

Underdetermination and Closure: Thoughts on Two Sceptical Arguments

Martin Smith

In this paper, I offer reasons for thinking that two prominent sceptical arguments in the literature – the underdetermination-based sceptical argument and the closure-based sceptical argument – are less philosophically interesting than is commonly supposed. The underdetermination-based argument begs the question against a non-sceptic and can be dismissed with little fanfare. The closure-based argument, though perhaps not question-begging per se, does rest upon contentious assumptions that a non-sceptic is under no pressure to accept.

I INTRODUCTION

Let P be some quotidian proposition that we would ordinarily take ourselves to know – say the proposition that I am currently seated at my desk. Here is an uninteresting sceptical argument to the effect that I don't know P:

- (S1) If I know P then I must have rational support that favours P over any proposition that is inconsistent with P.
- (S2) The only way that one can have rational support that favours a proposition X over any proposition that is inconsistent with X is if one's current subjective experiences entail that X is true.
- (S3) My current subjective experiences do not entail that P is true.
- (S4) Therefore, I don't know P (S1, S2, S3).

Call this the *simple* sceptical argument. What makes the argument uninteresting is, of course, S2 – a premise that obviously stacks the deck in favour of scepticism, and that no non-sceptic should be willing to accept. A couple of points of clarification regarding this premise: Experiences don't, strictly speaking, stand in any entailment relations. When we say that a proposition X is entailed by my subjective experiences what is meant is that X is entailed by a proposition E which provides a comprehensive description of those experiences – the proposition that it currently seems to me that I am seated at a desk, staring at an illuminated laptop screen, listening to the hum of distant traffic and so on... One's subjective experiences are meant to include things like perceptual seemings and apparent memories – and perhaps we could also add (occurrent) beliefs, desires, intentions etc. We could even let one's current subjective experiences include the totality of one's current non-factive mental states, if one finds this a better category to work with.

In any case, it's obvious that my subjective experiences, so characterised, are not going to entail anything about the external world, other minds, the past, the future etc. (modulo certain externalist constraints on content¹). If S2 holds, then the only propositions for which I could have favouring rational support are going to be propositions that directly describe my current subjective experiences. What this argument seeks to do, in effect, is derive the implausible claim that we cannot know any propositions about the external world, other minds, the past or the future from the even

¹ Such constraints may mean that certain very general propositions about the external world or other minds or the past are entailed by propositions about the contents of one's non-factive mental states. If the latter propositions are included as facts about one's 'subjective experiences' then S2 is consistent with the former propositions being known. Whatever the truth about this view, however, it won't offer any meaningful relief from the scepticism engendered by S2, and I won't consider it further here.

more implausible claim that we can't have rational support that favours any propositions about the external world, other minds, the past or the future over their alternatives. It may be that there are further, supplementary arguments that could be given for S2 – but, if so, then *this* is where the interest would lie. If no further argument is offered – if the sceptic just appeals to S2 as a bare premise – then this clearly begs the question against a non-sceptic.

By pinpointing S2, I don't mean to suggest that the other premises in this argument are necessarily beyond challenge. While S3 does seem obvious, given the way 'subjective experience' has been characterised, there may be certain theories of knowledge which would predict that S1 is false². In any case, whatever reasons there may be to reject S1, its role in this argument is certainly *not* amongst them. Given the question-begging nature of S2, the argument puts no pressure on S1.

II THE UNDERDETERMINATION-BASED SCEPTICAL ARGUMENT

Most epistemologists would, I imagine, regard the simple sceptical argument to be a rather poor example of the genre – and would insist that there are plenty of sceptical arguments that are more subtle and interesting than this. Indeed, one common argumentative technique in epistemology is to defend a substantial view (of knowledge, justification, evidence...) on the grounds that it is needed to answer some sceptical argument or other. I will be concerned here with two sorts of sceptical argument that have been particularly prominent in recent epistemology – the *underdetermination-based* sceptical argument and the *closure-based* sceptical argument – both of which have been used to prop up substantial philosophical views.

Underdetermination-based sceptical arguments have a long history, with Descartes' sceptical argument in the *Meditations* often interpreted as an argument of this kind (Yalçin, 1992). The closure-based sceptical argument is also, on occasion, associated with Descartes (see Stroud, 1984, chap. 1) but, in its familiar modern form, stems from the work of Dretske (1970) and Nozick (1981, chap. 3, section II). The contemporary distinction between these two forms of argument is a product of a literature on the logical structure of sceptical arguments from the 1990s and 2000s (Brueckner, 1994, Cohen, 1998, Byrne, 2004, Vogel, 2004, Pritchard, 2005) (the present paper owes much to this literature). The distinction plays a pivotal role in Pritchard's recent *Epistemic Angst* (2015), and the terminology that I will use in formulating these arguments is close to his (though little, I think, hinges on the particular terminology that we choose, and I will note some variants along the way).

