vatmeister – might deceive us about *everything*. For instance, *The Routledge Companion to Epistemology* describes the sceptical argument relying on an evil genius that is 'bent on universal deception' allowing the sceptic to 'infer you do not know anything about the external world' (Schantz 2011, 479-480). Also, though factivity's role is already clear from the standard discussions of vatted brain scenarios, it's even clearer when those scenarios are employed to target other knowledge. Take, for instance, Nozick's (1981) plot twist where he's a brain in a vat on Alpha Centauri and not at Van Leer Institute's Library (in Jerusalem on Earth, the place where he believed he was at the time of writing the plot twist). It would hardly make sense to tell this story about a vatted brain on a star if a vat at Van Leer's Library could do the sceptical work.

⁵ The role negative beliefs play here is similar to Maimonides's *via negatvia* theology – Halbertal (2014) persuasively defends *via negativa*, but also presents three other possible readings of Maimonides's (1190) *The Guide of the Perplexed* – that sees progression in religious excellence as gaining knowledge of categorical negations. These negations capture the misapplications to God of properties, relations, quantifiers (including the existential quantifier), analogies, and metaphors, indeed anything that could create an illusion of understanding and belief in the standard sense. (According to Maimonides, referring to God by a proper name carries no implicit attributions.) Unlike mystical traditions – among Halbertal's (2014) four is also a compelling mystical reading – that see thought and language as barriers to spiritual knowledge and experience of the transcendent, *via negativa* offers the study of the limits of religious language together with the study of nature (e.g., physics, biology, physiology, etc.) as the method of gaining knowledge of the utterly transcendent deity. Maimonides arguably relies on a special kind of 'categorical negation' that is distinct from the negation that one uses to negate a property, and so he appears to sharply distinguish the categorical negation 'It is not the case that God speaks' from 'God is such that he doesn't speak', where the former but not the latter is reckoned felicitous. (This, in effect, repudiates the ordinary language analogue of the rule of Beta conversion in the lambda calculus.) Our much less lofty endeavour doesn't take on board Maimonides's special treatment of categorical negation but is sympathetic to his general view that negated propositions offer a path to knowledge that transcends experience.

we are not deep under the ocean holding an armadillo owned by a Roman emperor. And so on. These propositions are strongly external. Moreover, many will be true in the evil demon scenario. This is even more obvious in the brain in a vat and dream narratives. Recall that you believe that you are not holding Caesar's armadillo deep in the ocean. (If you don't believe it, please break character and seek help.) No need to convince you that this proposition is true in the 'good case'. But it's also true in the 'bad case' because a handless brain in vat can't hold anything, let alone an illustrious armadillo. Similarly, even if your whole life is a butterfly's dream, you at most could have a dreamed lunch with Barack Obama, not the kind of encounter a butterfly might have with the former president that would render your belief false.⁶

Another plentiful supply of beliefs that are intuitively true even in stock sceptical scenarios are counterfactual beliefs. A particularly easy way (though by no means the only way) of generating various such truths is to exploit negated consequents: 'If I were a brain in a vat, I wouldn't be a brain in a vat holding an armadillo deep under the ocean'; 'If I were to eat a fish in the next hour, I wouldn't eat more than fifty fish in the next hour'; 'If I were in close proximity to a panda, I wouldn't do something evil'.

Thus, *IGNORANCE* cannot be underwritten by an observation to the effect that all our strongly external beliefs are false in the evil demon scenario.

Is there some other narrative that we can substitute for the evil demon scenario so that the corresponding *IGNORANCE VIA FACTIVITY* argument can go through? It is easy to see that there is no hope of this. Consider for example, the pair: I don't have exactly two pounds of salmon in my pockets; I don't have exactly three pounds of salmon in my pockets. No narrative can falsify both strongly external beliefs.⁷

We can state abstractly a sufficient condition for blocking any counterpart of the *IGNORANCE VIA FACTIVITY* argument. Let us say that a set of propositions is negation-compatible just in case: (i) an agent believes each of them and (ii) the negations of those propositions are compatible. Now there are many pairs of propositions that we believe that are *not* negation-compatible: if the negation of one of those propositions is true, the negation of the other proposition is false. If, for a given agent, there is a set of propositions that is not negation-compatible, then there will be no sceptical scenario that can vindicate the counterpart of *IGNORANCE VIA FACTIVITY*. Since the negations of the propositions in that set are incompatible,

⁶ This argument shows that the brain in a vat and Descartes' (or Zhuangzi's) dream arguments do not support global external world scepticism (at least not by anything like the factivity route). Unlike the barebones Cartesian evil demon argument, they include details relating to our bad case state (e.g., asleep in bed dreaming, having no body, and so on). And so certain negative beliefs will, thanks to those details, be inevitably true in the bad cases – e.g., a bodyless brain is not holding an armadillo.

