

RADICAL SCEPTICISM, EPISTEMOLOGICAL EXTERNALISM, AND CLOSURE

ABSTRACT. Many of the proposals put forward to answer the radical sceptical paradox in the recent literature have involved either blocking the sceptic's highly intuitive use of the 'closure' principle that knowledge is closed under known entailment or, equally implausibly, 'contextualizing' the knowledge operator. In contrast, it is argued here that an *externalist* theory of knowledge, properly understood, holds the key to meeting the sceptical paradox in a manner that keeps our basic epistemological intuitions intact. In particular, it is shown how, far from failing in sceptical arguments, the sceptic's use of the closure principle is in fact *validated* by an externalist epistemology. A diagnosis of the main argument for non-closure offered by Fred Dretske is then given in the light of these remarks.

1. RADICAL SCEPTICISM IN CONTEMPORARY DEBATE

There has been a dramatic resurgence of interest in radical sceptical arguments in the recent literature. Although, historically, the varieties of scepticism that have been proposed have been legion, the particular type of argument that the current debate has tended to focus upon can be roughly captured in terms of the following template:

- (S1) not: K [I am not the victim of a radical sceptical hypothesis].
- (S2) not: K [I am not the victim of a radical sceptical hypothesis] → not: K [Any one of a wide class of 'everyday' propositions (propositions that I ordinarily take myself to know)].
- (SC) not: K [Any one of a wide class of everyday propositions].¹

For example, suppose we plug the 'brain-in-a-vat' (BIV) sceptical hypothesis into this argument, the hypothesis that I am currently a disembodied brain sealed in a vat of nutrients being 'fed' my current 'experiences' by my captors (see Putnam 1981, chapter 1; Pollock 1984, 1-12). Furthermore, suppose that the paradigmatic everyday proposition in question is 'I have two hands'. If I know anything, then, or so it would seem at least, I must know this proposition. We thus get the following variation on the above template sceptical argument:

- (B1) not: K [I am not a BIV].
- (B2) not: K [I am not a BIV] → not: K [I have two hands].
- (BC) not: K [I have two hands].

This appears to be both an intuitively compelling argument and an intellectually devastating conclusion. The first premise seems perfectly acceptable because there is, *ex hypothesi*, no way of telling whether or not I am the victim of such a scenario. Moreover (and we will examine

this claim more closely in a moment), it also seems intuitive to argue that if I am unable to rule out such an error-possibility then I lack knowledge of much of what I believe, including the claim that I have two hands, since if I were a BIV then I would be radically mistaken about my relationship to the world. And given that these two premises hold, it follows, via the entirely uncontroversial logical principle of *modus ponens*, that I lack knowledge of this ‘everyday’ proposition. Moreover, since it makes no difference who the agent is here, and since one can run the very same argument with any number of ‘everyday’ propositions, it follows that radical scepticism—the view that I lack knowledge of much of what I believe—quickly ensues.

Most of the contemporary debate regarding radical scepticism has tended to take the first premise of the sceptical argument for granted and to focus instead upon either the status of the second premise or on the nature of the knowledge operator at issue in the argument. In their different ways, each of these proposals implicitly accept that what motivates the second premise of the sceptical argument is the so-called ‘closure’ principle that knowledge is ‘closed’ under known entailment:

$$(CK) \quad \{K[\varphi] \ \& \ K[\varphi \rightarrow \psi]\} \rightarrow K[\psi]^2$$

This is a highly intuitive principle. If I do indeed know that, say, the highest mountain in the world is Everest, then surely, given that I also know that if the highest mountain is Everest then it is not K2, I also know that the highest mountain in the world is not K2. The sceptic contraposes on this intuitive principle in order to derive the second premise of the sceptical argument. For if closure holds then, given that we might reasonably assume that we all know that (at least in the main) everyday propositions entail the falsity of radical sceptical hypotheses, it follows that a lack of knowledge of the falsity of a radical sceptical hypothesis will result in a lack of knowledge of everyday propositions.³ It is thus the closure principle that forges the connection between a lack of anti-sceptical knowledge and a lack of everyday knowledge that the sceptic needs in order to run her argument.

One prominent anti-sceptical camp has been those who, following Fred Dretske (1970) and Robert Nozick (1981), have argued for the rejection of this principle, and thus for the claim that knowledge of everyday propositions can happily coexist with a lack of knowledge of anti-sceptical knowledge. If successful, this would indeed be a strikingly simple way of dealing with radical scepticism, even if it does have the counterintuitive consequence of resulting in the denial of closure.

In essence, the element of their theory of knowledge which entails the rejection of closure is the following ‘modal’ condition that they place on our knowledge of contingent

propositions, what I will call (after its first exponent) ‘Dretskean Sensitivity’:

- (DS) A necessary condition of an agent’s knowledge of a contingent proposition, φ , is that she has a true belief in φ in the actual world and that, in the nearest possible world or worlds in which φ is false, she does not believe φ .⁴

As we will see below, a compelling case can be made for the condition on knowledge specified in (DS). Note, however, that one consequence of (DS) is that the range of possible worlds at issue in the evaluation of any particular knowledge claim will vary depending upon the proposition in question. When it comes to everyday propositions such as ‘I have two hands’, the possible worlds at issue will be those relatively near-by worlds where this is false, such as the world in which I lose one of my hands in a car accident. In contrast, the evaluation of the proposition ‘I am not a BIV’ requires one to consider the (hopefully) far-off possible worlds in which this is false, i.e., the nearest BIV-worlds. This difference in the range of worlds at issue is what secures the rejection of closure. Although one’s beliefs regarding everyday propositions might be sufficiently sensitive to near-by worlds to allow you to meet the standards imposed by (DS), it is extremely unlikely that one’s beliefs in anti-sceptical propositions will be able to do likewise. After all, one of the characteristics of, for example, BIV-worlds is that one continues to believe that one is not a BIV even though this is no longer true. With this condition in play one can thus explain how everyday knowledge can coexist with the absence of anti-sceptical knowledge, and hence why closure fails.