Let BIV be the proposition that I am a disembodied brain floating in a vat of nutrients wired up to a supercomputer that monitors electrical output and supplies electrical input, stimulating full and rich perceptual experiences, perhaps like those I am currently undergoing... Let P, once again, be the (incompatible) proposition that I am currently seated at my desk. The underdetermination-based sceptical argument, at a first pass, runs as follows:

- (U1) I don't have rational support that favours P over BIV.
- (U2) If I know P then I must have rational support that favours P over BIV.
- (U3) Therefore, I don't know P (U1, U2).

² There are, for instance, certain 'externalist' views on which knowledge doesn't require any kind of rational basis – views on which, say, a belief that is reliably formed, or a belief that appropriately 'tracks the truth' may qualify as knowledge, even if a rational basis for the belief is lacking (Pritchard, 2015, chapter 1, section 4, chapter 2, section 2). Externalists of this stripe may wish to reject S1.

In his initial discussion of this argument (in chapter 2 of *Epistemic Angst*), Pritchard's focus is largely upon U2. As well as offering a more detailed defence of the premise, he considers ways in which it might be modified in order to resist certain attempts to refute it³. His initial remarks on U1, however, are rather brief, and worth quoting in full⁴:

This claim seems undeniable. For given that the experiences had by the subject in the BIV case are subjectively indistinguishable from everyday experiences, then how is one to come by rational support for an everyday perceptual belief that epistemically favors this belief over an incompatible radical skeptical alternative?

(Pritchard, 2015, p30)

Although Pritchard does reject U1 in the end (a point I will come to), I want to look more closely at this short, but suggestive, rationale for the premise.

What Pritchard seems to be proposing is that because P and BIV are both *consistent* with my subjective experiences, it follows that I can't have rational support that favours the former over the latter. This suggests the following general principle, from which U1 might be derived:

(U0) For any two inconsistent propositions X and Y, if Y is consistent with one's current subjective experiences then one doesn't have rational support that favours X over Y.

Even the name 'underdetermination-based sceptical argument' very much suggests this kind of motivation for U1 – suggests that the reason I lack rational support favouring P over BIV is that the truth of the matter is 'underdetermined' by my evidence (construed subjectively).

But if U1 is really supposed to be supported by U0 then this exposes a serious problem in the argument. If I have rational support that favours X over an inconsistent proposition Y then, according to U0, Y must be inconsistent with my subjective experiences. As a result, if I have rational support that favours X over *every* proposition that is inconsistent with X then, according to U0, every proposition that is inconsistent with X would also have to be inconsistent with my subjective experiences, which is just to say that my subjective experiences would have to entail X. Though it may not look it at first, U0 is actually *logically stronger* than S2 – the question-begging premise from the simple sceptical argument.

This point is made quite clearly by Brueckner (1994, pp834-835) – though he uses a different terminology, speaking of 'sensory evidence' rather than 'subjective experiences' or 'rational support'. As Brueckner points out, if the sceptic is assuming that one's sensory evidence cannot favour one proposition X over an incompatible proposition Y in the event that it is consistent with both, then this is tantamount to assuming that one's sensory evidence can only favour those propositions that it

³ In his discussion of U2, Pritchard is largely concerned to thwart externalists of the kind discussed in n2 who claim that not all knowledge is rationally grounded, and who would seek to dismiss the premise on this basis. As Pritchard points out, we could simply replace the notion of knowledge with the notion of *rationally grounded knowledge* and run the argument as before. Having made this substitution, the externalist would no longer have any complaint against U2, but would still need to find some response to the argument, or face a wide-ranging scepticism about rationally grounded knowledge. To claim that not all knowledge is rationally grounded is one thing – to accept that knowledge is never, or almost never, rationally grounded would be an externalism too extreme for most (Pritchard, 2015, chapter 2, section 2).

⁴ Pritchard notes that this premise would be rejected by an abductivist – someone who holds, roughly, that ordinary quotidian propositions like P provide a *better explanation* of our evidence than their sceptical alternatives like BIV. Pritchard offers a series of criticisms of this view (Pritchard, 2015, chap 1, section 6) but, whatever we make of them, I think his focus on abductivism is somewhat misplaced. I will return to this in n11.

entails. To unearth this assumption in a sceptical argument is, as Brueckner puts it, an ‘embarrassment’ for the sceptic.

Any non-sceptic is committed to rejecting U0. In addition, it’s worth pointing out that U0 is subject to a number of compelling counterexamples. Suppose I throw a fair six-sided die and it lands just out of view. It seems obvious that I have rational support that favours the proposition that the die landed on a number between 1 and 5 over the proposition that the die landed 6, despite the fact that both propositions are *consistent* with my subjective experiences. Suppose someone tells me that it’s currently raining outside, and I have no reason to doubt their word. Once again, it seems obvious that I have rational support that favours the proposition that it’s raining over the proposition that it isn’t, even though these propositions are both consistent with my subjective experiences⁵.

