

EPISTEMIC DISJUNCTIVISM AND THE EVIDENTIAL PROBLEM*

José Zalabardo, UCL

1 Epistemic Disjunctivism

Duncan Pritchard's recent book (Pritchard 2012) provides a lucid articulation and defence of epistemological disjunctivism (ED). He offers a characteristically clear formulation of the view:

In paradigmatic cases of perceptual knowledge an agent, *S*, has perceptual knowledge that ϕ in virtue of being in possession of rational support, *R*, for her belief that ϕ which is both *factive* (i.e., *R*'s obtaining entails ϕ) and *reflectively accessible* to *S*. (13)

This factive, reflectively accessible rational support that postulated by ED consists in *S*'s seeing that ϕ :

So when one has paradigmatic perceptual knowledge of a proposition, *p*, one's reflectively accessible rational support for believing that *p* is that one *sees that* *p*. Seeing that *p* is factive, however, in that if it is the case that one sees that *p* then *p* must be true. (14)

By way of a declaration of interests, let me start by saying that I don't think ED provides a satisfactory account of perceptual knowledge. I think that epistemological disjunctivism would be highly implausible unless metaphysical disjunctivism were true, but metaphysical disjunctivism is false. However, this point is irrelevant to the assessment of Pritchard's contribution. He is very clear about what he hopes to achieve and what he doesn't hope to achieve, and the interaction between epistemological and metaphysical disjunctivism is one of the issues that he explicitly leaves to one side.

Pritchard has two main goals concerning ED. First, he puts forward solutions to three fundamental problems that the view may seem to face. Second, he argues that, if ED were true, we would have at our disposal a satisfactory strategy for dealing with the problem of radical scepticism. Here I am going to concentrate on the second issue—on the promised anti-sceptical dividends of ED. I am going to argue that, contrary to what Pritchard claims, ED doesn't sustain a successful anti-sceptical strategy. I shall contend, in particular, that the treatment of scepticism that Pritchard puts forward is unsatisfactory. My argument won't rely on my rejection of ED. I am going to argue that even if ED were true, the anti-sceptical strategy that Pritchard puts forward on behalf of the disjunctivist would have to be rejected.

2 The sceptical problem

Pritchard discusses a form of sceptical reasoning based on sceptical hypotheses. He concentrates on Hilary Putnam's brain-in-a-vat hypothesis, according to which your brain has been removed from the body and placed in a vat of nutrients which keeps the brain alive, and its nerve endings have been connected to a computer which causes you to have the illusion that everything is perfectly normal (Putnam 1981). The form of sceptical reasoning with which Pritchard engages derives the conclusion

* I am grateful to Mike Martin and Lee Walters.

that you don't know many of the things that you take yourself to know from two claims concerning the brain-in-a-vat hypothesis: first, there are many things that you take yourself to know that you wouldn't know unless you knew that the brain-in-a-vat hypothesis doesn't obtain; second, you don't know that the hypothesis doesn't obtain. The first claim is supported by the plausible principle that knowledge is closed under known entailment or, in Pritchard's preferred version, under competent deduction. If you know that p entails q , or if you have competently deduced q from p , then you can't know p unless you also know q . The knowledge claims that are threatened by the argument concern propositions of which you know that they entail that you are not a brain in a vat, or from which you can competently deduce this conclusion. Many propositions that you take yourself to know satisfy this condition. Take, for example, the proposition that you have hands. You know that this entails that you are not a brain in a vat (if you were one, you wouldn't have hands). Hence the sceptic can argue as follows:

1. You don't know that you are not a brain in a vat.
2. If you don't know that you are not a brain in a vat you don't know that you have hands.

Therefore:

3. You don't know that you have hands.

And parallel arguments can be deployed with respect to many other propositions that you take yourself to know. The premises of this argument logically entail the conclusion. Hence resisting the conclusion would require rejecting at least one of the premises.