A second anti-sceptical stance that has come to the fore in more recent discussion has been that of contextualism, as propounded by, for example, Keith DeRose (1995) and David Lewis (1996). This is the thesis that the sceptic raises the standards for knowledge by presenting the sceptical argument in such a way that it makes the conclusion of that argument true at this epistemic standard. Nevertheless, that this is so is meant to be consistent with the possession of knowledge at everyday epistemic standards, and thus with the falsity of the sceptical conclusion at those standards. In this way, closure is retained because, at any one epistemic standard, knowledge will be closed under known entailment. Either one knows both everyday propositions and anti-sceptical propositions at the ‘low’ everyday epistemic standard, or one lacks knowledge of both everyday propositions and anti-sceptical propositions at the ‘high’ sceptical epistemic standard. This proposal shares with the Dretske-Nozick line an implicit acceptance of the centrality of the closure principle to the sceptical argument because the thought is clearly that one can only make sense of the retention of closure by contextualizing one’s understanding of the knowledge operator. In effect, and this is simplifying somewhat, what the contextualist does is to maintain that the knowledge operator

should be construed as relative to the range of possible worlds at issue, such that the greater the breadth of worlds involved, then the more demanding the construal of the knowledge operator.⁵

In what follows I will take the centrality of the closure principle to the sceptical argument for granted.⁶ What I will argue is that even with closure and a non-contextualized construal of the knowledge operator in play one can still meet the radical sceptical argument head-on. In order to see how this is possible, however, it is necessary to examine more closely the motivation for the first sceptical premise, (S1).

2. EPISTEMOLOGICAL INTERNALISM AND THE SCEPTICAL PREMISE

What is granted (or at least, *ought* to be granted) by every party to the sceptical dispute is that if the actual world is, in fact, like one of the scenarios that the sceptic describes (if I am, for instance, a BIV being ‘fed’ experiences), then I will not know very much. I will certainly not know that I am *not* a victim of one of these scenarios because, apart from anything else, the truth condition does not obtain—my belief would, in this case, be false. Crucially, however, this reasoning is unable, in itself, to support the sceptical argument for the simple reason that it is not sufficient to motivate support for the first sceptical premise, (S1). After all, we might all agree that *if* the actual world is a *sceptical* world,⁷ *then* we lack the sort of knowledge specified in the sceptical premise (and, for that matter, specified in the conclusion, (SC)). But this claim is much weaker than that contained in the sceptical premise, since this in no way depends upon the agent being in any particular world, much less being in a sceptical world. What this point emphasises is that what the sceptic requires in order to destroy our knowledge is something stronger than the contention that we might be, say, BIVs and that if we were we would not know a great deal. Rather, she needs to make the much grander claim that even if we are not the victims of sceptical hypotheses—even if, indeed, the actual world is pretty much as we take it to be—still, we do not know very much.⁸

Typically, the kind of reasoning that the sceptic adduces in order to bridge this gap presupposes a commitment to the doctrine of epistemological internalism. For the purposes of this paper I will understand this doctrine in terms of the following core thesis, which is itself a slightly amended version of a characterisation of the basic component of the internalist position put forward by James Pryor (2001, §5):

Epistemic Internalism

The positive epistemic status of an agent’s belief is wholly determined by facts which she is in a

position to know by reflection alone.

Also following Pryor (2001, §5), I will understand the type of ‘reflective’ knowledge at issue in this formulation to include *a priori* reasoning, introspective awareness of one’s own mental states, and one’s memory of knowledge that has been gained in either of these ways. With this account in mind, we can offer the following minimal characterisation of the core internalist thesis regarding knowledge:

Simple Internalism

A necessary condition of an agent’s knowledge of a proposition, φ , is that her belief that φ be supported by a sufficient degree of positive epistemic status, where this is understood along epistemically internalist lines.

Although this leaves some room for interpretative manoeuvre (especially regarding the proper construal of the word ‘sufficient’), it does capture what is at least essential to the internalist position (and thus that component of the position which, as we will see, makes it conducive to scepticism).⁹ For although internalism is consistent with allowing an ‘external’ condition (over and above the truth-condition) to play a role in the possession of knowledge—that is, a condition which can obtain without it being required that the agent has reflective knowledge of the relevant facts—what is distinctive about the internalist view is that it will not permit knowledge to be possessed without the appropriate corresponding ‘internal’ condition being met. For example, an internalist could consistently demand that a necessary condition for knowledge be that the belief was formed in a reliable fashion. What the internalist could not do, however, was incorporate *just* this condition. Instead, since reliability would now be one of the factors that determines positive epistemic status, it would follow that the internalist would have to further demand that the agent needs to be able to have reflective knowledge of the relevant ‘reliability’ facts if she is to have knowledge of the proposition in question. Accordingly, whatever other constraints the internalist might put on her account of knowledge—and notwithstanding the manner in which the details of this characterisation are filled-in—this simple internalist thesis will form part of her final, considered, position.

For the purposes of this paper we will understand an externalist theory of knowledge to be any view which denies the simple internalist thesis sketched above and thus allows that one might know a proposition without being able to gain reflective knowledge of the facts which determine the positive epistemic status that makes up that knowledge.¹⁰ For example, the type of theory of knowledge that Dretske and Nozick outline is externalist because it only demands that the conditions for knowledge—and, in particular, the condition specified by (DS)—should be met, it does not further demand that the agent should be able to reflectively know the

relevant facts which support the epistemic status of her belief in each case. The standard form of contextualism is also configured along externalist lines. Typically, contextualist theories merely demand that the agent meet the relevant contextual standard without further demanding that the agent should be able to reflectively know the relevant epistemic facts as well.¹¹ Note that externalism, so construed, is consistent with a very mild view which accords internalist features of knowledge a central epistemological role. What is being explicitly denied is only their *necessity*, not also their centrality to our epistemic concepts. Accordingly, just as internalists can allow that reliability may be a factor in determining an agent's knowledge possession, so externalists can claim that the ability to gain certain sorts of reflective knowledge can, in certain cases, be crucial to knowledge possession.¹²