When expanded using U0, the underdetermination-based sceptical argument is every bit as uninteresting as the simple sceptical argument from section I. Perhaps we shouldn’t be too hasty in dismissing the argument, though – after all, there may be a better way to bolster U1. One thing that we might observe is that the BIV hypothesis is not only *consistent* with my current subjective experiences, it actually *entails* those experiences – or could, at any rate, be easily formulated in such a way that it does. That is, we could spell out BIV as the proposition that I am a disembodied brain floating in a vat of nutrients supplied with electrical signals... and it currently seems to me that I am seated at a desk, staring at an illuminated laptop screen, listening to the hum of distant traffic... So, rather than using U0 to support U1, we could appeal to the following principle instead:

(U0*) For any two inconsistent propositions X and Y, if Y entails one’s subjective experiences then one doesn’t have rational support that favours X over Y.

Something like this strategy is suggested by Cohen (1998, pp146-147), and may offer a more charitable interpretation of Pritchard’s remark as well (after all, he does say ‘...the experiences had by the subject in the BIV case are subjectively indistinguishable from everyday experiences’).

Ultimately, though, this second rationale for U1 is no more successful than the first – although U0* is logically weaker than U0, *it still commits us to the question-begging S2*. To see this, let E once again be a proposition describing the totality of my current subjective experiences. If there is a proposition Y that is inconsistent with X and consistent with E, then there is a further proposition – namely $Y \wedge E$ – that is inconsistent with X and which *entails* E. If I have rational support that favours X over every inconsistent proposition then, according to U0*, there are no propositions that are inconsistent with X and which entail E. By the above reasoning, it follows that there are no propositions that are inconsistent with X and consistent with E, which is just to say that E entails X. Substituting U0* for U0 fails, then, to inject any more interest into the underdetermination-based sceptical argument.

It’s also worth noting that the counterexamples I levelled at U0 can easily be adapted as counterexamples to U0*⁶: If I throw a fair, six-sided die that lands just out of view, it seems obvious that I have rational support that favours the proposition that I’m currently having an experience as of just having thrown a die etc. and the die landed on a number between 1 and 5, over the proposition that I’m currently having an experience as of just having thrown a die etc. and the die landed 6. If

⁵ Further counterexamples to U0 are described by Goldman (2007) – who also notes that the principle is taken for granted in the underdetermination-based sceptical argument – and by Markosian (2014, p170).

⁶ This is true also of Goldman and Markosian’s counterexamples to U0 as mentioned in the previous footnote (see also Markosian, 2014, p171).

someone tells me that it's currently raining outside and I have no reason to doubt their word, it seems obvious that I have rational support that favours the proposition that I'm currently having an experience as of being just told that it's raining etc. and it is raining, over the proposition that I'm currently having an experience as of being just told that it's raining etc. and it isn't raining. Though it may have an initial appeal, upon close scrutiny U0* is no more plausible a principle than U0.

We have now looked at two strategies to support U1 – using the principles U0 and U0* – and argued that they both beg the question against a non-sceptic. But what if the sceptic declines to offer *any* further support for U1 – and just puts it forward in the hope that we'll accept it as is? In one way, this looks like desperation on the sceptic's part – and could hardly serve to imbue the argument with greater interest. And yet, when putting forward an argument, a sceptic, like anyone else, has to rest content with *some* premises that are not derived from other things. So why not U1? Why should this need to be deduced from a general principle?

U1 states that our evidence fails to favour either one of two incompatible hypotheses – P and BIV. When we encounter rival hypotheses in day to day life – two people involved in a car accident give competing accounts of what happened, two websites give competing weather forecasts etc. – we can of course judge that our evidence favours neither one, and we can do this without necessarily committing to any general principles about evidential favouring. But the two propositions that feature in U1 are not like this; the BIV hypothesis is *artificially designed*, and it is designed to have one very particular characteristic – namely, to be consistent with all of our subjective experiences. In fact, in one important sense, *all that there is* to the BIV hypothesis is that it is inconsistent with many of our ordinary beliefs while being consistent with our subjective experiences – the rest is just filler that can vary from presentation to presentation⁷. If U1 is simply offered as a bare premise, then there is no reason to accept it unless one is also prepared to accept the general claim captured by U0.

Unless the sceptic is able to devise some better way of supporting U1, we are forced to conclude that the underdetermination-based sceptical argument is, in the end, no more philosophically interesting than the simple sceptical argument. In fact, the two arguments turn out to rest upon the very same question-begging assumption – it is just that the underdetermination-based argument succeeds in burying it deeper.