⁷ And if one is sympathetic to counterfactual excluded middle (as we are), the same goes for the pair 'If I were to have salmon in my pockets, I wouldn't have exactly three pounds of salmon in my pockets' and 'If I were to have salmon in my pockets, I wouldn't have exactly four pounds of salmon in my pockets'. Given counterfactual excluded middle, the only way for the first to be false is for the counterfactual 'If I were to have salmon in my pockets, I would have exactly three pounds of salmon in my pockets' to be true, but then the second counterfactual would be true. (In the case of counterfactuals with true antecedents, one merely needs the less controversial 'and-to-if inference' (secured by the strong centering assumption) to make the point. Suppose there are no pandas nearby *S* and *S* believes both 'If there were no pandas nearby, there would not be exactly thirty pounds of salmon within a few feet' and 'If there were no pandas nearby, there would not be exactly twenty pounds of salmon within a few feet'. Then, given the validity of and-to-if, one of the pair will be true.)

any scenario that falsifies some the beliefs will verify at least one of the others. Now one can find myriad sets of propositions that we believe that are not negation-compatible (indeed, we can find many two-membered sets of this sort). So there is no hope of a sceptical narrative that underwrites an *IGNORANCE VIA FACTIVITY* argument.

Of course, even if there is no path to global external scepticism from an evil genius hypothesis, it is still pretty bad news, epistemologically speaking, if we don't know whether we are the victims of an evil genius. Assuming epistemic closure for knowledge, if we don't know that, then a great many of the beliefs that matter to us fail to be knowledge. Still, global external scepticism looms large enough in epistemology that it is not of little importance to understand the standing of that thesis, even once the epistemic possibility of paradigm sceptical narratives is conceded. Our project here is to achieve some such understanding.

Before moving on, it is worth noting in passing an interesting result for at least one natural precisification of the evil genius hypothesis. Descartes imagined a supremely powerful evil demon that aimed to maximize deception. On a natural precisification, the truth of the hypothesis requires that Descartes be uniformly deceived. And when restricted to a genius hell-bent on external world deception, it is natural to precisify the hypothesis so as to require that the victim is deceived about all external world beliefs:

CARTESIAN EXTERNAL WORLD GENIUS: There is an evil genius who sees to it that, at all times, all of Descartes' external world beliefs are false.

Suppose there is a Cartesian external world genius. This can only happen if Descartes, at all times, had an external world belief set that is *negation-compatible*: If Descartes' beliefs at any time were not negation-compatible then it follows that there is no Cartesian external world genius that verifies the hypothesis. And, of course, Descartes' beliefs were not always negation-compatible. So the *CARTESIAN EXTERNAL WORLD GENIUS* hypothesis is false. Note that this also shows that the negation of this hypothesis is not a strongly external world proposition since possible inner life profiles conclusively establish its truth (at least assuming that some possible inner life profile entails the presence of beliefs that are not negation compatible).⁸

3 | OTHER AVENUES TO IGNORANCE

The stock sceptical narratives do not constitute scenarios where all our external world beliefs are false. So the *IGNORANCE VIA FACTIVITY* path to *IGNORANCE* is unavailable when it comes to such narratives. Is there some other route to global external world scepticism that proceeds via such scenarios that is at least *prima facie* promising? One thing to note is that it is not particularly obvious or intuitively compelling that the victims of deception in such scenarios know none of the propositions that we have been alluding to. It's not as if it's obvious that a

⁸ Similarly, it follows from our inner life that we have at least one strongly external true belief since, for example, our inner life entails that there are sets of strongly external propositions that we believe that are not negation-compatible in the sense given above. What that shows is that the proposition expressed by an utterance of 'I have at least one strongly external true belief' expresses a belief that is not itself strongly external.

brain in a vat doesn't know that it's not a brain in a vat being controlled by a team of thirty pandas or that a brain in a vat doesn't know that if it were to eat a pound of salmon, it wouldn't eat a panda as well. For various truths of this sort, it seems at best unclear whether the brain in a vat knows them. When knowledge cannot be ruled out on factivity grounds, its absence on these occasions is by no means intuitively obvious.

There is a slightly more theoretical tack that the global external world sceptic might try. Some have advocated for *counterclosure* as a structural constraint on knowledge: According to this principle, beliefs that are inferentially based on false beliefs are thereby not knowledge. With counterclosure in hand, ignorance spreads from those false beliefs had by, say, a brain in a vat to all the beliefs inferentially based on such falsehoods, including various true ones. Suppose, for example, that the brain in a vat believed that it is either sitting at a desk or will not eat pizza in the next week and that such a belief was inferentially based on the former disjunct. Then, even if the disjunction is true, counterclosure says it is not knowledge.

