VOL. 55 NO. 218 JANUARY 2005

The Philosophical Quarterly

CONTENTS

BLACKWELL PUBLISHING
FOR
THE SCOTS PHILOSOPHICAL CLUB
AND THE
UNIVERSITY OF ST ANDREWS

The Philosophical Quarterly, Vol. 55, No. 218 *January 2005*
ISSN 0031–8094

THE STRUCTURE OF SCEPTICAL ARGUMENTS

By Duncan Pritchard

It is nowadays taken for granted that the core radical sceptical arguments all pivot upon the principle that the epistemic operator in question is 'closed' under known entailments. Accordingly, the standard anti-sceptical project now involves either denying closure or retaining closure by amending how one understands other elements of the sceptical argument. However, there are epistemic principles available to the sceptic which are logically weaker than closure but achieve the same result. Accordingly the contemporary debate fails to engage with the sceptical problem in its strongest form.

I. SCEPTICISM AND CLOSURE

There has been a dramatic re-emergence of interest in radical sceptical arguments in the recent literature, and this has inevitably prompted the question of what the underlying structure of the standard radical sceptical argument is. A 'radical sceptical argument' is any argument which purports to call into question one's knowledge of a significant class of propositions which one would ordinarily take oneself to know ('everyday propositions'). The variety of radical scepticism on which much of the contemporary discussion focuses typically proceeds by highlighting a sceptical hypothesis which, it is claimed, one cannot know is false, and then argues that this lack of knowledge entails that one lacks knowledge of any one of a wide class of everyday propositions which are inconsistent with the sceptical hypothesis. This kind of sceptical argument, known as an argument from ignorance, has the following form, where '*sh*' refers to a sceptical hypothesis, and '*e*' refers to one of these everyday propositions which are inconsistent with *sh*:

S1. If *S* knows *e*, then *S* knows $\neg sh$
S2. *S* does not know $\neg sh$
S3. Therefore *S* does not know *e*.

And since this argument can be repeated with any one of a wide class of everyday propositions in the place of *e* (one would just have to vary the

sceptical hypothesis to suit), and for any subject in the place of *S*, it follows that radical scepticism ensues.

The most interesting element of the sceptical argument for my purposes is what supports (S1). The standard suggestion is that what underlies this component of sceptical reasoning is an implicit (or sometimes explicit) appeal to the highly plausible 'closure' principle that knowledge is 'closed' under known entailments. This can be roughly expressed as follows:

CK. For all *S*, *p*, *q*, if *S* knows *p*, and *S* knows that *p* entails *q*, then *S* knows *q*.

(CK) does seem, at first pass, to be entirely uncontentious. For example, one could imagine a detective reasoning as follows: I know that the butler was in the conservatory at the time of the murder, and I know that if the butler was in the conservatory at the time of the murder then he was not at the scene of the crime (the kitchen, say); hence I know that the butler was not at the scene of the crime.

The manner in which one could construe (S1) in the sceptical argument as resting upon this principle is evident. For example, suppose the sceptical hypothesis in question is that one is a brain in a vat being 'fed' experiences by a neuroscientist (so that 'BIV' stands for '*S* is a brain in a vat'). Given that *e* is taken to be a proposition that one knows is inconsistent with (BIV) (such as that one is currently seated), we can formulate the following instance of the (CK)-based sceptical argument, where (CK) is explicitly regarded as supporting (S1′):

S1′. If *S* knows *e*, then *S* knows ¬(BIV)
S2′. *S* does not know ¬(BIV)
S3′. *S* does not know *e*.

With (CK) in play, and since it can be taken as given that one knows that being seated entails that one is not a brain in a vat (because a brain in a vat does not *sit* anywhere), it follows that if one knows that one is currently seated, then one also knows that one is not a brain in a vat, just as (S1′) says. In effect, what (CK) achieves in this argument is to forge the necessary connection between knowledge of everyday propositions and knowledge of the denials of sceptical hypotheses (and therefore between lack of knowledge of the denials of sceptical hypotheses and lack of knowledge of everyday non-sceptical propositions). Once this connection is forged, justifying the sceptical conclusion becomes relatively straightforward.