Rejecting the second premise was the strategy favoured by Robert Nozick (Nozick 1981). This allowed him to concede to the sceptic that we don't know that we are not envatted brains while holding on to the everyday knowledge claims targeted by the argument. Pursuing this strategy requires abandoning the closure principle, but Nozick didn't see this as a problem, as the principle is incompatible with his preferred analysis of knowledge.

However, many others see the rejection of the closure principle as too high a price to pay, and are therefore forced to treat premise 2 of the sceptical argument as unassailable. For them (for us) resisting the sceptical conclusion requires rejecting premise 1. If we take this line, in order to vindicate your knowledge claims in the face of the sceptical challenge we need to provide an account of how you can know you are not a brain in a vat.

This is the point at which Pritchard expects ED to make a contribution to the anti-sceptical cause. If we accept ED, Pritchard claims, we have at our disposal a satisfactory account of how we can know that we are not envatted brains.

3 The evidential problem

Pritchard presents the ED-based account of how you know that you are not a brain in a vat as an answer to a specific challenge to the claim that this knowledge is possible. He introduces the challenge in connection with Dretske's zebra case (Dretske 1970), where closure also forces us to make a knowledge attribution that might seem hard to explain. Zula goes to the zoo and takes a good look at a zebra. Under normal circumstances, we would want to say that Zula can come to

know that the animal is a zebra by looking at it. But now closure poses a problem. We can assume that Zula knows that the proposition that the animal is a zebra entails that it isn't a mule cleverly disguised by the zoo authorities to look like a zebra, or that she has competently deduced the latter from the former. Hence, according to closure, she won't know that it is a zebra unless she knows that it isn't a cleverly disguised mule. But the claim that Zula knows that the animal isn't a cleverly disguised mule is problematic. Pritchard first raises a difficulty for this knowledge ascription based on the Discrimination Principle, but then argues that the difficulty is spurious, as the principle has to be rejected. He then raises a more enduring difficulty, to which he refers as the *evidential problem*:

[...] the closure principle forces us to regard Zula as knowing that what she is looking at is not a cleverly disguised mule, and yet she doesn't appear to have any good supporting evidence for this knowledge. Moreover, although we might not claim that all knowledge must be evidentially grounded, all will surely agree that knowledge of a proposition like this must be so grounded. (74)

The evidential problem arises from the claim that Zula wouldn't know that the animal is not a cleverly disguised mule unless she had good supporting evidence for this proposition. The problem consists in identifying the evidence that can play this role.

Pritchard argues that the evidence that Zula obtains just by looking at the zebra ("evidence regarding the visual scene before her" (77)) won't do the job, since this evidence "does not seem to speak at all to the cleverly disguised mule alternative" (75). Pritchard is certainly right about this. I think that the phenomenon can be fruitfully explained in terms of the notion of incremental confirmation. I take it that we can assume that in order to count as 'good supporting evidence' for a hypothesis H , a piece of evidence E has to provide a substantial amount of incremental confirmation for H . In order for E to provide any amount of incremental confirmation for H , E has to be more probable if H is true than if H is false. I.e. it has to be the case that $p(E|H) > p(E|\sim H)$.¹ In normal conditions the evidence that Zula obtains by looking at the animal will satisfy this condition with respect to the hypothesis that it is zebra. The animal is more likely to look like this if it is a zebra than if it isn't.² Crucially, however, the condition is not satisfied with respect to the hypothesis that the animal is not a cleverly disguised mule. The animal is not more likely to look like this if it isn't a cleverly disguised mule than if it is. In fact it is more likely to look like this if it *is* a cleverly disguised mule than if it isn't—this, after all, is what a cleverly disguised mule looks like.

Pritchard proposes to deal with the problem by considering other evidence that Zula can be expected to have:

One might reason, for instance, that there would be no point in such a deception, that it would be costly and time-consuming without bringing any comparable benefit, that it would be easily found out, and then the zoo-owner would be subject to penalties, and so on. (79)³

¹ This criterion is equivalent to Carnap's original criterion for incremental confirmation: $p(E|H) > p(H)$ (Carnap 1962).

² Notice that this is compatible with the existence of possibilities in which the animal isn't a zebra but still looks like this—e.g. cleverly disguised mules.