III AN ASIDE: EPISTEMOLOGICAL DISJUNCTIVISM

As I mentioned in the previous section, the underdetermination-based sceptical argument plays a key role in Pritchard's *Epistemic Angst*, where he uses it to argue his case for *epistemological disjunctivism* – the view that, in normal cases of perceptual knowledge, one's perceptual belief enjoys a kind of rational support that is both *factive* and *reflectively accessible* (Pritchard, 2012, part 1, chap. 1, 2015, chap. 5). According to the disjunctivist, when I know, on the basis of perception, that I'm currently seated at my desk, the rational support for my belief is provided by the fact that I *see* (and feel) that I'm seated at my desk – which entails that I'm seated at my desk and which is (supposedly) accessible via reflection. In this case, I do have rational support that favours P over BIV (indeed rational support that entails P), and U1 is false. If, of course, I was hallucinating or I really was a brain in a vat etc., then

⁷ While any worthwhile sceptical hypothesis will be consistent with one's subjective experiences, it is only 'global' sceptical hypotheses that are inconsistent with a broad range of ordinary beliefs. 'Local' sceptical hypotheses, in contrast, clash with only a few, targeted beliefs and leave many others untouched. While I only consider the global brain-in-a-vat hypothesis here, sceptical arguments can equally well be driven by local sceptical hypotheses (see Smith, 2016).

I wouldn't have this kind of rational support for my perceptual belief, even though my subjective experiences may be exactly the same. But if I'm not in this kind of situation then, according to the disjunctivist, my rational support need not be limited to my subjective experiences, or to the states that I share with my hallucinating or envatted counterpart⁸.

Whatever one makes of epistemological disjunctivism, to adopt this view *as a way of avoiding the underdetermination-based sceptical argument* would, I think, be a considerable overreaction. As suggested in the previous section, the most straightforward way for a non-sceptic to respond to the argument is to point out that one can have rational support that favours one proposition over another even when they are both consistent with one's subjective experiences – and to reject U0 and U1 on this basis. In making this point, we don't need to appeal to disjunctivism – or to any substantial philosophical theory. This is not to say, of course, that epistemological disjunctivism is necessarily mistaken. Pritchard offers some independent considerations in support of the view, which I won't consider here⁹. The point that I *would* make here is that independent considerations are very much needed – the underdetermination-based sceptical argument, in and of itself, generates no support for the view¹⁰.

Furthermore, even if one does accept disjunctivism (for independent reasons), one should still, I believe, respond to the underdetermination-based sceptical argument in the straightforward way suggested above. That is, even for an avowed disjunctivist, to invoke disjunctivism in responding to this argument is excessive. Suppose again that I've just rolled a fair six-sided die and it has landed out of view. Suppose a sceptic tells me that I have no rational support favouring the proposition that the die is showing between 1 and 5 over the proposition that the die is showing 6. Even an avowed disjunctivist, such as Pritchard, would presumably wish to resist this suggestion. But disjunctivism has no bearing on this kind of case – it isn't a case of putative perceptual knowledge, and I clearly don't have any rational support that entails the proposition that the die is showing between 1 and 5. If the disjunctivist accepts that I do nevertheless have rational support that favours the 1-5-proposition over

⁸ Sceptical concerns are also at the forefront of McDowell's case for epistemological disjunctivism (see, for instance, McDowell, 1982, 1995) – though McDowell is less forthcoming than Pritchard in explaining the sceptical problem(s) that disjunctivism is allegedly needed to answer.

⁹ According to Pritchard, disjunctivism is supported by our ordinary epistemic practices, since these practices often involve appealing to factive reasons when our beliefs are challenged. In particular, as Pritchard points out, 'I can see that...', 'I can hear that...' etc. can, under the right circumstances, seem like perfectly sensible answers to 'Why do you believe that...?' 'How do you know that...?' etc. While these considerations are mentioned in *Epistemic Angst* (2015, chap. 5, section 3), most of the dialectical weight still appears to be borne by the underdetermination-based sceptical argument. These considerations are perhaps given a more prominent role in Pritchard's *Epistemological Disjunctivism* (Pritchard, 2012, part I, section 2).

¹⁰ Logue (2011, section 1) suggests a variation on the underdetermination-based sceptical argument in which U1 is, in effect, broken down into two separate premises. Adapting Logue's terminology slightly, in her version of the argument U1 becomes: (U1a) If BIV were true then I would not have rational support that favours P over BIV, and (U1b) The rational support that is available to me is the same whether P or BIV is true. (While a few other minor adjustments are needed to get us to U3, the argument can, in effect, go through with U1a and U1b in place of U1). In a way, this reformulated argument seems tailor-made for a disjunctivist treatment, as U1b would immediately set off disjunctivist alarm bells. And yet, presenting the argument in this way can also highlight why the dispute between disjunctivists and their opponents may be something of a sideshow: Unless we can find a way to motivate U1a that does not presuppose U0 or U0* or something of this ilk, the anti-sceptic should refuse to follow the argument any further, and will never even get so far as to consider the anti-disjunctivist premise U1b.

the 6-proposition, then they already have all that they need to reject U0 and U1 – epistemological disjunctivism is surplus to requirements¹¹.