It thus appears that the standard sceptical argument is itself dependent upon the correctness of (CK), the closure principle for knowledge. This, at least, is how it has seemed to a number of prominent commentators. In particular, there has been an influential group of epistemologists, led by

Fred Dretske and Robert Nozick, who have argued that denial of the principle would be sufficient to block radical sceptical arguments of this form.[1] Indeed, Nozick (p. 242) goes so far as to endorse the following conditional: '... if our notion of knowledge was as strong as we naturally tend to think (namely, closed under known entailment) then the sceptic would be right'.

Others, most notably Keith DeRose and David Lewis, have argued that one can only evade radical scepticism whilst retaining (CK) by 'contextualizing' one's understanding of the knowledge operator.[2] Again, then, we find the view that denying (CK) would suffice to meet the radical sceptical argument, albeit as part of a position which advocates the retention of this principle. It is this claim, that the status of the radical sceptical argument is inextricably intertwined with the status of (CK) in this way, which I shall consider in more detail in this paper.

II. SCEPTICISM AND UNDERDETERMINATION

The most prominent attack on this picture of the role of (CK) in radical sceptical arguments has been due to Anthony Brueckner.[3] In effect, Brueckner's argument is that the sceptic's use of (CK) relies in turn upon a certain epistemic underdetermination principle, and thus that one could dispense with (CK) in setting up a sceptical argument, and simply employ the epistemic underdetermination principle instead. The underdetermination principle which Brueckner proposes, a version of which can plausibly be found in ancient Pyrrhonian sceptical writings, can be expressed as follows:

UP. For all *S*, *p*, *q*, if *S*'s evidence for believing *p* does not favour *p* over some hypothesis *q* which *S* knows to be incompatible with *p*, then *S*'s evidence does not justify *S* in believing *p*.[4]

This principle has a great deal of appeal, because it is far from clear how one's evidence could possibly perform a supporting role, let alone a

[1] See F. Dretske, 'Epistemic Operators', *Journal of Philosophy*, 67 (1970), pp. 1007–23; R. Nozick, *Philosophical Explanations* (Oxford UP, 1981).

[2] K. DeRose, 'Solving the Skeptical Problem', *Philosophical Review*, 104 (1995), pp. 1–52; D. Lewis, 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74 (1996), pp. 549–67.

[3] A. Brueckner, 'The Structure of the Skeptical Argument', *Philosophy and Phenomenological Research*, 54 (1994), pp. 827–35; but see also Ü. Yalçin, 'Sceptical Arguments from Underdetermination', *Philosophical Studies*, 68 (1992), pp. 1–34; J. Vogel, 'Varieties of Skepticism', *Philosophy and Phenomenological Research*, 68 (2004), pp. 1–37.

[4] Brueckner's own formulation of this principle is slightly but unimportantly different from (UP). For a very different formulation of the underdetermination principle, see Vogel, §1.

supporting role that would suffice to support a justification, if that evidence does not prefer the target proposition over alternatives that are known to be incompatible. If, for example, my evidence for believing that I am currently in the town's Odeon cinema does not favour this proposition over the competing scenario (which I know is incompatible) that I am in the town's other cinema, the Multiplex, perhaps because I have just woken up from a drunken sleep with only a dazed recollection of the chain of events that led to my being in a cinema, then it is hard to see how that evidence could possibly support a justification for my belief that I am at present in the Odeon cinema.

Crucially, however, one of the key claims made by the sceptic is that the evidence we have for our everyday beliefs does not favour those beliefs over incompatible sceptical hypotheses. For example, the sceptic will claim that the (largely perceptual) evidence I possess for believing that I am presently sitting at my computer does not favour this hypothesis over the incompatible sceptical hypothesis that I am a brain in a vat being 'fed' the phenomenology of sitting at my desk, when in fact I am not *sitting* anywhere. Let 'sh_k' indicate that *sh* is known (to *S*) to be incompatible with *e*. Accordingly, using (UP), the sceptic can reformulate the argument as follows:

B1. *S*'s evidence for believing *e* does not favour *e* over sh_k
B2. If *S*'s evidence does not favour believing *e* over sh_k, then *S*'s evidence does not justify *S* in believing *e*
B3. Therefore *S*'s evidence does not justify *S* in believing *e* (and therefore *S* is not justified in believing *e*, and therefore *S* does not know *e*).