In the main, the relationship between scepticism and internalism is forged with reference to the sceptical premise—the claim that we lack knowledge of the denials of sceptical hypotheses. As noted above, it is an entirely uncontentious point (acceptable to both internalists and externalists) that, if we were the victims of a sceptical scenario, then we would not know very much. What makes this uncontentious thought such that it can support the sceptical premise—and thus act as a spur for scepticism—is the contention that, whatever the nature of the actual world (sceptical or otherwise), in order to know that we are not the victims of a sceptical hypothesis we must be able to *reflectively* know those facts which determine this knowledge. The problem is that sceptical hypotheses are carefully chosen such that, from the agent's point of view, there will be no detectable difference between these scenarios and everyday 'non-sceptical' circumstances—*ex hypothesi*, there is no phenomenological difference between a subject's experience of being a BIV and her experience of, say, sitting at her computer typing. Accordingly, there is not going to be anything that is reflectively available to the subject that would indicate to her that she was not the victim of such a scenario and thus it is impossible for the agent to possess the kind of reflective knowledge demanded by the internalist account. It follows that on the classical internalist view one is *never* able to know that one is not a BIV, regardless of whether or not one is actually a victim of this scenario (or whether it is even unlikely that such a scenario could ever occur). And what applies in the BIV case will hold, *mutatis mutandis*, for other sceptical hypotheses. A commitment to the internalist doctrine will thus be sufficient to motivate the sceptical premise in the unrestricted form that it appears above no matter what world one inhabits, sceptical or otherwise. And once the sceptical premise has been motivated then one is immediately face-to-face with the harsh choice of either accepting scepticism or adopting some form of epistemological revisionism, whether in the guise of contextualism or the rejection of closure.

In contrast, on an *externalist* account it is at least a further issue as to whether or not one knows the sceptical premise. As we just noted, an externalist account of knowledge can exploit the distinction between a lack of knowledge, internalistically construed, and a lack of knowledge *simpliciter*. For although it is certainly true that one lacks knowledge, internalistically conceived, of the denials of sceptical hypotheses, this does not entail, for an externalist, that one lacks knowledge, externalistically conceived, of the sceptical premise. This latter contention may depend upon the actual facts of the situation.

Indeed, since, on the externalist view, the internalist makes an extra demand on knowledge, we can usefully distinguish between two sorts of knowledge operators in play here.¹³ On the one hand, there will be knowledge that merely meets the essential external conditions, such as a reliability condition of some sort. We will represent this type of knowledge via the subscripted ‘K’ operator, ‘K_E’. On the other, there will be knowledge which also meets the inessential, though perhaps desirable, condition universally demanded by the internalist—the ability to reflectively know those facts which determine the positive epistemic status at issue in that knowledge. We will represent this type of knowledge via the subscripted ‘K’ operator, ‘K_I’. In effect, what an externalist will argue (especially those who advocate the milder versions of this thesis) is that whilst all knowledge meets the externalist criteria, not all knowledge meets the internalist rubric (though some of it might). *Prima facie*, then, the externalist is in a position to qualify her understanding of the sceptical premise as follows:

(S1’) not: K_I [I am not the victim of a radical sceptical hypothesis].

This explicitly states that the knowledge that one lacks of anti-sceptical propositions is of a specifically internalist variety. And since one could, on the externalist account, lack knowledge internalistically construed of a proposition and yet still (externalistically) know it, it follows that this premise would not suffice to support the sceptical argument as it stands. On an externalist theory of knowledge, then, if the sceptic wishes to advance her unrestricted sceptical premise then she will have to do more than merely establish a case for (S1’). The issue thus comes down to whether there is a plausible way of spelling-out the externalist thesis such that it can account for our knowledge of the sceptical premise. And, indeed, there is.

3. SIMPLE EPISTEMOLOGICAL EXTERNALISM AND CLOSURE

Consider again the key externalist intuition—that one can have knowledge, at least in certain

cases, merely because one meets some external condition or conditions. Typically, this thought will be cashed-out in terms of some sort of reliabilist thesis to the effect that what is fundamental to the possession of knowledge is our mere responsiveness, via the processes we utilise to form our beliefs, to the truth across a range of relevant situations—regardless of whether we are able to reflectively know those facts which determine the positive epistemic status at issue in this knowledge. There are, of course, going to be many different ways of spelling-out this thought. In general, though, I think we can get a neat, albeit incomplete, characterisation of the core externalist thesis as follows:

Simple Externalism

Penumbral cases aside, an agent knows a contingent proposition, φ , if, and only if, she has a true belief in φ and her belief in φ is *sensitive* to the truth as to whether φ across a wide-range of near-by possible worlds (i.e., where φ is true, she believes it and where φ is false, she does not believe it).

This characterisation captures the heart of the position in that it conveys the key externalist intuition that knowledge is essentially dependent only upon being so related to the world that one not only believes the truth in actual conditions, but that one's belief in this respect is also responsive to a wide-range of relevantly similar conditions.¹⁴ Moreover, although this is an externalist thesis in the sense that it makes no demand that the subject should be able to reflectively know the relevant facts regarding her 'sensitivity', it is entirely consistent with the sort of 'mild' externalism noted above which allows internal factors to exercise a pivotal role. After all, it may well be that, in a wide range of cases, what makes it such that the subject's belief sufficiently tracks the truth as to whether φ is simply that she is responding to facts which she is able to know by reflection alone.¹⁵

Note that one key difference between this characterisation of knowledge and the condition placed on knowledge by (DS) is that it keeps the range of possible worlds at issue—what we will henceforth call the 'domain of sensitivity'—fixed. That is, the range of worlds at issue does not vary in response to the proposition at issue as it does with (DS), such that sometimes it concerns near-by worlds and sometimes far-off worlds. Instead, the range of worlds at issue only concerns near-by worlds. In this sense it is both more and less demanding as a constraint on knowledge than (DS). More demanding, because it requires one to consider a *range* of near-by worlds, regardless of whether the proposition in question is false in those worlds. Less demanding, because it *only* demands consideration of those worlds. Whatever the proposition at issue, far-off possible worlds are not relevant to the determination of knowledge. Likewise, it differs from the contextualist thesis because it makes no mention of variable contextual standards. As far as the simple externalist thesis is concerned, there is just *one*

epistemic standard that we need to meet. We will consider the relative merits of these three accounts of knowledge (simple externalism, the Dretske-Nozick proposal, and contextualism) in more detail below. For now, we will concentrate on exploring the consequences of this simple externalist thesis.