Is counterclosure a promising basis for claiming that victims in the stock sceptical scenarios know nothing? We should first remind ourselves that counterclosure is a controversial principle in epistemology, rather more controversial than a closure principle for knowledge. (For a range of worries about counterclosure, see Warfield 2005 and also Hawthorne and Rabinowitz 2017.) But in any case, it does not seem like an adequate basis for arguing that each belief in the stock scenarios is knowledge free. For while certain of the true beliefs in those scenarios are clearly inferentially based on falsehoods, that does not seem true of all of them. Consider, for example, beliefs such as 'If I were a brain in a vat, I would not be a brain in a vat deep under the ocean holding an armadillo'. Or 'If I were to eat a whole salmon today, I would not eat more than fifteen salmon today'. If asked to form a judgment about such propositions, one tends to respond affirmatively on the basis of supposing the antecedent in one's imagination and then developing that supposition in imagination. (See Williamson 2022 for details.) There is no clear sense in which the counterfactuals are inferentially arrived at from false propositions. Similarly, for various negated propositions. One might believe that one is not holding an armadillo belonging to a Roman emperor on the grounds that one is not holding an armadillo at all. But one might instead think to oneself that even if one were, despite appearances, holding an armadillo, it wouldn't be one belonging to a Roman emperor. And even more obviously, one might form a belief that there is no panda in front of one simply on the basis of there seeming to be no panda in front of one. Here, the proposition that it seems that there is no panda in front of one is a truth, and so there is no ground on which to indict the belief for counterclosure theoretic reasons. Counterclosure is too restricted an instrument to vindicate *IGNORANCE*, and so, again, the very first step in the standard sceptical path to global external world scepticism – a claim along the lines of *IGNORANCE* – is in question.

It may be helpful to frame the dialectical situation in terms of the standard accessibility framework for representing knowledge. Let's say that a world is epistemically accessible to an agent just in case it is compatible with what that agent knows that the world obtains. The external world sceptic can allow that there are worlds that are epistemically inaccessible to a given agent since for all this sceptic says, agents can know certain facts about their inner life. And if they do, worlds where their inner life does not match in the known respects are

epistemically inaccessible. But the external world sceptic cannot allow for there to be epistemically inaccessible worlds where the agent's inner life is a perfect match to their actual inner life. So what the external world sceptic needs to do is to motivate a model according to which every world where the agent has a matching inner life is epistemically accessible. Now, of course, one can write down such a model. The question is how to motivate it. We noted at the outset that such a model could be motivated by simply taking for granted that one can only know that which is entailed by one's phenomenal evidence. But we also noted that this is so highly tendentious a starting point as to be very underwhelming. And what we have seen is that the stock sceptical narratives do not go that far when it comes to motivating it. Indeed, it is hard to see that any single sceptical narrative is going to go that far by way of motivating such a model. No such narrative can motivate the counterpart of *IGNORANCE VIA FACTIVITY*. And once this is conceded, it is hard to see how a single narrative can do the work.

The external world sceptic may very well despair of appealing to a single grand sceptical narrative to do the argumentative work. A natural recourse is some kind of divide-and-conquer strategy, where one tells a multitude of narratives, each designed to cast sceptical doubt on certain of our beliefs so that for each external world belief, there is a sceptical narrative that problematises it. In the words of David Lewis (1996), we 'let our paranoid fantasies rip'. The sceptic hopes that a sufficiently rich exercise of our imagination will do the job even if no single narrative will suffice. The hope is that for each possible world where a given agent's inner life is a match, some narrative provides a compelling explanation for why that world is epistemically accessible to that agent.

Remember that the external world sceptic we are interested in needs to be doing more than simply reminding us that our phenomenal evidence does not entail this or that proposition. That is something that we were already perfectly well aware of. What sceptical narratives standardly do is have the effect of making the falsity of certain beliefs seem especially live by making salient detailed narratives that, once entertained, are *prima facie* narratives that for all we know we are in. But counterpart detailed narratives for certain other beliefs seem by comparison somewhat idle and pointless. Each of us believes that there is no panda in Doncaster that is glued to six lawnmowers. We know already perfectly well that the proposition that is here believed is not entailed by our phenomenal evidence nor, more generally, by facts about our inner life. Now one can tell an elaborate tale in which a panda is transported to Doncaster and glued to six lawnmowers, and that such fact is also deemed insufficiently newsworthy as to be worth relaying in news sources that we are privy to. The world of the story, rather than massively deceptive, will in many respects be very much like the world we think we are in, with certain anomalous goings on in Doncaster involving a panda. But the details of the story are not going to do very much. It's not as if, once the details are filled in, one reacts by saying, 'Oh, before I thought I knew that there was no panda in Doncaster glued to six lawnmowers, but now that you have told me an elaborate tale in which a panda is so glued, I'm a lot less sure about that!' Similarly, for an elaborate tale in which some people conspire to force-feed us fifty salmon in a day. We knew already that the propositions expressed by 'If I were to eat a salmon today, I wouldn't eat fifty salmon' were not entailed by our inner lives. The elaborate tale of force-feeding does not seem to add anything to that elementary observation. And thus,

it is far from clear how useful elaborate sceptical narratives really are if the aim of one's game is global external world scepticism. (Note that for some of our beliefs, the details of a 'bad case' will inevitably be filled with lots of good news. Consider Tamar's belief that it's not the case that things are as she thinks they are with the exception that a teapot recently came into existence by a quantum fluctuation on Pluto. Any detailed narrative that falsifies this belief will be one in which things are just as Tamar thinks they are everywhere else.⁹)