³ A solution along these lines was originally proposed by Jonathan Vogel (Vogel 1990). I have endorsed this solution in (Zalabardo 2012: 141). The only difference between Pritchard's treatment and the approach that

General evidence of this kind will not be afflicted by the problem that invalidates the evidence that Zula obtains just by looking, and can in principle be invoked to solve this instance of the evidential problem.

4 The evidential problem and scepticism

The claim that you know that you are not a brain in a vat faces another instance of the evidential problem. If knowing that you are not a brain in a vat requires an evidential basis, then the claim that you have this knowledge needs to be backed up by an account of which evidence can play this role. We are owed, Pritchard writes, “an explanation of the evidential basis for this anti-sceptical knowledge (or else an explanation of how it could feasibly lack an evidential basis)” (117).

For the non-disjunctivist, the proposition that you are not a brain in a vat cannot inherit the requisite evidence from the everyday propositions by which you know it to be entailed. Take for example your visual evidence for the proposition that you have hands. For the non-disjunctivist, the evidence that you obtain by looking at your hands is the same evidence that you would have if you were a brain in a vat under the false impression of looking at your hands. Under normal circumstances this evidence can be expected to provide incremental confirmation for the proposition that you have hands. Things are more likely to look the way they do when you seem to look down at the ends of your arms if you have hands than if you don't.⁴

But this visual evidence doesn't provide incremental confirmation for the hypothesis that you are not a brain in a vat. The probability that things look the way they do when you seem to be looking down at the ends of your arms is just as high if you are a brain in a vat as if you are normally embodied. The computer to which the brain's nerve endings are connected ensures that this is so.

So far the situation that you face with respect to the proposition that you are not a brain in a vat is parallel to Zula's situation with the proposition that the animal is not a cleverly disguised mule. However, Pritchard argues, the move that solves the evidential problem for Zula is not available to you. None of the evidence at your disposal will support the hypothesis that you are not a brain in a vat. Here is how Pritchard characterises the disanalogy:

For example, just as we appealed to Zula's background knowledge about the plausibility of the cleverly disguised mule hypothesis in order to accord her with the appropriate favouring evidence in support of her beliefs, suppose we appealed to similar background knowledge that our agent in the radical sceptical case might be thought to have in order to accord her with the requisite favouring evidence.⁵ After all, we might naturally suppose that the agent in the sceptical case possesses all sorts of good reasons for thinking that the BIV hypothesis

Vogel and I recommend is that Pritchard needs Zula's knowledge that the animal is not a cleverly disguised mule to be supported by evidence that also supports her knowledge that the animal is a zebra. This is in order to satisfy his Evidential Transmission Principle (75). Neither Vogel nor I support this restriction, and I believe that the Evidential Transmission Principle is false. Zula can know that the animal is a zebra on the basis of the evidence that she obtains by looking at it, but this evidence cannot ground knowledge that the animal is not a cleverly disguised mule, even if Zula has competently deduced this from the premise that the animal is a zebra.

⁴ Once again, this is compatible with the existence of possibilities in which things look like that but you don't have hands—e.g. envatment scenarios.

⁵ Pritchard's distinction between favouring and discriminating evidence won't concern us here.

is implausible (e.g., that this scenario is incompatible with the current state of technology, that it is not plausible that someone would have a motive to pursue such ‘envatments’, and so on).

The problem, however, should now be manifest. We can unproblematically appeal to such background knowledge in the case of Zula precisely because the error-possibility at issue does not call into question this background knowledge. The same is not true when we focus on the radical sceptical hypothesis. For if the agent were indeed a BIV, then she would lack such background knowledge. Accordingly, it would be contentious to appeal to such evidence as a means of showing that the agent’s beliefs have the requisite epistemic support. (85)

According to Pritchard, the evidential problem concerning the hypothesis that you are not a brain in a vat cannot be solved by appeal to the background knowledge that you think you have. This is because if the hypothesis were false—if you were a brain in a vat—you wouldn’t have this knowledge. I don’t think this is the real reason why background evidence is disqualified, but I agree with Pritchard that this instance of the evidential problem cannot be solved in this way. More generally, I believe that on a non-disjunctivist conception of perceptual knowledge, the problem has no solution. If ED sustained a satisfactory solution, this would have to count as a major advantage for the view.