What is interesting about such an account of knowledge is that it actually *licenses* the closure-type inferences that the sceptic makes. In order to see this, consider what a possible ‘sceptical’ counterexample to closure would consist in by the lights of the simple externalist account of knowledge described above. We would require an agent both to have a sensitive belief in an everyday proposition which she knew entailed the denial of a sceptical hypothesis and for that agent to have an *insensitive* belief in the denial of the sceptical hypothesis. Moreover, in order for this to be a decisive counterexample to the principle, it had better be clear-cut both that the subject *lacks* knowledge of the consequent proposition (the denial of the sceptical hypothesis) and that she *possesses* knowledge of the antecedent proposition (the everyday proposition).

For example, suppose the agent has clear-cut knowledge of the everyday proposition, P, “I am wearing my black shoes”. In the actual world this proposition is true, and the agent’s belief tracks the truth across a wide range of near-by possible worlds. This means that in close-by worlds in which she is wearing her black shoes, she believes that she is, whereas in other close-by worlds in which she is not wearing her black shoes (but, say, her brown ones), she does not believe that she is (because, for example, she just looked).

Clearly, since counterfactual sensitivity is necessary for knowledge, the range of possible worlds included in the determination of the agent’s knowledge of the everyday proposition cannot include *sceptical* worlds—worlds in which the agent is the victim of a widespread deception—because in such worlds the agent’s beliefs will not be appropriately sensitive to the facts. If there were possible worlds in the domain of sensitivity in which, for instance, the radical sceptical BIV hypothesis is true, then there would be worlds in the domain of sensitivity in which the agent believes P even though it is false, or does not believe P even though it is true. The agent would thus clearly lack knowledge of P.

Indeed, given that the agent’s knowledge of P is meant to be clear-cut, hence even *borderline* cases of sceptical worlds will be ruled-out. That is, even if there are no sceptical possible worlds in the domain of sensitivity, as long as there are such worlds *near* to the borderline of the domain of sensitivity then this will be sufficient to destroy the *clear-cut* knowledge of P required for the counterexample. The agent’s clear possession of knowledge of P thus entails that there are no sceptical possible worlds either in, or near-by, the domain of

sensitivity.

Now consider the following ‘anti-sceptical’ proposition, which is entailed by P and which the agent also believes: “I am not a BIV” (not: BIV). Given that, as we have just seen, the sensitivity of the agent’s belief in P requires that there be *no* sceptical possible worlds either in or near to the domain of sensitivity, how would it be possible for the agent to *not* have a (true) belief in not: BIV which was sensitive? Since such a sceptical scenario is, *ex hypothesi*, far-fetched, it follows that all the agent needs to do in order to have a sufficiently sensitive belief in not: BIV (and thereby have externalist knowledge of not: BIV) is simply retain her belief in not: BIV in all of the resolutely non-sceptical worlds contained within the domain of sensitivity.

Of course, one might object that the flaw in this reasoning is that it is perfectly conceivable that there may be some possible worlds in the domain of sensitivity in which the agent does not believe not: BIV. If this were so then she would lack a sensitive belief in not: BIV and thus externalist knowledge of not: BIV. Now one might question this claim on the basis that it is at least dubious to suppose that anyone could fail to believe the denials of sceptical hypotheses,¹⁶ but we need not rest our response on this contentious manoeuvre. This is because whatever grounds that might be offered to support the claim that the agent lacks the belief in not: BIV—perhaps, for example, she would be convinced by a sceptical argument if she heard one—would be such that they would interfere with the sensitivity of the subject’s belief in P in the first place. If there are near-by possible worlds in which the agent is convinced by scepticism, for instance, and so does not believe not: BIV, then, accordingly, there will be near-by worlds in which she does not believe P either, even though P is true in those worlds. It thus won’t be possible to find a case in which one has a sensitive belief in an everyday proposition and one knows that this entails the denial of a sceptical hypothesis, and yet one lacks a sensitive belief in the denial of the sceptical hypothesis.

This result—that closure-type inferences hold in sceptical scenarios—should not surprise us. If one were to ask the question of whether or not a subject knows not: BIV independently of stipulating what else the subject knows, then the orderings of the possible worlds could be different in such a way as to allow for the subject to not have a sensitive belief in not: BIV and therefore not know not: BIV. But that is not the question. Rather, we need to ask whether the subject has a sufficiently sensitive belief in not: BIV (and therefore knows not: BIV), *given* that the subject has a sufficiently sensitive belief in P (and therefore knows P), and knows that P entails not: BIV. In such a case, the orderings of the possible worlds are set such that the relevant sceptical possible worlds must, perforce, be well beyond the domain of

sensitivity. As a result, it should not astound us in the slightest to find that closure holds in these ‘sceptical’ examples. Moreover, this result also confirms to a deeply held intuition that we have that if we know anything much at all, then we must know the denials of sceptical hypotheses, despite the fact that this is not the sort of knowledge which, as the sceptic emphasises, we can reflectively recover.¹⁷

Prima facie, then, there is an account of knowledge available that both licences anti-sceptical closure-type inferences and which can also explain, in principle, how it is that we know the denials of sceptical hypotheses. Moreover, unlike contextualism, this thesis objects to the first sceptical premise *simpliciter*, rather than just relative to low standards of knowledge. We thus have the beginnings of the non-contextualist anti-sceptical thesis that retains the sceptic’s use of closure which was promised at the outset of this paper.