It perhaps bears emphasis here that certain narratives compatible with our experience seem far more promising as a basis for sceptical doubt than others. If someone claims to know that Sally is coming to the party, narratives according to which she is sick or run over by a bus seem much more telling as a basis for sceptical doubt than narratives according to which she has turned into a can of baked beans and is locked in the pantry or according to which she has been abducted by a roll of baby armadillos.¹⁰ If someone claims to know that they have hands, the standard brain in a vat narrative is far more compelling as a basis for sceptical doubt than a narrative according to which they are a handless frying pan with a rich experiential life. If the global sceptic throws out this kind of distinction – so that the handless frying pan is reckoned just as helpful as any other 'sceptical scenario' for their purposes, then they would seem to simply be falling back on treating as axiomatic the claim that whatever is external (in our sense) is not known. Meanwhile, if they, like many of us, see merit in the distinction between worthy and stupid sceptical hypotheses, then the path to global scepticism will be compromised: even when worthy sceptical hypotheses are treated as epistemically possible, that will not suffice to establish the epistemic possibility of the stupid ones and hence will not suffice to refute the claim that we know the negation of various stupid ones (e.g., 'I am not a handless frying pan'; 'Sally has not been abducted by a roll of baby armadillos').

The remarks of the previous sections bear not only on external world global scepticism. They also bear on those contextualists who think that at least *certain contexts* are such that characteristic speeches of the global external world sceptic, when made in those contexts, express truths. The work of David Lewis is a highly relevant case study here. His well-known 'Elusive Knowledge' purports to provide a framework in which certain contexts vindicate the speeches of the external world sceptic. With the preceding sections in mind, let us see whether he really does this.

⁹ Note also that Tamar's beliefs about how things are include such beliefs as that it's not the case both that there is a teapot on Pluto and that she knows that there is no teapot on Venus. So a 'bad case' for this belief will have to be one where she knows there is no teapot on Venus. For negated conjunctions like this, the strategy of inducing sceptical doubt by telling a narrative where the content of the belief is false seems to break down altogether. For if the narrative has it that, for example, she knows that there is no teapot on Venus, then it will, if anything, reinforce our conviction that if in fact there is no teapot on Pluto, that is something we know. (Thanks to Daniel Berntson for helpful discussion here.)

¹⁰ In case the reader was unaware, a roll is to armadillos as a flock is to sheep and a herd is to cows.

4 | ELUSIVE SCEPTICISM

David Lewis's (1996)¹¹ discussion of scepticism in his 'Elusive Knowledge' serves as an instructive case study and will helpfully illuminate various of the points that we already made. Let us briefly recap some of the essential moves of that paper. It starts by juxtaposing two competing frames of mind. The first is enunciated by the opening sentence of the paper, 'We know a lot'. It is a frame of mind where we have all sorts of 'everyday knowledge'. The other frame of mind is one we get into when we 'let our paranoid fantasies rip'. Such imaginative adventures bring into salience all sorts of 'uneliminated possibilities' (i.e., possibilities compossible with the facts of our actual experience).¹² The key motivator for this frame of mind is that we meet a compelling argument that 'we know next to nothing': 'If you grant that you cannot eliminate a certain possibility in which not-P, it certainly seems as if you have granted that you do not know P'. The upshot, when in this frame of mind, is that we know hardly anything, except, perhaps, 'a few simple axiomatic necessary truths; and [some knowledge] of our own present experience' (Lewis, 549).

In the face of this competition, Lewis advocates for a kind of contextualism that is designed to effect a compromise. Claims involving 'know' express different propositions in different contexts of use (even across contexts where the contribution of other vocabulary does not vary). What varies across contexts is which possibilities are properly ignored, where a 'possibility' in his sense is not just any old possibly true proposition but rather a proposition that is 'specific enough' for the conversational purposes at hand and which thus cannot be split into interestingly different subcases (Lewis, 552).¹³ He then lays down various rules that are supposed to jointly comprise a reasonable articulation of when possibilities are properly ignored. Most relevant to our discussion are: (i) The rule of *Actuality* – if a possibility actually obtains, then it cannot be properly ignored. (ii) The rule of *Resemblance* – if a possibility saliently resembles a possibility that cannot be properly ignored (in virtue of some of the other rules), then the former possibility cannot be properly ignored;¹⁴ and (iii) The rule of *Attention*: if a possibility is attended to, then it is not ignored and so, *a fortiori*, not properly ignored.

Lewis is thinking that in ordinary mundane conversational contexts, many knowledge ascriptions come out true. Suppose one of us says, 'It is raining outside' in a mundane context where it is raining. There are plenty of non-raining possibilities where our experiences match the experiences had during the speech. But such possibilities are properly ignored and so, in that context, 'know' expresses a relation that the speaker does have to the proposition that it is raining outside. But Lewis is also thinking that when we 'let our paranoid fantasies rip', we are thrown into a context where sceptical speeches are vindicated. The picture is one

¹¹ Page numbers below refer to Lewis's (1996).

¹² The definition of 'uneliminated': A 'possibility *W* is uneliminated iff the subject's perceptual experience and memory in *W* exactly match his perceptual experience and memory in actuality' (Lewis, 553).

¹³ Lewis does not want to restrict 'possibility' to individual possible worlds as they are so specific as to never be fit candidates for cognitive attention. Nor does he want them to be so course-grained as to count someone as attending to a sceptical alternative to *p* simply by considering whether or not *p*.