Pritchard does, in fact, explicitly accept that there is such a thing as non-conclusive favouring – that it is possible to have rational support that favours a proposition X over a proposition Y without entailing X (Pritchard, 2015, chap. 5, section 4, see also, Pritchard, 2012, part II). He argues, however, that non-conclusive favouring is simply not available when it comes to P and BIV, and that the only rational support that could favour P over BIV is rational support that entails P (such as that supplied by disjunctivism). If this argument were effective, then it could be used to bolster U1 in such a way that it really would take a substantial philosophical theory, like disjunctivism, to resist it.

Pritchard's argument begins from the claim that one cannot appeal to any background evidence or beliefs in order to argue against the BIV hypothesis – because the hypothesis calls all of this evidence 'into question' (Pritchard, 2015, pp137-138). But what does it mean exactly for a hypothesis like BIV to call a piece of evidence or a background belief 'into question'? In explaining this, Pritchard falls back on something close to the idea that all of our subjective experiences are consistent with, or entailed by, BIV. He writes, for instance: 'Raising the consideration that, for example, one has good reason for thinking that current technology could not support BIVs is clearly inappropriate in the context of the BIV sceptical hypothesis precisely because one would also believe that one had reasons of this sort if one were a BIV.' (Pritchard, 2015, p138). If this is all that is at work in the argument then, clearly, we are still going to need something like U0 or U0* in order to get to the desired conclusion. I won't discuss this further here – and, in any case, I don't mean to definitively rule out the possibility of there being some kind of non-question-begging argument in support of U1. But if such an argument were offered, then this is the argument that would be doing the philosophical 'heavy lifting' – and not the underdetermination-based sceptical argument.

IV THE CLOSURE-BASED SCEPTICAL ARGUMENT

The closure-based sceptical argument, like the underdetermination-based sceptical argument, trades on the possibility of concocting hypotheses – like BIV – that are consistent with our subjective experiences, but inconsistent with propositions that we would ordinarily take ourselves to know – like P. Sticking with these two propositions, the closure-based argument runs as follows:

- (C1) I cannot know \sim BIV.
- (C2) If I know P, then I can know \sim BIV.
- (C3) Therefore, I don't know P. (C1, C2)

This argument has been widely discussed since the work of Dretske and Nozick and, from the very beginning, the primary focus has been on C2, which is motivated by a *closure principle* for knowledge – a principle to the effect that one can always extend one's knowledge by drawing competent (single

¹¹ This helps to illustrate why Pritchard is wrong to focus on abductivism in his initial defence of U1 (as noted in n4). U1 should be rejected by anyone who thinks it possible for evidence to favour a proposition without entailing it. It doesn't matter whether this favouring relation is spelled out in abductivist terms or in some other way. In fact, the non-sceptic is under no obligation to spell this relation out *at all* in order to respond to the underdetermination-based sceptical argument – its existence is strongly supported by examples, and the sceptic has provided no reason whatsoever to doubt it.

premise) deductive inferences. While there is some controversy over how to precisely formulate this closure principle, many epistemologists have settled on something like the following:

If one knows a proposition X and knows that X entails Y and one competently deduces Y from X, while retaining knowledge of X, then one comes to know Y.

(Hawthorne, 2003, p34, Pritchard, 2015, p13)

Since I know that P entails \sim BIV and am perfectly capable of deducing the latter from the former, if I know P (and can make the deduction without losing this knowledge) then, according to the closure principle, I must be in a position to know \sim BIV, giving us C2.

C1, on the other hand, is not usually discussed at length. Some who present the argument simply regard it as 'intuitive' that I can't know that I'm not a brain in a vat and leave it at that. Those who do say more to motivate the premise generally just labour the familiar point that the BIV proposition is consistent with all of our current subjective experiences. As Pritchard writes:

The initial plank in the case for skepticism comes from the contention that one cannot know that one is not a BIV. Such a claim seems entirely compelling. After all, since the BIV scenario is ex hypothesi subjectively indistinguishable from normal perceptual conditions, it is hard to see how one might come to know such a thing.

(Pritchard, 2015, p12)

In a similar vein, Nozick writes:

You think you are seeing these words but could you not be hallucinating or dreaming or having your brain stimulated to give you the experience of seeing these marks on paper even though no such thing is before you? More extremely, could you not be floating in a tank while super-psychologists stimulate your brain electrochemically to produce exactly the same experiences as you are now having, or even to produce the whole sequence of experiences you have had in your lifetime so far? If one of these other things was happening, your experience would be exactly the same as it now is. So how can you know that none of them is happening?