4. TWO TYPES OF CLOSURE PRINCIPLE

We are a long way from being out of the woods yet, however. On the one hand, we need to give some account of why it is we should make the type of externalist proposal put forward above central to our epistemology rather than the competing Dretske-Nozick proposal, (DS), which has the consequence that we lack knowledge of the denials of sceptical hypotheses and thus results in the rejection of the sceptic’s use of closure. After all, although we have the intuition that if we know anything at all then we must know the denials of sceptical hypotheses, there is also the intuition that we could not know such anti-sceptical propositions. On the other hand, we need to explain how it is that the externalist account sketched here does not fall foul of a related radical sceptical argument. As we will, see, the answers to these two problems are closely related.

So far we have shown that there is an externalist account of knowledge available that can, in principle, accord us with knowledge of the denials of sceptical hypothesis and which accepts the sceptic’s use of the closure principle. Nevertheless, we have also conceded that a mild version of externalism ought to be able to allow that at least some knowledge meets the internalist rubric, and thus is ‘internalist’ knowledge, and that we could never know the denials of radical sceptical hypotheses in this internalist fashion. Given these two claims, however, the sceptic appears to be able to run a very similar radical sceptical argument to the one that we considered in §1. Unlike that argument, however, this revised sceptical argument is directed at *internalist* knowledge:

- (S1') not: K_I [I am not the victim of a radical sceptical hypothesis].
 (S2') not: K_I [I am not the victim of a radical sceptical hypothesis] \rightarrow not: K_I [Any one of a wide class of 'everyday' propositions].
 (SC') not: K_I [Any one of a wide class of everyday propositions].

If this argument were conceded, however, then it would appear to jeopardise the externalist programme by forcing it into the stronger, and less palatable, position of arguing for the complete lack of any knowledge that met the internalist rubric.

The important thing to note about this argument, however, is that it does not use the closure principle at all, at least as it is described above, but rather a more demanding sister principle to closure that subscriptes the knowledge in question to that which meets the internalist rubric, an inessential component of knowledge on the externalist account. Accordingly, this holds out the prospect that we might be able to deny this sister closure principle—and therefore block this 'internalised' version of the sceptical argument—whilst keeping our commitment to the relevant instances of the core closure principle (and thus our basic anti-sceptical thesis) intact.

To begin with, we need to be clear about the differences between these two principles. Accepting the sceptic's use of the closure principle means allowing the following type of inference, where 'E' refers to an everyday proposition, and 'not: SH' refers to the denial of a relevant radical sceptical hypothesis:

$$\{K [E] \ \& \ K [E \rightarrow \text{not: SH}]\} \rightarrow K [\text{not: SH}]$$

However, *prima facie* at least, for an externalist this is not the same as accepting the following type of inference that features in the 'internalist' sceptical argument above:

$$\{K_I [E] \ \& \ K_I [E \rightarrow \text{not: SH}]\} \rightarrow K_I [\text{not: SH}]$$

We will call the principle that underlies the 'internalist' sceptical argument '(CKI)'. It can be represented formally as follows:

$$(CKI) \quad \{K_I [\varphi] \ \& \ K_I [\varphi \rightarrow \psi]\} \rightarrow K_I [\psi]$$

Whereas (CK) only demands that knowledge transfer across known entailments, (CKI) makes the stronger demand that the *type* of knowledge at issue should transfer across known entailments.¹⁸ Accordingly, it is a logical possibility that one could motivate a denial of the stronger principle, (CKI), and therefore block the internalist sceptical argument considered above, whilst leaving (CK) intact. What one would need to show is that although knowledge

simpliciter transfers across known entailments (in line with (CK)), this does not mean that the *type* of knowledge always transfers across known entailments (*contra* (CKI)).

Making a distinction is one thing, however, drawing dialectical consequences from that distinction quite another. After all, or so the thought might run at any rate, if the relevant closure-type inferences are allowed then surely one has no choice but to sanction the analogous instances of internalist closure as well. If this thought is sound then this contrast between (CK) and (CKI) is simply a distinction without a difference. It is this thought that I want to challenge here. Indeed, as we will now see, the standard argument against closure is, in fact, when properly understood, an argument against *internalist* closure. Accordingly, by showing how the standard argument for non-closure is in fact an argument against (CKI), we both motivate the denial of the latter principle whilst offering the requisite grounds for preferring the basic externalist account of knowledge sketched here over the competing Dretske-Nozick thesis. Moreover, we also explain why the ‘intuition’ that we do not know the denials of sceptical hypotheses can, at least sometimes, seem so strong.

5. DRETSKE’S ARGUMENT FOR NON-CLOSURE

Dretske (1970; 1971; cf. Nozick 1981 *et al*) was the first to offer a sustained attack on the standard closure principle (CK), and his critique remains the most persuasive.¹⁹ I will thus focus my remarks on his account. *Mutatis mutandis*, however, my critique will apply in equal measure to any account that follows the basic ‘Dretskean’ line of endorsing (DS), including Nozick’s.

In ‘Epistemic Operators’ (1970), Dretske pointed out that it need not be surprising that epistemic operators should fail to ‘penetrate’ across entailments (even known ones) when there are many other operators that also fail in this way.²⁰ Examples that Dretske offers include operators such as “It is extraordinary that ...”, or “It is a mistake that ...”. That I bumped into my next door neighbour on the summit of mount Vesuvius is certainly an extraordinary occurrence, but it is not extraordinary at all (given the extent of tourism in that region) that I bumped into *someone* at this summit, and this latter claim is entailed by the former. In any case, Dretske argues, it is not as if the knowledge operator *completely* fails to penetrate across known entailments since there are only a certain class of cases in which closure fails. As a consequence, he can stay true to our intuition that we often gain new knowledge in this way.²¹ In essence, the class of entailments over which knowledge will not penetrate are those where

the truth of the entailed empirical proposition is already taken for granted in the agent's knowledge of the antecedent empirical proposition. The intuition is that one cannot lift oneself up by one's own bootstraps—one cannot gain empirical knowledge of what is simply presupposed in another item of empirical knowledge whilst staying in the same epistemic position. Either one had the knowledge of the consequent proposition in the first place, in which case this proposition was not being presupposed after all, or else one didn't, in which case it was being presupposed and further empirical investigation will be required in order to come to know it.