5 Disjunctivism to the rescue

Pritchard claims that ED can solve the evidential problem. We can appreciate the change that ED brings about if we compare the standard construal of the Zula case with the ED-based alternative. On the construal of the case that we have assumed until now, the evidence that Zula obtains, just by looking, for the proposition that the animal is a zebra is compatible with the animal not being a zebra—she could have the evidence that she has even if the animal were not a zebra. But on an ED-based construal of the case, if we think of it as a paradigmatic case of perceptual knowledge, that Zula sees that the animal is a zebra provides her with factive rational support for believing that the animal is a zebra. She couldn’t have the evidence that she has if the animal weren’t a zebra. This, according to Pritchard, makes all the difference:

[...] because her reflectively accessible rational basis for this belief is factive this itself suffices to directly exclude the target error-possibility. (125)

Pritchard is highlighting an important difference between deductive and non-deductive evidence. As we have seen both in Zula’s case and with the brain-in-a-vat hypothesis, if E provides adequate but non-deductive evidential support for H1, E might not provide adequate support for H2 even if H1 entails H2. The way the animal looks provides adequate support for the proposition that it is a zebra, but not for the proposition that it isn’t a cleverly disguised mule; how things look to you when you seem to be looking towards the ends of your arms provides adequate support for the proposition that you have hands, but not for the proposition that you are not a brain in a vat. However, with deductive evidence this is not a possibility. If E provides deductive evidence for H1, and H1 entails H2, then E also provides deductive evidence for H2. On the ED-based construal of the case, Zula has factive perceptual evidence for the proposition that the animal is a zebra. Since this proposition

entails that the animal is not a cleverly disguised mule, it follows that her perceptual evidence also provides deductive support for the latter proposition.

Pritchard concludes from this that on the ED-based construal of the case, the evidential problem concerning Zula's knowledge that the animal is not a cleverly disguised mule can be solved in terms of her perceptual evidence. The (factive) evidence that she obtains just by looking can be invoked to explain the evidential basis for her knowledge that the animal is not a cleverly disguised mule.

The same situation obtains with respect to the brain-in-a-vat hypothesis. According to ED, when you look to the ends of your arms, you obtain factive rational support for your belief that you have hands—namely that you see that you have hands. Since your seeing that you have hands provides deductive evidence for the proposition that you have hands, and this proposition entails that you are not a brain in a vat, it follows that your seeing that you have hands also provides deductive evidence for the proposition that you are not a brain in a vat. This is how Pritchard characterises the situation of a subject, John, in the circumstances that I have described:

John has reflectively available to him rational support for his belief that p which (he is fully aware) *entails* p and hence which entails not-BIV as well. (123)

Hence, Pritchard claims, ED puts at our disposal a solution to the evidential problem with respect to your anti-skeptical knowledge. Your factive evidence for, say, the proposition that you have hands can be invoked to explain the evidential basis for your knowledge that you are not a brain in a vat.⁶

6 A gap in the argument

I've claimed in the previous section that Pritchard is right in endorsing the following principle:

- A. If E provides deductive evidence for H1, and H1 entails H2, then E provides deductive evidence for H2.

And by virtue of this principle, ED entails:

- B. Your seeing that you have hands provides deductive evidence for the proposition that you are not a brain in a vat.

If this is to amount to a solution to the evidential problem, we have to be able to conclude the following:

- C. Your seeing that you have hands can constitute the evidential basis for your knowledge that you are not a brain in a vat.

This conclusion would be unproblematic in the presence of the following principle:

⁶ Pritchard embeds his proposal for dealing with the evidential problem in a 'two-tiered' relevant-alternatives theory of knowledge. On this theory, the kind of evidence that is needed to rule out a hypothesis that is incompatible with your belief (discriminating or supporting evidence) will depend on whether the hypothesis is epistemically motivated. None of this will affect the points that I want to make.