(Nozick, 1981, p167)

For a final, more recent, example Kelp remarks:

You don't know that you are not a handless BIV. How could you? After all, everything would seem to you exactly as it would were you to be a normal handed person.

(Kelp, 2021, p376)

On one interpretation, Pritchard, Nozick and Kelp are appealing to the following principle:

(C0) For any proposition X, if \sim X is consistent with one's current subjective experiences then one cannot know X.

But it follows immediately from C0 that the only way in which one can know a proposition X is if X is entailed by one's subjective experiences (see Vogel, 2004, pp426-427). That is to say, C0 commits us to an especially demanding kind of *infallibilism* about knowledge, and should be immediately rejected by any non-sceptic¹².

¹² There are several different ways of defining infallibilism about knowledge, not all of which will make scepticism inevitable. On one definition, for instance, infallibilism is the view that one can only know a proposition X if X is entailed by one's *evidence*. If we adopted a generous view of when a proposition could count as part of one's evidence – one that allowed propositions about the external world, other minds, the future, the past etc. to qualify – then we could accept this kind of infallibilism and still be non-sceptics (for discussion see Brown, 2018,

As we've seen, it's obvious that my subjective experiences are not going to entail anything about the external world, or other minds, or the past, or the future etc. If C0 holds, then the only propositions which I could know are going to be those propositions that directly describe my current subjective experiences. C0 allows us, then, to derive a sceptical result directly, without the need for C2 or the closure principle. The idea that the closure principle may turn out to be an idle wheel in motivating scepticism is suggested by Brueckner (1994) and Byrne (2004) and defended in detail by Schönbaumsfeld (2017). Schönbaumsfeld also notes that C1 is usually motivated by citing the 'subjective indistinguishability' of propositions such as P and BIV, and argues that this motivation implies an underlying assumption that would effectively lead to scepticism all by itself¹³.

When C1 is based upon C0, the closure-based sceptical argument is devoid of interest. It would be convenient, in a way, if we could simply leave our analysis here – and put the closure-based sceptical argument into the same category as the underdetermination-based sceptical argument (and simple sceptical argument). But there are some further complications associated with the closure-based argument – complications which do, in the end, give it more philosophical interest.

As observed in section II, the BIV hypothesis could be formulated in such a way that it *entails*, and is not merely consistent with, one's subjective experiences. This suggests another way of motivating C1:

(C0*) For any proposition X, if $\sim X$ entails one's current subjective experiences then one cannot know X.

This represents an alternative – and perhaps even more faithful – interpretation of Pritchard, Nozick and Kelp's remarks. C0* is logically weaker than C0 and does not commit one to infallibilism about knowledge. Is a non-sceptic under pressure to accept it? One thing that we might observe right away is that C0* is somewhat reminiscent of U0* which, as we've seen, a non-sceptic is committed to rejecting. That is, a non-sceptic must affirm that one can have rational support that favours a proposition X over an inconsistent proposition Y, even if one's subjective experiences are entailed by Y. From here, it seems a relatively small step to accept that one could know a proposition X even if one's subjective experiences are entailed by $\sim X$.

As discussed in sections I and II, it's plausible that rational favouring is one necessary condition for knowledge – if one knows a proposition X then one must have rational support that favours X over $\sim X$. But even if $\sim X$ entails one's current subjective experiences, once we have rejected U0*, we should remain open to the idea that this condition could *still be met* – that it could still be the case that one has rational support that favours X over $\sim X$. As a result, unless there is some *other* condition on knowledge which is necessarily precluded in this kind of situation, we should also be open to the idea that one could still know X – and, thus, we should reject C0*.

chap. 1). The kind of infallibilism forced by C0, however, is less forgiving, and leaves no scope for manoeuvres of this kind.

¹³ The underlying assumption that Schönbaumsfeld identifies is what she calls the 'default view' of perceptual reasons – the anti-disjunctivist position that our perceptual reasons are limited to how things *appear* and, thus, are the same in both 'good' cases of veridical perception and 'bad' cases of radical deception. In light of the discussion in the last section, I don't think that this default view is enough, by itself, to motivate C1 (or a more general sceptical conclusion). A non-sceptic can perfectly well accept Schönbaumsfeld's default view (or stay neutral) and still reject C1.

While I think it is a small step for a non-sceptic to reject $C0^*$, given their prior rejection of $U0^*$, I wouldn't go so far as to say that the use of this principle to support $C1$ *begs the question* against a non-sceptic. Furthermore, on this reconstruction of the closure-based sceptical argument, the closure premise $C2$ will turn out to play a crucial role in deriving the sceptical conclusion – and, as a result, a non-sceptic could accept $C0^*$, provided they are willing to reject the closure principle. This particular combination of views is, in fact, characteristic of one version of the *theory of relevant alternatives*.