There are two obvious assumptions in play in the parenthetical part of the conclusion of this argument. The first is that justification is essentially evidential, so that a lack of appropriate evidence entails a lack of justification. The second is that justification is a necessary component of knowledge. Whilst both of these assumptions are *prima facie* plausible, they have each been denied at some point by sections of the epistemological community. For now, however, I shall let them stand and return to consider their status later on. With these assumptions in play, and given that, like the (CK)-based sceptical argument, this argument can be repeated with just about any everyday proposition in the place of *e* (one would just have to vary the sceptical hypothesis to suit), and any subject in the place of *S*, radical scepticism results. Seemingly, then, I have formulated a sceptical argument which is based upon a principle different from (CK), but issues in the same sceptical result.

III. UNDERDETERMINATION AND CLOSURE

The obvious question to ask at this juncture is: what exactly is the relationship between the (CK)-based and the (UP)-based sceptical arguments, and thus between the principles that underlie them? Brueckner's own view on this topic is that (UP) is, at least in effect, equivalent to a version of closure formulated in terms of justification.[5] That is, (UP) is equivalent to the principle that justification is closed under known entailments:[6]

CJ. For all *S*, *p*, *q*, if *S* is justified in believing *p*, and *S* knows that *p* entails *q*, then *S* is justified in believing *q*.

(CJ) seems just as plausible a principle as (CK). If one has evidence which is sufficient to support a justification for belief in a proposition, and one knows that this proposition entails a second proposition, then it is hard to see how one could fail to have evidence thereby that would suffice to support a justification for one's belief in the entailed proposition. If we grant, as I have so far, that justification is necessary for knowledge, then we would expect (CJ) and (CK) to stand or fall together, in that the considerations which lead to the downfall of the one would similarly affect the other.[7] In order to assess the relationship between (CJ) and (UP), I shall simplify matters by considering a concrete example, where one believes the everyday proposition *e* that one is presently seated, which one knows is inconsistent with the (BIV) sceptical hypothesis. Assuming that one knows the relevant entailment, a simplified instance of the closure principle for justification can be given as follows:

A. If *S* is justified in believing *e*, then *S* is justified in believing ¬(BIV).

Similarly, given that one knows the relevant entailment, and that justification is essentially evidential, the underdetermination principle can be simplified as follows:

> If *S*'s evidence does not favour *e* over (BIV), then *S* is not justified in believing *e*.

[5] In his text Brueckner (p. 832) confines himself to saying that the closure principle for justification and (UP) are 'virtually equivalent', but in a footnote on that page he outlines his grounds for thinking that they are actually equivalent.

[6] Brueckner (pp. 832–3) actually makes the stronger claim that (UP) is equivalent to a version of closure for justification where the entailment is *not* known, but since the examples I shall be dealing with are such that one might reasonably suppose that the entailment is always known, I can utilize without loss this modified version of the closure principle for justification.

[7] As Vogel (pp. 8–9) points out, we need to tread carefully here.

Contraposed, this becomes

B. If *S* is justified in believing *e*, then *S*'s evidence favours *e* over (BIV).

With conditionals (A) and (B) as simplified representatives of (CJ) and (UP), the logical relationship between the two principles can be assessed.

First, the putative entailment from (CJ) to (UP). I shall begin by assuming the common antecedent of conditionals (A) and (B):

1. *S* is justified in believing *e*.

Using conditional (A), it follows that

2. *S* is justified in believing ¬(BIV).

All that is needed now is to assume the uncontentious principle that the possession of a justification for belief in a proposition entails that you *lack* a justification for belief in the *negation* of that proposition (I shall consider this principle in more detail below). If that principle is granted, then one's possession of a justification for believing that one is *not* a brain in a vat entails that one *lacks* a justification for believing that one *is* a brain in a vat:

3. *S* is not justified in believing (BIV).

I am assuming, for the time being, that justification is essentially evidential. Accordingly, that *S* is justified in believing *e*, (1), but not in believing (BIV) which *S* knows is incompatible with *e*, (3), entails that *S*'s evidence supports a justification for *S*'s belief in *e* but not for his belief in (BIV). That is, it supports the conclusion that *S*'s evidence favours the hypothesis *e*, that he is sitting here now, over the hypothesis (BIV), that he is a brain in a vat. It therefore follows that

4. *S*'s evidence favours *e* over (BIV).