¹⁴ The parenthetical is important, since otherwise, the rule of resemblance threatens to be excessively sceptical in all contexts. (Assuming that one can get from any possibility to any other by the transitive closure of salient resemblance, the rule without the parenthetical qualification would, tentatively, generate scepticism in all contexts.)

where each of the competing frames of mind articulated at the outset of the paper has its day. The 'Moorean' frame of mind is perfectly correct in a mundane context. But in a sceptical context, Moorean speeches are incorrect. Rather, in those contexts, the gloomy – 'we know next to nothing apart from perhaps certain basic axiomatic truths and certain truths about our current experience' – is the appropriate speech to make.

Lewis's framework is widely considered as very concessive towards global external world scepticism. Rather than rejecting external world scepticism *tout court*, it makes the familiar speech of the external world sceptic come out true in various contexts where our epistemological imagination runs free. We want to argue, however, that when one carefully thinks through Lewis's framework in the light of our previous discussion, it does not, after all, vindicate the idea that there are myriad contexts of epistemological conversation in which global external world scepticism is the winner.

We should begin by noting that Lewis's framework provides the beginning of a story as to why richly described sceptical possibilities might play a distinctive epistemological role that is not played by the mundane observation that external world propositions are not entailed by the facts of our experience. That last observation does not bring any possibility – in Lewis's sense – into view. And so that last observation does not provide fodder for the rule of Attention. Just pointing out that the facts of our experience do not necessitate the proposition that we have hands does not, *ipso facto*, bring into view any (uneliminated) possibility – in Lewis's sense – in which we lack hands. And so that speech does not, by itself, throw us into a context where 'I know I have hands' expresses a falsehood. By contrast, the familiar brain in a vat narrative does make salient an uneliminated possibility in which we have no hands. And so, by the rules of Lewis's game, by making that possibility salient, one is thrown into a context where the speech 'I know I have hands' expresses a falsehood. So far, so good. But does Lewis's framework point to contexts where the distinctive speeches of the global external world sceptic are correct?

One point that we have already made is pertinent here. It is at least somewhat intuitive that attending to a detailed brain in a vat narrative puts one in a context where various mundane knowledge ascriptions are inappropriate. But suppose instead one attends to a detailed narrative in which things are roughly as we think they are but where a panda is transported to Doncaster and glued to six lawnmowers. It is far less plausible that one is thereby thrown into a context where one is hesitant to say, 'I know there is no panda glued to six lawnmowers in Doncaster'. If one had not already embraced the general thesis that nothing is known unless our phenomenal evidence entails it, the detail that upgrades the barebones panda-stuck-to-lawnmowers hypothesis to something rich enough to count as a possibility in Lewis's sense does not seem to have the kind of context-shifting effect that Lewis's framework predicts.

But even pretending that this tiresome exercise of the imagination did have the context-shifting effect, Lewis's framework falls considerably short of pointing to contexts where the speech of the external world global sceptic comes out true. Let us return to the opening sentence of Lewis's paper, one that adumbrates the anti-sceptical point of view: 'We know a lot'. Just to make clear that axiomatic truths and truths about experience are neither here nor there as far as external world scepticism goes, let us consider the refinement 'We know a lot about

the external world'. This is a quantificational claim – it says, in effect, that there are a lot of propositions about the external world such that we know them. Suppose some epistemologists let their fantasies rip. They attend to brain in vat possibilities, demon possibilities, dream possibilities. You know the drill. What becomes of the sentence 'We know a lot about the external world' in such contexts? Clearly, Lewis thinks that such contexts make for the falsity of that speech. But it is quite striking that his framework does not secure any such result. Some of our earlier examples can be very instructive here. Suppose Jack and Jill have been to epistemology class and exercised their sceptical imagination in familiar ways. We can get a fix on the status of 'There are a lot of propositions p about the external world such that I know p ' at a context by looking at the status of the open sentence 'I know p ' relative to various assignments of propositions to the variable ' p ' at that context. Let us do just that. Consider the position that Jack is not holding an armadillo deep under the ocean. We take it that Jack and Jill will not have attended to possibilities in which their experience matches the way it actually is but where they are holding an armadillo deep below the ocean. So the rule of Attention cannot kick in directly. There are, of course, possibilities where, say, Jill's experience is exactly as it is in the actual world but where she is holding an armadillo deep under the ocean. But, within Lewis's framework, the key issue is whether or not those possibilities are properly ignored. Similarly, for the proposition expressed by 'I am not having lunch with Barack Obama right now' and so on. And since possibilities incompatible with the relevant negated proposition are not attended to, the rule of Attention can do nothing to refute the hypothesis that they are properly ignored, even if Jill has let her fantasies rip in familiar ways. Note, moreover, that even if Jill were a brain in a vat, the same would be true. By the rule of Actuality, if Jill is a brain in a vat, a certain kind of familiar sceptical possibility – one in which she is a brain in a vat – cannot be properly ignored. But insofar as she is a brain in a vat, as we noted above, she is certainly not holding an armadillo deep under the ocean. And so, on the hypothesis that she is a brain in a vat, neither the rule of actuality nor the rule of Attention can do anything to indict the hypothesis that the ocean-armadillo scenario is properly ignored.