According to the theory of relevant alternatives, in order to know a proposition X , one must have ruled out all of the *relevant alternatives* to X – that is, all of the relevant propositions that are inconsistent with X . Relevant alternatives theorists disagree over what exactly is involved in ruling an alternative out, but may take the view that an alternative which entails one's subjective experiences can never count as ruled out. Relevant alternatives theorists also disagree over what it is for an alternative to be relevant, but may take the view that $\sim X$ always counts as a relevant alternative to X . These commitments would deliver us $C0^*$ – if $\sim X$ entails one's subjective experiences then one cannot rule it out and, if it is a relevant alternative to X , then one is prevented from knowing X . This view would predict, then, that I cannot know $\sim BIV$, but it need not lead to a general sceptical result. Consider again proposition P (that I'm currently seated at my desk). Since $\sim P$ does not entail my subjective experiences, it is consistent with the present view that I have successfully ruled it out and, provided BIV and any other alternatives that do entail my subjective experiences are deemed irrelevant, it is consistent with my knowing P . It is for this reason, of course, that closure can fail on the theory of relevant alternatives¹⁴.

There are yet further ways that one could try to support $C1$. It seems clear, for instance, that I haven't conducted any kind of inquiry or investigation into the possibility that I might be a brain in a vat – I can't point to a time at which I acquired justification for believing that this possibility doesn't obtain, and I can't offer any reasons that would reassure someone who takes the possibility seriously (Wright, 2004, Smith, 2013). If we were prepared to endorse a principle like the following, then this would give us another motivation for $C1$:

($C0^{**}$) For any proposition X , if one has not acquired justification for believing X and/or one can offer no reasons that would potentially convince a doubter, then one cannot know X .

Once again, is a non-sceptic under pressure to accept such a principle? For any proposition X , if $\sim X$ entails all of one's subjective experiences then it seems plausible that one could not legitimately claim to have acquired justification for believing X or to have convincing reasons that one could offer a person who doubts X . As a result, if we deny $U0^*$, then lacking acquired justification for believing X and being unable to offer convincing reasons in favour of X is still compatible with one having rational support that favours X over $\sim X$. But if we have rational support that favours X over $\sim X$, why should it matter whether we acquired our justification, or are able to convince a doubter? As above, unless there is some other condition on knowledge, aside from rational favouring, which is supposed to be

¹⁴ Some relevant alternatives theorists have rejected the assumption that $\sim X$ must always count as a relevant alternative to X (Stine, 1976, Blome-Tillmann, 2015). This style of relevant alternatives theory can't be used to underwrite $C0^*$ and may, on the contrary, provide a rationale for rejecting it – and a way of responding to the closure-based sceptical argument that doesn't involve rejecting closure. In any case, as I suggest in the main text, a non-sceptic has no obligation to offer a 'rationale' for rejecting $C0^*$ – drawing the connection with the suspect $U0^*$ is enough to shift the dialectical onus onto the sceptic.

blocked in this circumstance, we should be open to the idea that one could still know X – and, thus, we should reject CO**¹⁵.

While CO* and CO** may not exactly beg the question against a non-sceptic, they should be regarded as contentious assumptions and, absent any other reasons for accepting C1, the premise should be regarded as weak. In spite of this, the argument does carry philosophical lessons. If one is attracted to principles like CO* and CO** (and wishes to avoid scepticism) then the price of this is to deny C2 and the principle of closure. As discussed, there are theories of knowledge – like the theory of relevant alternatives – which could offer this package of views. If, on the other hand, one wishes to preserve closure (and avoid scepticism) then one will have to give up CO* and CO**. This too can serve as an impetus for developing a compliant theory of knowledge¹⁶.

V CONCLUSION

The best sceptical arguments, one might think, expose a deep or hidden tension within our ordinary thinking or orthodox philosophical theorising about knowledge (Wright, 1985, pp429-430). Whether there are any sceptical arguments that accomplish this is not something that I take myself to have settled here – but the two arguments that I have surveyed do not I think live up to this kind of description. In the end, the primary tension at work in these arguments is the tension between the idea that we know a great deal, and the idea that the standards for knowledge or for rational support are extremely demanding. While that tension is, to some extent, present in our ordinary thinking and philosophical theorising, it is neither hidden nor deep.

ACKNOWLEDGEMENTS

Thanks to Dominik Kauss and to an anonymous referee for helpful comments on earlier drafts of this material.