From the assumption that *S* is justified in believing *e* we can thus infer, using conditional (A) and a further uncontentious principle regarding the notion of justification, that *S* has evidence which favours *e* over (BIV). We are thus entitled to conclude

5. If *S* is justified in believing *e*, then *S*'s evidence favours *e* over (BIV).

Crucially, of course, (5) is just the simplified version of (UP) which I formulated above as conditional (B). (CJ) thus entails (UP).

Now to consider the other direction of entailment, from (UP) to (CJ). This time I shall not only make use of the conditional (B), but also (for reasons which will become apparent in a moment) the following conditional (B*), a variation on (B):

B*. If *S* is justified in believing (BIV), then *S*'s evidence favours (BIV) over *e*.

As before, I begin by assuming the common antecedent of both (A) and (B):

1. *S* is justified in believing *e*.

Using conditional (B), it follows that

2. *S*'s evidence favours *e* over (BIV).

That one's evidence favours belief in *e* over (BIV) must entail that one's evidence does *not* favour belief in (BIV) over *e*:

3. *S*'s evidence does not favour (BIV) over *e*.

Using conditional (B*), however, it follows only from this that one is not justified in believing (BIV):

4. *S* is not justified in believing (BIV).

This is a much weaker conclusion than was aimed at, however, in that *lack* of a justification for believing that one is a brain in a vat is much weaker than *possession* of a justification for believing the *negation* of this proposition, that one is *not* a brain in a vat; and it is the latter conclusion which is licensed by (CJ). Given the uncontentious principle employed above, that if one is justified in believing a proposition, then one is not justified in believing the negation of that proposition, then that one is justified in believing ¬(BIV), as (CJ) licenses, entails that one is not justified in believing (BIV), which is what (UP) licenses. Crucially, however, it does not follow from the fact that one is not justified in believing (BIV) that one is justified in believing ¬(BIV). Indeed, it is entirely possible that one lacks both a justification for believing (BIV) *and* a justification for believing ¬(BIV). Accordingly, (CJ) is a logically stronger principle than (UP), and thus, *contra* Brueckner, (UP) and (CJ) are not logically equivalent. In short, all (UP) gets you is relative evidential supremacy, when what (CJ) demands is a supremacy of justifications; and pending further arguments, the latter claim is much stronger than the former.[8] We are thus faced with two logically distinct epistemic principles, and hence with two different sceptical arguments which employ these principles. That is, the following argument, expressed in terms of (CJ), poses a sceptical challenge different from the (UP)-based sceptical argument (B1)–(B3) which I considered in §II:

[8] S. Cohen, 'Two Kinds of Sceptical Argument', *Philosophy and Phenomenological Research*, 58 (1998), pp. 143–59, argues for a similar conclusion, though as I shall show in a moment, the line of criticism I advance here against Brueckner is very different from that put forward by Cohen. See also Vogel, §2.

S1´´. If *S* is justified in believing *e*, then *S* is justified in believing ¬*sh*
S2´´. *S* is not justified in believing ¬*sh*
S3´´. Therefore *S* is not justified in believing *e* (and therefore *S* does not know *e*).

Moreover, since (CJ) entails, but is not entailed by, (UP), it follows that it is the sceptical argument expressed in terms of (CJ) that is the more demanding. *Prima facie* at least, this has the interesting consequence that even if one were able to impugn (CJ) successfully, perhaps via an argument against (CK), one would still be faced with a sceptical argument characterized in terms of (UP) which would need a response.

IV. INTERLUDE: COHEN ON UNDERDETERMINATION AND CLOSURE

Of course, that one epistemic principle is logically stronger than a second does not in itself mean that any argument formulated in terms of the weaker epistemic principle will survive the refutation of the corresponding argument formulated in terms of the stronger. There could, after all, be other premises in the argument which tip the logical balance back the other way. This, at any rate, is Stewart Cohen's contention in his analysis of Brueckner's argument. For whilst he admits that (UP) is logically weaker than the corresponding closure principle (CJ), he argues that any refutation of the (CJ)-based sceptical argument would entail a refutation of the (UP)-based sceptical argument, but not *vice versa*. If this were so, then the *status quo* would be preserved and the focus of the current sceptical debate could legitimately return to the relevant closure principle.

Cohen begins by taking the first two premises of the (UP)-based sceptical argument, (B1) and (B2), as premises:

C1. *S*'s evidence for believing *e* does not favour *e* over sh_k
C2. If *S*'s evidence does not favour believing *e* over sh_k then *S*'s evidence does not justify *S* in believing *e*.