- D. If S knows p and S knows that p entails q, then p can constitute the evidential basis for S's knowledge that q.

You know, by reflection, that you see that you have hands, and you know that your seeing that you have hands entails that you are not a brain in a vat. Hence, in the presence of D, C follows directly.

I think, however, that intuition is firmly against the universal validity of D. There are cases in which S knows p and S knows that p entails q, but it seems intuitively wrong to say that p can constitute the evidential basis for S's knowledge that q.

We encounter intuitive counterexamples to D in the cases that we have been considering. Zula knows that the animal is a zebra, and she knows that this entails that it isn't a cleverly disguised mule, and yet it seems intuitively wrong to say that the animal being a zebra can constitute the evidential basis for Zula's knowledge that it isn't a cleverly disguised mule. And you know that you have hands, and you know that this entails that you are not a brain in a vat, but once again it seems intuitively wrong to say that your having hands constitutes the evidential basis for your knowledge that you are not a brain in a vat.⁷

It seems to me that the clash with intuition is undeniable.⁸ I realise, however, that this doesn't amount to a refutation of D. We need to be open to the possibility of accepting epistemic principles that don't provide a perfect match for our intuitions, and it might turn out that D provides the optimal account of knowledge based on deductive evidence, and that recalcitrant intuitions will have to be discarded as mistaken.

Whether or not we should take this line will depend to a large extent on whether we can provide a plausible diagnosis of the shortcoming exhibited by the cases that intuition doesn't want to treat as genuine instances of evidence-based knowledge—a necessary condition for evidence-based knowledge that is violated in these cases but satisfied in intuitively acceptable instances of D. At this point we have a wide range of options at our disposal. They can be found in the extensive recent literature on the transmission of warrant by inference.⁹ This is not the place to review these proposals. What I want to do is to present what I regard as the most satisfactory option, and to consider how it affects the anti-sceptical power of ED.

The necessary condition for evidence-based knowledge that I want to consider derives from a condition imposed by Nozick on when inferring q from p yields knowledge of q: according to Nozick, this requires that if q were false you wouldn't believe p (Nozick 1981: 231). If it is the case that if q were false you would still believe p, then even if you know p, and p supports or even entails q, and you know that it does, inferring q from p won't enable you to know q. I favour a version of this condition formulated in terms of conditional probabilities instead of counterfactuals: what evidential

⁷ The intuition that I am invoking here is in principle compatible with Pritchard's Evidential Transmission Principle. One could maintain that your knowledge that you are not a brain in a vat cannot be based on the evidence provided by the proposition that you have hands but that it can be based on the same evidence on which your knowledge of this proposition is based. In any case, as I've indicated above, I think that the Evidential Transmission Principle is false.

⁸ For Pritchard's purposes, it would be enough if D held in cases in which the challenge to q has not been epistemically motivated. But the intuitive counterexamples that I have presented would also invalidate this restricted version of the principle.

⁹ See, e.g. (Wright 2000, 2002, 2003; Davies 1998).

knowledge of a hypothesis requires is that you are more likely to believe the evidence if the hypothesis is true than if the hypothesis is false (Zalabardo 2012: 98). The result is a necessary condition for knowing q on the basis of the evidential support provided by p :

- E. S can know q on the basis of the evidential support provided by p only if S is more likely to believe p if q is true than if q is false.¹⁰

This principle succeeds in ruling out the intuitively unacceptable instances of D that we have considered.¹¹ Zula would be just as likely to believe that the animal she is looking at is a zebra if it were a cleverly disguised mule as if it were a zebra. This is what the cleverness of the disguise amounts to. It follows that, according to E , Zula can't know that the animal is a cleverly disguised mule on the basis of the evidential support provided by her belief that it is a zebra. And you are as likely to believe that you have hands if you are a brain in a vat as if you are normally embodied, as it is an essential feature of the brain-in-a-vat scenario, as it figures in recent discussions of scepticism, that envatted brains have the same perceptual beliefs as their embodied counterparts.¹² It follows that according to E , you can't know that you are not a brain in a vat on the basis of the evidential support provided by your belief that you have hands.