What about the rule of resemblance? As far as we can tell, this is the most obvious rule to turn to if one is trying to extract sceptical results for the ocean-armadillo or the lunch-with-Barack-Obama scenario within the framework that Lewis proposes. But we really don't see how this is supposed to go. Consider the scenario where Jill is having lunch with Barack Obama despite appearing to be in intense conversation with Jack in a room and where it appears to Jill that she is in a room where there are neither waiters nor food. That does not seem to be similar at all to a scenario in which Jill is a brain in a vat. Of course there is one obvious dimension of relevance. Both the relevant brain in a vat scenario and the ocean armadillo scenarios are ones where Jill's experience matches her actual experience. (That's what makes them uneliminated possibilities in the first place.) But Lewis is very explicit that this kind of resemblance is not intended to be enough to make one possibility salient resemble another in the sense intended by the rule. He writes: 'Actuality is a possibility uneliminated by the subject's evidence. Any other possibility W that is likewise uneliminated by the subject's evidence thereby resembles actuality in one salient respect: namely, in respect of the subject's evidence. That will be so even if W is in other respects very dissimilar to actuality – even if, for instance,

it is a possibility in which the subject is radically deceived by a demon' (Lewis, 556). Lewis then proceeds to note that if one allows this aspect of resemblance to count as salient resemblance then every context will be one in which sceptical speeches are true. After all, by the rule of Actuality, the actual situation is relevant. And any possibility in which one's evidence matches one's actual evidence will be similar in this respect to actuality. And so, if one is so liberal to allow that matching experience counts as resemblance enough for the rule of Resemblance to kick in, then no possibility in which one has matching experience can be properly ignored in any context.¹⁵

Counting a match of experience as salient resemblance in the sense intended by the rule would thus radically misconstrue Lewis's framework. But once that misconstrual is set to one side, it is very hard to see how an uneliminated possibility in which one is a brain in a vat is supposed to saliently resemble an uneliminated possibility in which one is embodied and, despite appearances, is in a restaurant having lunch with Barack Obama. (Indeed, as noted above with regard to the ocean-armadillo example, if the question is whether one of us knows that they are not having lunch with Barack Obama, brain-in-a-vat scenarios are very much not to the sceptical point since they are possibilities in which one is, indeed, not having lunch with Barack Obama.) In sum, it is very hard to see how the rule of Resemblance is supposed to make epistemic trouble for our focal negated propositions once, say, brain-in-a-vat possibilities are attended to. (And this is so even if one happens to be a brain in a vat.) Very well. Perhaps Jack and Jill got to properly ignore the ocean-armadillo scenario. But we can raise it to salience right now. (We could, after all, describe an uneliminated possibility in which one is underwater holding an armadillo.) But that piecemeal approach will do nothing to vindicate the quantificational speech distinctive of global external world scepticism. Even having attended to the ocean-armadillo scenario, there will still it seems be all sorts of external world beliefs (like the Obama-lunch one) where there is no sceptical counter-possibility that is being attended to. And so, *prima facie*, for all the rule of Actuality and the rule of Attention say, extending one's fantasies to the ocean-armadillo scenario can leave all sorts of other uneliminated possibilities properly ignored. And if the rule of Resemblance doesn't make scepticism spread from brain-in-a-vat scenarios to ocean-armadillo scenarios, it surely does not make scepticism spread from ocean-armadillo scenarios to Obama-lunch scenarios. And so on. In sum, even having let our fantasies rip in familiar ways, and even after throwing a few more kooky fantasies for good measure, it may well be that the open sentence '*p* is an external world proposition and I know *p*' is true relative to myriad assignments to the variable '*p*'. Thus, perhaps surprisingly, it seems that the sentence 'There are many propositions about the external world that I know' is true relative to any such context and so the distinctive global sceptics claim 'I know nothing about the external world' will be false relative to any such context. (And so on a perfectly reasonable construal of 'a lot', the opening sentence of Lewis's paper 'We know a lot' will be true relative to any such context.)

¹⁵ Note that this liberal approach to resemblance is tantamount to the view that 'know' context invariably expresses a relation that one can only hold to propositions entailed by the facts of one's inner experience. Adopting such an approach as a starting point would be tantamount to the very kind of tendentious strategy we set to one side at the outset.