An example will clarify matters. Consider Dretske's famous (putative) counterexample to closure where the relevant propositions are as follows:

- (A) The animals in the pen are zebras.
- (B) The animals in the pen are not mules cleverly disguised to look like zebras.

Now consider the scenario of someone at the zoo looking into the zebra cage who has no special knowledge about zebras or mules. If she is in a good position to judge that the animal is a zebra, and it is in fact a zebra, then, by any usual standard of knowledge, we would be willing to ascribe to her knowledge of (A). However, despite the fact that (A) entails (B), it is (according to Dretske) clearly the case that the subject in our example does not have adequate evidence to support knowledge of (B), this proposition merely being something that she is taking for granted (quite legitimately as far as Dretske is concerned) in gaining her knowledge of (A) in the first place. She is not, for instance, an expert at picking out the sort of deception at issue in (B) (a cleverly disguised mule and a zebra look exactly the same to her). Dretske thus concludes that one's belief can be appropriately reliable with respect to a certain proposition, and so known, even though this proposition entails a second proposition which the subject is not in a reliable relation towards and so does not know.²² And this will be so regardless of whether the agent in question knows that the entailment holds.²³

For Dretske, then, knowledge of everyday propositions will always involve assuming that certain far-fetched error-possibilities do not obtain.²⁴ Typically, however, although such error-possibilities are assumed to be false, they are not *known* to be false, and Dretske's point is the extremely plausible one that we cannot come to know them merely by inferring such knowledge from our everyday knowledge which takes the falsity of these error-possibilities for granted. And what applies to zebras in zoos will apply with equal force to BIVs and *malin génies*. Knowing an everyday proposition will require the denials of sceptical hypotheses to be true, but whatever grounds I have for my knowledge of the everyday proposition, they cannot

sufficiently transfer through to the presuppositions of my knowledge and thereby give me knowledge of what is presupposed in my knowing (i.e., the denials of sceptical hypotheses). Thus, according to the Dretskean line, we can both know everyday propositions whilst failing to know the denials of sceptical hypotheses. Closure is thus denied and the sceptical paradox is therefore blocked.

6. RECONFIGURING THE ARGUMENT FOR NON-CLOSURE

Contra the Dretskean proposal, however, we have found that on the simple externalist account of knowledge sketched here if closure-type inferences are acceptable anywhere, then they are acceptable in sceptical contexts. What then should we make of Dretske's putative 'sceptical' counterexample to closure? Let us take it for granted that the agent in Dretske's zebra example does indeed know the antecedent proposition, (A). Moreover, given that this is meant to be everyday knowledge (and that we are content to endorse a mild form of externalism), it should not be problematic to concede that the knowledge possessed is of an internalist variety. Since we are here retaining closure-type inferences in sceptical contexts, it follows that the agent's internalist knowledge of (A)—conjoined with her (no doubt internalist) knowledge of the entailment—*will* result in her also knowing the consequent proposition. Does Dretske offer grounds for thinking that the subject does not know the consequent proposition? Not exactly. If one looks closely at how Dretske motivates his counterexample to closure, it is clear that what he is actually arguing is that the subject lacks *internalist* knowledge of the consequent proposition, which is a different claim entirely (especially if, like Dretske, one is an externalist). As Dretske explicitly puts the matter:

If you are tempted to say [*that the agent does know (B) ...*], think for a moment about the reasons that you have, what evidence you can produce in favour of this claim. The evidence you *had* for thinking them zebras has been effectively neutralised, since it does not count toward their *not* being mules cleverly disguised. Have you checked with the zoo authorities? Did you examine the animals closely enough to detect such a fraud? (Dretske 1970, 1016)

In other words, in whatever sense we may regard the agent as knowing this proposition, it cannot be in terms of the evidence that the agent *is in a position to adduce*. And to characterise knowledge in these terms can only mean that Dretske has an internalist notion of evidence in mind here and thus an internalist notion of knowledge. That is, Dretske is construing the relevant facts which determine the epistemic status of the agent's beliefs in such a way that it is being implicitly presupposed that the agent must be in a position to reflectively know these

facts—hence the stress on evidence that the agent is in a position to adduce. As he intimates, had the subject undertaken more strenuous tests of her knowledge so that there was more and better evidence that she could reflectively access, then she might have been in a position to know this proposition after all. Despite being an externalist, then, Dretske is nevertheless employing internalist standards in his characterisation of the agent's knowledge of the consequent proposition.

In any case, it is rather odd to suppose that the agent does not know, even internalistically, the consequent proposition (B) given that she does internalistically know (A). After all, given that she does indeed know (A), it follows that the kind of sceptical worry highlighted in (B) must be far-fetched for otherwise it would undermine the knowledge the agent has of (A) in the first place. It is irrelevant, then, to demand that the subject should check with the zoo authorities before she could be said to know (B), since if this is required in order to have internalist knowledge of (B) then it ought to be likewise demanded in the case of internalist knowledge of (A). Accordingly, and especially given that the subject knows that the entailment holds, what could be wrong with the subject's reflectively accessible grounds for belief in (A) such that they are not also reflectively accessible grounds for belief in (B)? Seemingly, if the agent does have internalist knowledge of (A), then she has internalist knowledge of (B) also.²⁵

It would appear then that Dretske has something stronger in mind as regards (B) since, so defined, (B) is unable to do the job that he asks of it. The problem with (B) is that it is only a *local*, rather than a *radical*, sceptical hypothesis. Given that the subject does indeed have sufficient reflectively accessible grounds for her belief in (A), it follows, (given her internalist knowledge of the entailment and granted the remarks just made), that she will have sufficient reflectively accessible grounds for (B) also. What Dretske requires, in contrast, is something more akin to a radical sceptical hypothesis—an hypothesis the belief in the denial of which *no* reflectively accessible grounds could be found to support. Only then would Dretske get the required tension between internalist knowledge of (A) and a lack of (internalist) knowledge of (B). Whereas the local sceptical hypothesis is unable, given one's internalist knowledge of (A), to call one's reflectively accessible grounds into question, a radical sceptical hypothesis challenges the very idea that one has reflectively accessible grounds in the first place. Of course, in the hands of the sceptic, this kind of claim is meant to be enough to destroy internalist knowledge of the antecedent proposition, (A). Dretske, however, wants to use the same reasoning in order to try to motivate instead the denial of the closure principle that supports this move. In effect, he wants to play on the appeal of scepticism as a means of

arguing to a *reductio* on closure.