Furthermore, I want to claim that E is not simply extensionally correct. It's not just that it succeeds in ruling out the instances of D that intuition rejects. In addition, it provides a very appealing diagnosis of what is wrong with these cases. The reason why we don't want to say that Zula can know that the animal is not a cleverly disguised mule on the basis of the evidential support provided by her belief that it is a zebra is that if the animal were a cleverly disguised mule Zula would be no less likely to believe that it is a zebra—it wouldn't be a zebra, but Zula would be just as likely to believe that it is one; and the reason why we don't want to say that you can know that you are not a brain in a vat on the basis of the evidential support provided by your belief that you have hands is that if you were a brain in a vat you would be just as likely to believe that you have hands—you wouldn't have hands, but you would be just as likely to believe that you do.¹³

Let's consider now how E bears on the ED-based solution to the evidential problem. Consider Zula first. In order to determine whether the proposal complies with E , we need to ask the following question: would Zula be less likely to believe that she sees that the animal she's looking at is a zebra if it were a cleverly disguised mule than she would be if it were an actual zebra? And concerning your knowledge that you are not a brain in a vat we need to ask: would you be less likely to believe

¹⁰ If D were universally valid, it would provide a very neat explanation of why knowledge is closed under known entailment. By limiting the scope of D , we lose this explanation, but we are not thereby committed to rejecting closure. One could maintain that whenever I know that p and that p entails q I also know q , even if there are cases in which my knowledge of q can't be based on the evidence provided by p .

¹¹ In addition, the intuitively acceptable cases of inferential knowledge pass the test. In particular, the condition is satisfied by the general background evidence on the basis of which Pritchard, Vogel and I think that Zula can know that the animal is not a cleverly disguised mule (Zalabardo 2012: 142). It follows from this that endorsing E doesn't amount to the demand that knowledge requires discriminating evidence. Zula's background evidence is favouring evidence that satisfies E .

¹² Putnam's original intention was to show, precisely, that envatted brains don't have the same perceptual beliefs as their embodied counterparts (Putnam 1981). This approach takes the discussion of scepticism in an entirely different direction.

¹³ E can also be used to explain why background knowledge cannot be treated as the evidential basis for your knowledge that you are not a brain in a vat. If you were a brain in a vat, you would be just as likely to believe that the brain-in-a-vat scenario is incompatible with the current state of technology, etc.

that you see that you have hands if you were a brain in a vat than you would be if you were normally embodied? If we answer these questions in the affirmative, we'll have to conclude that the ED-based solution to the evidential problem is incompatible with E. Hence if E is correct, as I think it is, the ED-based solution to the evidential problem will have to be rejected.

I think that these questions have to be answered in the affirmative. I take this to be a direct consequence of the introspective indistinguishability of paradigmatic cases perceptual knowledge, in which, according to ED, you have factive reasons, and parallel scenarios in which you lack factive reasons but you continue to blamelessly suppose that you possess them. Pritchard thinks that this point is "an undeniable truth" (20). I agree. If we characterise Zula's case in these terms, we'll have to say that if the animal she is looking at were a cleverly disguised mule, Zula would be just as likely to believe (falsely, in this case) that she sees that the animal is a zebra. Hence our question concerning Zula has to receive an affirmative answer. It follows that the ED-based solution to this instance of the evidential problem is incompatible with E. And if we apply this template to the brain-in-a-vat case, we will have to say that if you were a brain in a vat, you would be just as likely to believe (falsely, in this case) that you see that you have hands. Hence our question, again, has to receive an affirmative answer. It follows that the ED-based solution to this instance of the evidential problem is also incompatible with E.

7 Conclusion

I have argued that Pritchard's ED-based solution to the evidential problem is incompatible with the principle (E) that knowing q on the basis of evidence provided by p requires that you are more likely to believe p if q is true than if q is false. I think that the principle is correct, and that Pritchard's solution will have to be rejected on these grounds, but I haven't provided an adequate defence of the principle, so I cannot claim to have established that Pritchard's solution is unsatisfactory.