He then assumes, for *reductio*, what he takes to be the negation of the major premise of the (CJ)-based sceptical argument (S1´´), though with the relationship between justification and evidence made explicit:

C3. *S*'s evidence justifies *S* in believing *e*, but *S*'s evidence does not justify *S* in believing ¬*sh*.

Given the way in which he has understood the denial of the major premise

of the (CJ)-based argument, Cohen is in a position to detach the first conjunct so as to get

C4. *S*'s evidence justifies *S* in believing *e*.

Crucially, however, he can complete the argument by inferring from (C1) and (C2) that

C5. *S*'s evidence does not justify *S* in believing *e*.

What Cohen is trying to do here is to show that if one grants the two premises of the (UP)-based sceptical argument along with the negation of the major premise of the (CJ)-based sceptical argument, then one is directly led to a contradiction. He therefore concludes that this argument shows that the premises of the (UP)-based sceptical argument entail the major premise of the (CJ)-based argument, *viz* the relevant instance of closure for justification that if one's evidence justifies one's belief in *e*, then one's evidence also justifies one's belief in ¬*sh*. Accordingly, if one is able to reject (CJ), and thus to reject the major premise of the (CJ)-based sceptical argument, then one can thereby evade the sceptical argument expressed in terms of (UP) and effective logical parity between the two principles is restored.

Despite the superficial plausibility of this line of reasoning, however, it is fatally flawed. To begin with, (C1), (C2) and (C5) of this argument just *are* the (UP)-based sceptical argument. Accordingly, it should come as no surprise to find that assuming anything that entails that you have a justification for your beliefs should lead to a contradiction, since this is bound to conflict with the radical sceptical conclusion. Indeed, the assumption of *any* conjunction at (C3), where one of the conjuncts implies a justification for belief in *e*, would entail a contradiction, no matter how otherwise outlandish this assumption might be. What is suspicious about Cohen's argument is thus that it only establishes its conclusion (that the premises of the (UP)-based sceptical argument entail the major premise of the (CJ)-based sceptical argument) on the assumption that anyone who denies (CJ) must affirm an instance of justified belief, and this seems highly implausible.

So where does the argument go wrong? The mistake is to allow that the denial of the major premise of the (CJ)-based argument (that a justification for belief in *e* entails a justification for belief in ¬*sh*) is equivalent to the conjunction of a justification for belief in *e* and the absence of a justification for belief in ¬*sh*. It is this move that commits anyone who denies (CJ) to arguing that a justification *does* exist for belief in *e*, so contradicting the sceptical conclusion of the (UP)-based sceptical argument. What Cohen should have contended instead is that the denial of this premise is equivalent to the *possibility* that a justification for belief in *e* could co-exist with the lack

of a justification for belief in $\neg sh$. After all, it could be that a justification for belief in *e* might not entail a lack of justification for belief in $\neg sh$ (and thus that (CJ) might fail) even though there is *never* a case in which a justification for belief in *e* is possessed because of the success of the (UP)-based sceptical argument. In such a situation (CJ) would fail, and yet the (UP)-based sceptical argument would go through without any inconsistency.

Indeed, if Cohen's argument were right, anyone who argued for scepticism on the basis of the (UP)-based argument, and yet nevertheless thought that the relevant (CJ) inference was invalid, would be committed to the implausible position of maintaining both that scepticism is true *and* that the denial of this (CJ)-inference entailed a justification for the target proposition. Clearly this is not the case. Rather, such a person would only be claiming that scepticism is true and that the relevant (CJ)-inference would not have gone through had there been a justification for the antecedent proposition (which there is not, because scepticism is true). It is only by failing to pay careful attention to this point that Cohen gets the initially surprising result that he does.

V. THE SOURCE OF SCEPTICISM

I can thus return to the contention made above, that (UP) and (CJ) can be used to construct different sceptical arguments in such a way that a response to the (CJ)-based argument is not thereby a response to the (UP)-based argument. I have argued, *contra* Brueckner, that these two epistemic principles are logically distinct, and I have also maintained, *contra* Cohen, that the (CJ)-based argument is indeed logically stronger than the (UP)-based argument. There is thus a new sceptical argument to consider, one not obviously affected by arguments against closure.