In order to see this, consider the following analogous scenario. Patrick meets Jesper coming out of Michael's office and asks if Michael is in there. Jesper says (truly) that he is, and, let us say, he has excellent reflectively accessible grounds to support this claim (for example, he just this moment saw him in there). Imagine now that Patrick challenges Jesper by asking how it is that he can be sure, given that Michael might have just jumped out of his (third-storey) window. Radical scepticism aside, to such an objection Jesper is quite within his rights to be dismissive since, after all, the excellent reflectively accessible grounds he has to support his belief that Michael is in his office right now ought to also suffice to support his belief that he hasn't just jumped out of the window. That is, insofar as Jesper does internalistically know that Michael is in his office, then it seems that he also internalistically knows that he has not just jumped out of the window. (Conversely, were it to transpire that Jasper lacks good reflectively accessible grounds for supporting his belief that Michael has not just jumped out of his window—and thus that he lacks internalist knowledge of this proposition—then this would appear to entail that he does not internalistically know that Michael is in his office after all). Now suppose that Patrick asks how Jesper can be sure that Michael is in his office given that he might be a BIV. Of course, Jesper cannot claim to have sufficient reflectively accessible grounds to dismiss this error-possibility, since no such grounds exist. Accordingly, he must concede defeat on this point.

This is the juncture at which the sceptic usually argues for an extension of the defeat by querying the agent's (in this case Jesper's) knowledge of the antecedent proposition (that Michael is in his office). For what has happened now is that we have moved from a situation in which one grants that reflectively accessible grounds are available and simply challenges the conclusions drawn from those grounds, to one where one attacks the very sufficiency of the reflectively accessible grounds themselves. As a result, within an internalist framework at least, insofar as the agent really does lack internalist knowledge that he is not a BIV, then, argues the sceptic, neither can he know that Michael is in his office.

Dretske, however, wants it both ways. He both wants the reflectively accessible grounds that support knowledge of the antecedent proposition to stand, *and* also to challenge the very possibility of such grounds existing. The manner in which he tries to effect this feat is by using an example which is sufficiently like a radical sceptical scenario that it can play upon radical sceptical intuitions, whilst being different enough to enable him to adduce conclusions outwith scepticism. As I have just argued, however, this manoeuvre does not work. Either Dretske keeps with the locally sceptical scenario, in which case the agent's internalist

knowledge that (A) will transfer to internalist knowledge that (B) and thus he does not have a counterexample to closure (or (CKI) for that matter), or he goes with the radical sceptical example, and so does battle with the sceptic.

One possible formulation that might capture the consequent proposition that Dretske has in mind could be (B*):

(B*) There is not some wholesale, yet undetectable, conspiracy to deceive going on.

Clearly the agent in Dretske's example (indeed, *any* agent) lacks sufficient reflectively accessible grounds to support internalist knowledge of this proposition. Of course, this might be considered to be a good reason to endorse scepticism about the agent's knowledge of (A), especially if one endorses an internalist epistemology and the closure principle. The Dretskean line, on the other hand, takes the opposite tack of maintaining that an agent's (internalist) knowledge of (A) can coexist with that very same agent's lack of (internalist) knowledge of the (B*)-type proposition, even if the agent knows that (A) entails the (B*)-type proposition. Dretske thus implicitly uses a (B*)-type proposition in order to motivate support for the denial of what he takes to be the closure principle.

Crucially, however, these considerations have no capacity in themselves to militate against *externalist* knowledge of (B*)-type propositions, and thus against the sceptic's use of the closure principle, provided we keep (CK) and (CKI) well apart. Indeed, as we saw above with our account of the core externalist thesis, since we have granted that the agent does know the antecedent proposition it is going to be impossible to find a possible world in which the agent has a sensitive belief about the antecedent proposition, and so has externalist knowledge of this proposition, but in which the agent does not have a sensitive belief about (and thus externalist knowledge of) the consequent proposition. Simply in virtue of having a sensitive belief about (A), the agent must, perforce, have a sensitive belief about an anti-sceptical proposition like (B*) and therefore have externalist knowledge of (B*).

Dretske obscures this point by doing two things. The first is that, despite being an externalist, when he comes to considering the agent's knowledge of the consequent proposition he focuses on the grounds that the agent can adduce to support her knowledge, and thus subtly shifts his demand from knowledge *simpliciter* to specifically *internalist* knowledge. A related issue in this respect is that Dretske plays upon the fact that the knowledge is question, if it is possessed at all, certainly could never be properly *claimed* to be possessed. He does this by first asking us whether we are "tempted to say" that the agent knows, and then proceeding to put us in the agent's shoes, thereby implicitly querying whether, if *we* were in such a position,

we would think it appropriate to self-ascribe knowledge. The answers one gets in the two situations may come radically apart, however. It certainly seems plausible to suppose, for example, that one can never properly claim knowledge when the knowledge possessed is merely of an externalist variety because what matters when it comes to the propriety of knowledge claims is not the truth of what is claimed so much as the grounds that one can offer in support of that claim. And since the possessor of mere externalist knowledge cannot give very good grounds at all, it follows that it may be improper to claim such knowledge even if one does indeed have it. The implicit appeal to the propriety of a claim to know is thus extremely unhelpful.²⁶

The second thing that Dretske does which obscures the issue is that, as the quotation above illustrates, he asks us the *wrong* question. Rather than asking us whether the agent knows the consequent proposition *given that he knows the antecedent proposition*, which would make clear that the ordering of the possible worlds in question was resolutely in the anti-sceptic's favour, he instead asks, independently of any thing else that the subject may know, whether or not the agent knows that proposition. Accordingly, he leaves the ordering of possible worlds open. However, we can perfectly well concede that, *considered independently of whatever else a subject may know*, there is no reason for thinking that she knows that she is not the victim of a sceptical scenario, without thereby conceding that, *given that the subject knows a number of relevant everyday propositions* (and thus that sceptical possible worlds do not figure in the domain of sensitivity), she does not know that she is not the victim of a sceptical hypothesis. As we saw above, the answer we give under the latter supposition could well be entirely different.