Notice, however, that the anti-sceptical power of ED cannot be saved by upholding the universal applicability of D. If D is universally applicable, and ED is correct, then you can know that you are not a brain in a vat on the basis of the evidence provided by your seeing that you have hands. Hence, in the presence of D, ED would sustain a satisfactory solution to the evidential problem. However, if D is universally applicable, we will have a satisfactory solution to the evidential problem even if we reject ED. As we have seen, it follows from D that you can know that you are not a brain in a vat on the basis of the evidence provided by your having hands, independently of how we construe your evidence for your belief that you have hands.

In sum, if, on the one hand, E is true, Pritchard's ED-based solution to the evidential problem is incorrect. If, on the other hand, D is universally valid, then Pritchard's ED-based solution to the evidential problem is redundant. The solution will only be relevant and correct if the validity of D is restricted by a principle other than E—one that rules out your having hands, but not your seeing that you have hands, as a suitable evidential basis for your knowledge that you are not a brain in a vat. A vindication of Pritchard's solution will have to be based on a principle with these features.

In section 4 I expressed my agreement with Pritchard's claim that if ED is not the right account of perceptual knowledge, then the evidential problem for anti-sceptical knowledge cannot be solved, either in terms of perceptual evidence or by reference to background knowledge. Now I'm claiming,

although I haven't quite shown this, that endorsing ED won't give us a solution to the evidential problem either.

It may seem that this outcome would make scepticism unavoidable. If the evidential problem can't be solved, the thought would go, then the first premise of the sceptical argument is correct: you don't know that you are not a brain in a vat, and if, as I've also claimed, the second premise is also unassailable, we seem to be left with no option but to accept the conclusion that you don't know that you have hands.

I want to close by pointing out a gap in this line of reasoning. I am claiming that, with or without ED, we won't be able to find a suitable evidential basis for your knowledge that you are not a brain in a vat. This is a problem, however, only on the assumption that this knowledge needs an evidential basis. If we could provide a non-evidential account of how you know you are not a brain in a vat, we would be able to reject the first premise of the sceptical argument even if the evidential problem had no solution.¹⁴ This, I submit, is the best way forward.

REFERENCES

- Carnap, Rudolf. 1962. *Logical Foundations of Probability*. 2nd ed. Chicago and London: University of Chicago Press.
- Davies, Martin. 1998. "Externalism, Architecturalism, and Epistemic Warrant". In *Knowing Our Own Minds*, edited by C. Wright, B. Smith and C. Macdonald. Oxford: Oxford University Press.
- Dretske, Fred. 1970. "Epistemic Operators". *Journal of Philosophy* 67:1007-23.
- Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge, Massachusetts: Harvard University Press.
- Pritchard, Duncan. 2012. *Epistemological Disjunctivism*. Oxford: Oxford University Press.
- Putnam, Hilary. 1981. *Reason, Truth and History*. Cambridge: Cambridge University Press.
- Vogel, Jonathan. 1990. "Are There Counterexamples to the Closure Principle?". In *Doubting. Contemporary Perspectives on Skepticism*, edited by M. D. Roth and G. Ross. Dordrecht: Kluwer Academic Publishers.
- Wright, Crispin. 2000. "Cogency and Question-Begging: Some Reflections on McKinsey's Paradox and Putnam's Proof". *Philosophical Issues* 10:140-63.
- . 2002. "(Anti-)Sceptics Simple and Subtle: G. E. Moore and John McDowell". *Philosophy and Phenomenological Research* 65:330-48.
- . 2003. "Some Reflections on the Acquisition of Warrant by Inference". In *New Essays on Semantic Externalism and Self-Knowledge*, edited by S. Nuccetelli. Cambridge, Massachusetts: MIT Press.
- Zalabardo, José L. 2012. *Scepticism and Reliable Belief*. Oxford: Oxford University Press.

¹⁴ I have put forward a position along these lines in (Zalabardo 2012: 136-39).