Interestingly, however, this particular issue about the logical structure of the sceptical argument still has a final surprise in store. Earlier I showed that (CJ) entailed (UP), given only the following principle, a principle I took to be an uncontentious conceptual truth regarding the notion of justification:

J. For all *S*, *p*, if *S* is justified in believing *p*, then *S* is not justified in believing $\neg p$.

So formulated, the principle seems entirely unremarkable. What is interesting about it, however, is that it appears to license a much weakened version of the closure principle for justification, which one might refer to as itself a closure-type principle, and which can be formulated as follows:

WCJ. For all *S*, *p*, *q*, if *S* is justified in believing *p*, and *S* knows that *p* entails *q*, then *S* is not justified in believing ¬*q*.

This is a weaker principle than closure for justification, in that it only demands that a justification for belief in a proposition should entail a *lack* of justification for belief in the *denials* of known entailments of that proposition (rather than entailing a justification for belief in the known entailments). The relationship between (J) and (WCJ) is that they are both based on essentially the same idea, that one cannot have justifications for believing inconsistent propositions, at least where the inconsistency is known. (Indeed, one could construe (J) as being a simplified instance of (WCJ) where the relevant entailment, from *p* to *p*, is so obvious that it is trivially known and so not worth stating.) Accordingly, if one is persuaded on this basis that (J) is true, then one should equally be persuaded of the truth of (WCJ).

One can, however, formulate a sceptical argument in terms of (WCJ) alone which has a conclusion that is just as potentially devastating as the (UP)-based sceptical argument. In order to simplify matters, I shall consider an instance of (WCJ) where *p* is *e*, as defined above, and *q* is the (BIV) hypothesis. Moreover, I shall take it as given that the agent in question knows the relevant entailment. We thus get

C. If *S* is justified in believing *e*, then *S* is not justified in believing (BIV).

I shall begin by assuming, for *reductio*, that *S* is justified in believing *e*:

1. *S* is justified in believing *e*.

Using conditional (C), it follows that

2. *S* is not justified in believing (BIV).

Since I have granted, for the sake of argument, the assumption that justification is essentially evidential in character, it follows from the fact that one has a justification for *e* but not for (BIV) (and that one knows that these two propositions are incompatible) that one's evidence must favour *e* over (BIV). Thus it follows that

3. *S*'s evidence favours *e* over (BIV).

Crucially, however, the sceptical contention is that it is impossible to have evidence for one's beliefs in everyday propositions which prefers them over sceptical alternatives. The sceptic will thus add

4. *S*'s evidence does not favour *e* over (BIV).

And since there is a contradiction here, one can therefore infer the denial of the assumption and conclude

5. Therefore *S* is not justified in believing *e* (therefore *S* does not know *e*).

Clearly, such an argument can be repeated with just about any everyday proposition (one would just have to vary the sceptical hypothesis to suit), and any subject, and thus radical scepticism quickly ensues.

Furthermore, the way in which this sceptical argument proceeds indicates that (WCJ) entails (UP). After all, instead of moving to line (4) above one could instead have inferred

4*. If *S* is justified in believing *e*, then *S*'s evidence favours *e* over (BIV).

That is, it follows from the assumption (1), that one has a justification for believing *e*, that (3) one's evidence favours *e* over (BIV). But of course (4*) is just the simplified version of (UP) which I expressed above as conditional (B). Accordingly, (WCJ) entails (UP). As a result, it should come as no surprise that (CJ) entails (UP) when the principle (J) appealed to in deriving that entailment appears to license a related principle (WCJ) which entails (UP) on its own.

Indeed, (WCJ) and (UP) are, plausibly, logically equivalent, since one can also derive (WCJ) from (UP). I shall begin by assuming

1. *S* is justified in believing *e*.

Employing conditional (B), it follows that

2. *S*'s evidence favours *e* over (BIV).

But if one's evidence favours *e* over (BIV), then it must also be true that one's evidence does not favour (BIV) over *e*:

3. *S*'s evidence does not favour (BIV) over *e*.

Given conditional (B*), however, this means that

4. *S* is not justified in believing (BIV).

So from the assumption that *S* is justified in believing *e*, it follows that *S* is not justified in believing (BIV):

5. If *S* is justified in believing *e*, then *S* is not justified in believing (BIV).

And (5) is, of course, just the simplified instance of (WCJ), conditional (C). Accordingly, (UP) entails (WCJ).