Are there *any* contexts in which the global external world sceptic wins? Let us first think about what it would take for the combination of the rule of Actuality and the rule of Attention to all by itself vindicate a global external world sceptical speech. Take all one's external world beliefs that are true at actuality. What would be needed was an exercise in attention so prodigious that, for any p that is the content of such a belief, one attended to an uneliminated possibility in which not- p . This, we take it, is an exercise in attention so demanding that no human being has come even close to accomplishing it – one would have to simultaneously attend to an insane number of uneliminated possibilities in order to vindicate global external world scepticism on the grounds of the Actuality and Attention alone.¹⁶

Now, as we noted, the rule of Resemblance can be of some help to the sceptic. But here is what has to happen for the global sceptic to win (at some context of discourse) on the grounds of Actuality, Attention *and* Resemblance: For any external world proposition p that is true and for which one does not attend to a sceptical counterpossibility, one has at least to attend to a sceptical possibility that saliently resembles a sceptical possibility in which not- p . And assuming (as Lewis intends) that salient resemblance is not understood in a lax way, this is still an extremely demanding task. As we have seen, attending to brain in a vat possibilities does not yet amount to attending to a possibility that saliently resembles an uneliminated possibility in which one is holding an armadillo deep under the ocean. Meanwhile, attending to the latter does not yet amount to attending to a possibility that saliently resembles an uneliminated possibility in which one is having lunch with Barack Obama. It is hard to get a fully precise fix on what Lewis's framework would require here because – as he was aware – the ideology of salient resemblance is left rather underdeveloped. But it seems already clear enough that familiar sceptical conversations in epistemology are ones in which the protagonists have not done anywhere near enough to vindicate global external world sceptical speeches. And, at least on many natural precisifications of 'salient resemblance' it is at least not clear whether we even have the attentional capacity to get into a context where global external world sceptical speeches come out true.

We have focused on three of Lewis's rules. There is one more that is potentially relevant and instructive, namely, the rule of belief. We focus here on Lewis's more refined version of the rule, one that is articulated in terms of degrees of belief (i.e., credence): If one gives or ought to give sufficiently high credence to a possibility, then it is not properly ignored. Now the sceptic may try to get some sceptical leverage out of this rule. Let's suppose that the threshold for 'high enough' is not set that high. Plausibly, even if one doesn't believe p (in the binary sense of belief), in a setting where one takes p seriously – or ought to take p seriously – p is not properly ignored. If one's car is spinning out of control and one takes seriously that one

¹⁶ It may be argued that Lewis is aware of the limits of his attention rule when he suggests that epistemologists on a bushwalk attending to sceptical possibilities nevertheless know 'where they are and where they are going' (Lewis, 565). To resolve the apparent conflict with the main thrust of the paper, Lewis proposes a compartmentalization resolution where one compartment knows while the other (epistemological compartment) remains ignorant. (He brushes aside, as not altogether felicitous, the question of what the compartmentalized agents know.) Whether this idea can be fully worked out is not clear – see Spectre 2019 – but note that even if there is a path for Lewis's compartmentalization idea, it will not vindicate global scepticism even within the bushwalkers' epistemological compartments. Not by the Attention rule or even, as we argue more generally, by a combination of Lewis's rules.

is about to die, one has not properly ignored the possibility that one dies, and this is so even if one gives around .2 credence to the proposition that one will die. With this 'ought to take seriously' gloss on 'ought to give high enough credence' in mind, one can imagine the sceptic trying to find an entering wedge via the rule of belief so construed. People really do have vivid dreams. So perhaps one really ought to give non-negligible credence to the hypothesis that one is dreaming right now. And perhaps, given that one ought to take Boltzmann brain cosmology seriously, one ought to give non-negligible credence to the hypothesis that one is a boltzmann brain (a hypothesis somewhat analogous to the brain in a vat hypothesis). Now one might fight these particular judgments. But what bears emphasis is that even if they are conceded they do nothing to vindicate global external world scepticism. Even if one concedes that one ought to give non-negligible credence to the hypothesis that one is dreaming, it hardly follows that one ought to give non-negligible credence to – or in any reasonable sense take seriously – the hypothesis that one is holding an armadillo deep under the ocean.¹⁷ Note that so long as the probability of having hands is less than one, it is easy to construct propositions that are entailed by the proposition that one has hands that have a higher probability than that proposition: just negate one of the subcases of lacking hands. E.g., it is not the case that I am a handless person having lunch with a panda in Doncaster. Set the probabilistic threshold for 'ought to take seriously' at any non-zero threshold you like. There will be subcases of lacking hands that one ought not to take seriously and so their negations will not be threatened by a rule of sufficient credence indexed to that threshold.¹⁸ As with the rule of Attention and the rule of Resemblance, the rule of 'sufficient credence' can do some sceptical work, but neither this rule in isolation, nor in combination with the other rules make for identifiable contexts in which the pessimistic speeches of the global external world sceptic come out correct. The case study of Lewis was not chosen arbitrarily. Lewis is often thought of as someone who offers a framework that is highly concessive to the sceptic. What is striking is that having absorbed the lessons of our earlier discussion, one can see that Lewis does not clearly provide even a single context in the history of epistemological conversations where the global external world

¹⁷ Similar remarks apply to the idea that the better candidate explanations of the data cannot be properly ignored. Now of course one can – like Russell 1912, Vogel 1990, and others – reasonably take the position that the vatted brain hypothesis is much worse an explanation than more mundane competitors. But what bears emphasis is that even if one is so concessionary as to think dreaming and the brain in a vat hypotheses are right up there by explanatory standards, it hardly seems that ocean-armadillo theories and so on will pass muster. And so, once again, we see that even highly concessionary views are nowhere near concessionary enough to vindicate global external world scepticism. This kind of theme can be replicated for all sorts of other epistemological ideas. Suppose one thought that sufficiently normal scenarios cannot be properly discounted – cf. Goodman and Salow forthcoming – and was concessionary enough to regard the boltzmann brain scenarios as sufficiently normal. Still, ocean-armadillo scenarios should hardly thereby count as normal.