¹⁵ Instances of CO** that involve the negations of *local* sceptical hypotheses, of the kind discussed in n7, may be more intuitively compelling than instances that involve the negations of global sceptical hypotheses like BIV. While global hypotheses would typically lie beyond any kind of empirical investigation, it may be straightforward to describe a procedure for investigating a local hypothesis – and for potentially acquiring justification or convincing reasons for rejecting it. If such a procedure exists, it might be particularly awkward to insist that one can know the hypothesis to be false without having carried the procedure out. I won't discuss this further here – but see Smith (2016).

¹⁶ Another kind of response to the closure-based sceptical argument, favoured by some contextualists and subject sensitive invariantists, involves conceding that 'I don't know that I am not a brain in a vat' (usually) expresses a truth when uttered, while maintaining that 'I know that I am currently seated at my desk' can also express a truth when sceptical considerations are far from one's mind. One advertised advantage of this approach is that it can explain why premise C1 seems intuitive and the argument seems compelling, while still avoiding scepticism (see, for instance DeRose, 1995). I think it is a mistake, however, to regard the 'intuitiveness' of C1 as some sort of datum that an adequate treatment of scepticism needs to account for. For the kinds of reasons given at the end of section II, I suspect that one will only find C1 intuitive to the extent that one is willing to accept certain general principles – principles that are very much in need of philosophical scrutiny. To give these issues the attention they deserve would, however, take me beyond the scope of this paper. None of this is to say, in any case, that either contextualism or subject sensitive invariantism is necessarily incorrect – merely that they are not needed for a satisfying response to the closure-based sceptical argument.

BIBLIOGRAPHY

- Blome-Tillmann, M. (2015) 'Solving the Moorean puzzle' *Philosophical Studies* v172(2), pp493-514
- Brown, J. (2018) *Fallibilism: Evidence and Knowledge* (Oxford: Oxford University Press)
- Brueckner, A. (1994) 'The structure of the skeptical argument' *Philosophy and Phenomenological Research* v54(4), pp827-835
- Byrne, A. (2004) 'How hard are the sceptical paradoxes?' *Noûs* v38(2), pp299-325
- Cohen, S. (1998) 'Two kinds of skeptical argument' *Philosophy and Phenomenological Research* v58(1), pp143-159
- DeRose, K. (1995) 'Solving the skeptical problem' *Philosophical Review* v104(1), pp1-52
- Dretske, F. (1970) 'Epistemic operators' *Journal of Philosophy* v67(24), pp1007-1023
- Goldman, A. (2007) 'The underdetermination argument for brain-in-the-vat scepticism' *Analysis* v67(1), pp32-36
- Hawthorne, J. (2003) *Knowledge and Lotteries* (Oxford: Oxford University Press)
- Kelp, C. (2021) 'Theory of inquiry' *Philosophy and Phenomenological Research* v103(2), pp359-384
- Logue, H. (2011) 'The skeptic and the naïve realist' *Philosophical Issues* v21(1), pp268-288
- Markosian, N. (2014) 'Do you know that you are not a brain in a vat?' *Logos and Episteme* v5(2), pp161-181
- McDowell, J. (1982) 'Criteria, defeasibility and knowledge' *Proceedings of the British Academy* v68, pp455-479
- McDowell, J. (1995) 'Knowledge and the internal' *Philosophy and Phenomenological Research* v55(4), pp877-893
- Nozick, R. (1981) *Philosophical Explanations* (Cambridge, MA: Harvard University Press)
- Pritchard, D. (2005) 'The structure of sceptical arguments' *Philosophical Quarterly* v55(218), pp37-52
- Pritchard, D. (2012) *Epistemological Disjunctivism* (Oxford: Oxford University Press)
- Pritchard, D. (2015) *Epistemic Angst* (Princeton, NJ: Princeton University Press)
- Schönbaumsfeld, G. (2017) 'The 'default view' of perceptual reasons and 'closure-based' sceptical arguments' *International Journal for the Study of Skepticism* v7(2), pp114-135
- Smith, M. (2013) 'Entitlement and evidence' *Australasian Journal of Philosophy* v91(4), pp735-753
- Smith, M. (2016) 'Scepticism by a thousand cuts' *International Journal for the Study of Skepticism* v6(1), pp44-52
- Stine, G. (1976) 'Skepticism, relevant alternatives and deductive closure' *Philosophical Studies* v29(4), pp249-261
- Stroud, B. (1984) *The Significance of Philosophical Scepticism* (Oxford: Oxford University Press)
- Vogel, J. (2004) 'Skeptical arguments' *Philosophical Issues* v14, pp426-455

Wright, C. (1984) 'Facts and certainty' *Proceedings of the British Academy* v71, pp429-472

Wright, C. (2004) 'Warrant for nothing (and foundations for free)?' *Aristotelian Society Supplementary Volume* v78(1), pp167-212

Yalçın, U. (1992) 'Skeptical arguments from underdetermination' *Philosophical Studies* v68(1), pp1-34