If these arguments are allowed to stand, and there seems no good reason available for thinking that they should not, then it appears that one could

just as well run the sceptical argument in terms of (WCJ). The import of this result is that (WCJ) is an epistemic principle more fundamental to the sceptical argument than (UP) is, since it is more obviously a conceptual consequence of the notion of justification in play here, being, as it is, a principle which is closely related to the highly plausible principle (J). Moreover, the advantage of (WCJ) is that it brings out the relevant sense in which there is a closure-type principle at work in the sceptical argument, albeit one which is weaker than either (CK) or (CJ) (though no less effective).

VI. RESPONDING TO SCEPTICISM

The project that Brueckner instigated of examining more closely the logical structure of the sceptical argument thus reveals two epistemologically significant results. The first is that the sceptical argument can plausibly be understood as driven by an epistemic underdetermination principle (UP) which is logically distinct from the corresponding closure principle (CJ), so that any response to scepticism which merely denied (CJ) would be ineffective. The second is that this logically distinct epistemic principle is, seemingly, itself logically equivalent to a third principle (WCJ), which is conceptually more fundamental.

Before closing, I shall consider briefly how these conclusions affect the sceptical debate. That is, given these two conclusions, how should we go about responding to the sceptical problem? Perhaps the most direct response would be to deny that the notion of justification we are interested in is as robust as it would need to be if (J) or (WCJ) were true. That is, perhaps there is a perfectly respectable sense of justification which allows that one can simultaneously be justified in believing incompatible propositions, even when the incompatibility is known. If this were the case, then one could straightforwardly reject the sceptical argument just formulated which trades on (WCJ). Moreover, since I employed principle (J) in the argument to the effect that (CJ) entails (UP), it follows that the rejection of this principle could be cited as grounds for questioning whether (CJ) is logically stronger than (UP) after all. Accordingly, advocating a conception of justification which is inconsistent with (J) could be used as part of a wider strategy to reestablish the *status quo*, that an adequate response to the (CJ)-based sceptical argument would thereby also be an adequate response to the (UP)-based one. Hence provided that one could deal with the (CJ)-based sceptical argument, then one could bypass sceptical concerns which utilize (UP).

The pitfall of adopting this strategy should, however, be obvious. After all, it is now absolutely transparent just how limited a notion of justification

is being rescued from the sceptic's grasp. In particular, what is now manifest is that the only style of justification that evades the sceptic's grasp, if any, is one so weak as to permit *S* to have justified beliefs in propositions which are known to be incompatible. Such a victory against the sceptic is thus Pyrrhic at best. The more substantive issue, then, is how one is to respond to the sceptical problem as it is now configured, given that one grants the force of (J), and therefore (WCJ). Three main strategies suggest themselves.

The first is to deny that there is any essential connection between knowledge and justification. If one could make this move, then the original sceptical argument I considered in §I, which was construed in terms of the closure principle for knowledge (CK), would stand independently of any of the other sceptical arguments which I have been considering, whether they are expressed in terms of (UP), (WCJ) or, for that matter, closure for justification (CJ). Accordingly, a response to the former, such as an argument which denies (CK), would be unaffected by its failure to deal simultaneously with any of the latter.

It is certainly true that there is a brand of hardcore externalism about knowledge in the literature which is happy to incorporate a claim of this sort. The problem for such a proposal *qua* anti-sceptical strategy, however, is that it is not an intellectually satisfying resolution of the problem in hand. There are two interrelated reasons for this. The first is that if knowledge has no essential connection with justification, then it is unclear just what epistemic property for our beliefs we are rescuing from the sceptic, since, at the very least, the conception of knowledge seems to include a justification component. The second reason is more serious, and relates to how the conclusions of the sceptical arguments expressed in terms of justification considered above, even when they are restricted so as only to make a claim about our lack of justification, appear just as devastating as the conclusion to a sceptical argument which is expressed in terms of knowledge, such as the (CK)-based argument. It is thus little comfort to be told that the latter sceptical argument is blocked whilst simultaneously granting the conclusions of the former that the possession of widespread justification for one's beliefs is impossible.