¹⁸ Analogous remarks apply to Stewart Cohen's (1988) version of contextualism. According to his vision, the key contextual parameter is a shiftable threshold of evidential support for knowledge. Accordingly brain in a vat discussions tend to raise the threshold in such a way that the proposition expressed by 'I have hands' does not get over the high threshold set by those contexts. But note that so long as the threshold falls short of simply requiring entailment by the facts of one's inner life, setting a high threshold that various mundane beliefs fail to surpass is quite compatible with such a threshold still being surpassed by the belief that one is not a handless person having lunch with a panda in Doncaster. None of this is a problem for Cohen's framework. It merely reminds us that his framework does not entail that contexts where 'I don't know whether or not I am a brain in a vat' comes out true are contexts where global external world scepticism is vindicated.

sceptic wins the day.¹⁹

References

- Cohen, Stewart (1988). 'How to Be a Fallibilist'. *Philosophical Perspectives* 2, pp. 91–123. DOI: [10.2307/2214070](https://doi.org/10.2307/2214070).
- Descartes, René (1641). *Meditations on First Philosophy*. Trans. by John Cottingham, Robert Stoothoff, and Dugald Murdoch. in *The Philosophical Writings of Descartes: Volume 2*, Cambridge University Press (1984).
- Dretske, Fred I. (1970). 'Epistemic Operators'. *Journal of Philosophy* 67.24, pp. 1007–1023. DOI: [10.2307/2024710](https://doi.org/10.2307/2024710).
- Goodman, Jeremy and Bernhard Salow (forthcoming). 'Epistemology normalized'. *Philosophical Review*.
- Halbental, Moshe (2014). *Maimonides*. Trans. by Joel Linsider. Princeton University Press.
- Hawthorne, John and Dani Rabinowitz (2017). 'Knowledge and false belief'. *Explaining knowledge: New essays on the Gettier problem*, pp. 325–344.
- Hickey, Lance P. (n.d.). 'The Brain in a Vat Argument'. In: *The Internet Encyclopedia of Philosophy*. URL: <https://iep.utm.edu/brain-in-a-vat-argument/>.
- Lewis, David K. (1996). 'Elusive Knowledge'. *Australasian Journal of Philosophy* 74.4, pp. 549–567. DOI: [10.1080/00048409612347521](https://doi.org/10.1080/00048409612347521).
- Magidor, Ofra (2018). 'How Both You and the Brain in a Vat Can Know Whether or Not You Are Envatted'. *Aristotelian Society Supplementary Volume* 92.1, pp. 151–181. DOI: [10.1093/arisup/aky009](https://doi.org/10.1093/arisup/aky009).
- Maimonides (1190). *The Guide of the Perplexed*. Trans. by Shlomo Pines (1963) from Jewish Arabic. University of Chicago Press.
- Markosian, Ned (2014). 'Do you know that you are not a brain in a vat?' *Logos & Episteme* 5.2, pp. 161–181.
- Nozick, Robert (1981). *Philosophical Explanations*. Harvard University Press.
- Pryor, James (2000). 'The Sceptic and the Dogmatist'. *Noûs* 34.4, pp. 517–549. DOI: [10.1111/0029-4624.00277](https://doi.org/10.1111/0029-4624.00277).
- Russell, Bertrand (1912). *The Problems of Philosophy*. London, England: William & Norgate.

¹⁹ For helpful correspondence, we would like to thank Daniel Berntson, Tim Button, Kevin Dorst, David Enoch, Moshe Halbental, and Ofra Magidor. Also, for much worthy feedback, many thanks to Matti Eklund and the participants at Uppsala University's Higher Seminar in Theoretical Philosophy and to Kathrin Glüer and the participants at Stockholm University's CLLAM SEMINAR in Logic, Language, and Mind.

Funding: Levi Spectre's research was supported by the Swedish Riksbankens Jubileumsfond as part of the 'Knowledge Resistance: Causes, Consequences, and Cures' research project (Reference Number: M18-0310:1).

- Schantz, Richard (2011). 'Skepticism and anti-realism'. In: *The Routledge Companion to Epistemology*. Ed. by S. Bernecker and D. Pritchard. Routledge, pp. 477–487.
- Spectre, Levi (2019). 'Compartmentalized knowledge'. *Philosophical Studies* 176.10, pp. 2785–2805.
- Vogel, Jonathan (1990). 'Cartesian Skepticism and Inference to the Best Explanation'. *Journal of Philosophy* 87.11, pp. 658–666. DOI: [10.5840/jphil1990871123](https://doi.org/10.5840/jphil1990871123).
- Warfield, Ted A (2005). 'Knowledge from falsehood'. *Philosophical Perspectives* 19, pp. 405–416.
- Williamson, Timothy (2022). *The Philosophy of Philosophy*. John Wiley & Sons.