The second potential anti-sceptical response to the (WCJ)- or (UP)-based sceptical arguments is to contend that there is no essential connection between justification and evidence. If one could defend this claim, then one could block these arguments by rejecting the assumption they both make, that it follows from the fact that one lacks evidence which would justify one's belief that one's belief is therefore not justified. This move might be initially appealing, because there certainly seem to be cases where justification is possessed in the absence of evidence (basic perceptual belief is an obvious

example in this regard). Getting this claim to stick across the board rather than just in a restricted range of cases, however, is a little more difficult. After all, would it really be plausible to suppose that justification could always (i.e., in all, or at least most, cases) be possessed in the absence of supporting evidence? Moreover, on the supposition that we could make sense of a widespread justification for our beliefs in the absence of supporting evidence, it is important to recognize how unsatisfying a response to the sceptic this would be, since it would mean that all of our justification is now on a par with the kind of 'brute' justification putatively possessed in, for example, the basic perceptual case. When we sought a response to justification scepticism it was surely not a response of this sort that we were looking for.

This leaves a third style of anti-scepticism which tries to argue that there is a sense in which we *do* have evidence which favours our beliefs in everyday propositions over sceptical hypotheses. In its boldest form, as found for example in recent work by Timothy Williamson, this view will straightforwardly defend the claim that our evidence is not to be identified with what is phenomenologically available to us in experience.[9] Accordingly, one's evidence when one is the victim of the (BIV) hypothesis is not identical with the evidence one possesses when one is not so deceived, even if the phenomenology of the two scenarios is exactly the same. On this view, the moral of scepticism is that we should radically reconfigure our conception of the notion of evidence.

A less radical, though similar, anti-sceptical response to the sceptical challenge based on underdetermination is to offer a contextualist account of evidential priority, so that what counts as evidence, and for what, varies in such a way as to undermine the sceptical argument. On this view, it *is* possible to have evidence in favour of one's everyday beliefs which prefers them over sceptical alternatives in normal contexts in which the epistemic standards are low, even though it is also true that in sceptical contexts where the epistemic standards are higher the evidential superiority is lacking. One way of spelling out this view, found in recent work by Ram Neta, is to claim that what the sceptic achieves by raising the epistemic standards is a restriction on what counts as evidence.[10] In normal contexts, our evidence can include, for example, our knowledge of certain facts regarding the external world (such as that I see that there is a chair before me), and so it is little wonder that our evidence in this context can favour our everyday

[9] See, e.g., T. Williamson, 'Scepticism and Evidence', *Philosophy and Phenomenological Research*, 60 (2000), pp. 613–28.

[10] See R. Neta, '*S* Knows that *p*', *Noûs*, 36 (2002), pp. 663–81, and 'Contextualism and the Problem of the External World', *Philosophy and Phenomenological Research*, 63 (2003), pp. 1–31.

beliefs over sceptical alternatives. In sceptical contexts, however, only non-world-involving evidence is permitted, and this is why our evidence no longer favours our everyday beliefs over sceptical alternatives.

Although this type of view faces the usual problems directed at contextualism, such as how to explain just what constitutes a context, it does have the advantage over Williamson's account of being able to give a diagnosis of why scepticism can seem so plausible. This advantage brings with it further problems, however. Once one grants that the sceptic is employing more austere epistemic standards, in the light of which it is true that our evidence does not favour our everyday beliefs over sceptical alternatives, then this prompts the thought that *strictly speaking* we do not have evidence in favour of our everyday beliefs, even though in looser and epistemically less demanding quotidian contexts we talk as if we had.

In any case, it is not my aim here to decide between these competing styles of response to the sceptical challenge that I have identified, only to set out what they are. The key point is that these responses to scepticism proceed in a very different way from their counterpart anti-sceptical proposals, which take the closure-based sceptical argument as their template. If we are to make any progress against the spectre of scepticism then it will be essential to refocus our efforts towards proposals of this sort, since it is only these types of views that hold out the hope of dealing with the very different kind of sceptical argument that I have identified here.[11]

University of Stirling

[11] Thanks to Michael Brady, Tony Brueckner, Antony Duff, Andrew McGonigal, Alan Millar, Chris Thornton and Jonathan Vogel, who all read and commented on drafts of this paper, and to the Leverhulme Trust for the award of a Special Research Fellowship, which has enabled me to conduct research in this area. I am also grateful to Jessica Brown and Ram Neta for discussion on related topics. Earlier versions of this article were presented at the Society for Sceptical Studies session at the 2002 Pacific APA in Seattle, and at the 2002 Joint Session of the *Mind* Association and the Aristotelian Society in Glasgow. I am grateful to the audiences on these occasions.