There is thus nothing in Dretske's reasoning which need trouble a proponent of the sort of simple externalism outlined above and who hence maintains that if the agent in question does know the antecedent proposition then she must know the consequent proposition. By the lights of this view, what Dretske has shown is not that one can know everyday propositions whilst lacking knowledge of the denials of sceptical hypotheses, but rather that one can internalistically know everyday propositions whilst lacking *internalist* knowledge of the denials of sceptical hypotheses. In effect, he has given us an argument against (CKI), not (CK). Indeed, only an internalist would conclude from a lack of internalist knowledge of a certain proposition that it therefore follows that the proposition is, *ipso facto*, unknown. And since Dretske is explicit in his endorsement of externalism, we can conclude that his argument against closure is based on a confusion. Moreover, with a motivation for denying internalist closure in the bank, the onus is back on to the sceptic to provide further grounds for her doubt.

We can thus account for the phenomenon that Dretske and the sceptic claim to discover without either rejecting the sceptic's use of closure or denying that we know many of the propositions which we believe, including the denials of sceptical hypotheses. It is this last point that is crucial to Dretske's overall account of knowledge, because it is only by implicitly privileging an internalist notion of knowledge and then taking one's epistemic position as regards 'sceptical' propositions (such as (B*)) in isolation that he is left with an epistemology which straightforwardly accepts the sceptical premise that we do not know the denials of sceptical hypotheses. As we have just seen, if he had endorsed a thorough-going externalist epistemology then Dretske could have respected the distinction between internalist and externalist knowledge in such a way as to deny this premise in an unqualified form and thereby concede the sceptical use of closure whilst evading scepticism. Once Dretske has argued on this basis for the denial of closure, however, he is thereafter committed to endorsing an epistemology that is able to incorporate these epistemological 'facts' (i.e., the denial of closure and the acceptance of the sceptical premise).

Dretske's synthesis of these disparate epistemic elements is, however, to say the least, ingenious. For what Dretske adduces is the condition on knowledge that we saw above specified in (DS), and it is this condition which produces the very consequences required. Take (A) as an example. According to (DS), in order for the agent to know (A), all that is required is that she have a true belief that (A) and, in the nearest possible world in which (A) is false, such as the possible world in which the animal she is looking at in the zebra enclosure is an ape, she does not believe that the animal before her is a zebra (as presumably she does not). In contrast, consider whether the agent knows the anti-sceptical proposition, (B*). Although her belief in this proposition may well be true in the actual world, in the nearest possible world in which this proposition is false (i.e., the sceptical not: (B*)-world), she still continues to believe that it is true because she is the victim of the wholesale, and undetectable, conspiracy to deceive noted in (B*). As a result, closure will fail on this view because the range of possible worlds at issue in determining knowledge of the antecedent proposition need not be the same as that which is at issue in determining knowledge of the consequent proposition. Using (DS), then, Dretske can thus explain why it is that we both know everyday propositions whilst failing to know the denials of sceptical hypotheses, and thus why, accordingly, closure fails.

It is important to remember, however, the dialectical ordering at play here. Dretske *first* notices (in Dretske (1970)), erroneously for an externalist as it turns out, that closure fails and *then* (in Dretske (1971)), he develops a modal theory of knowledge which would explain that fact.²⁷ Put the dialectic the other way around, however, and the position begins to lose its

attraction. For now we start with an independently plausible account of knowledge which entails that the relevant sceptical closure-type inferences must be legitimate and then determine that internalist closure must fail. Although in both cases we deny something in the face of scepticism, the dialectical gain is far greater in the latter case. For whereas Dretske is forced to argue that the sceptic's use of closure is unsound—despite our firm conviction that we cannot know everyday propositions without also knowing the denials of sceptical hypotheses—all we need argue is that the sceptic's use of a stronger sister principle is unsound. Hence, by running the dialectic in reverse Dretske finds himself endorsing an unnecessarily excessive form of epistemological revisionism in the face of radical scepticism.

Nevertheless, it may be that one has, in the meantime, become persuaded by the Dretskean account of knowledge and therefore feels that the order of the reasoning does not matter. On this view, (DS) provides at least good reason for thinking both that the sceptic's use of the closure principle is illegitimate, and that the sceptical premise is, unrestrictedly, true. Let me pause for a moment, then, to outline the manner in which the theory of knowledge offered here is superior to the Dretskean account.

To begin with, note that although the account in play here is less demanding than (DS) in sceptical contexts—since it does not demand that I need to consider far-off sceptical possible worlds in order to know that I am not a BIV—it is also in other respects more demanding. Knowing an everyday proposition, *E*, will, on this view, require more than merely being sensitive to the truth as to whether *E* in the nearest not: *E* world and the nearest *E* world. Rather, one must be sensitive to the whole *range* of worlds in the domain of sensitivity, and that can include a number of non-nearest not: *E* worlds and *E* worlds.

Moreover, the manner in which the account offered here is weaker than the Dretskean account should not unduly concern us. After all, the externalist thought is that if sceptical possible worlds are indeed counterfactually remote (which they must be if we are to know the antecedent propositions of the closure-type inferences in the first place), then surely they should not be able to influence our possession of knowledge. Rather, only actual circumstances and relevantly similar counterfactual possibilities should be relevant to knowledge. Indeed, this intuition drives both this response to scepticism and, surprisingly, that which is offered by Dretske and his followers. After all, Dretske has the same intuition when he argues that everyday knowledge should not be influenced by a failure to discount manifestly *irrelevant* sceptical error-possibilities. Where Dretske goes wrong is inferring from the fact that we can never have internalist knowledge that these error-possibilities have not obtained to concluding—in line with the unrestricted sceptical premise—that we just cannot know the

denials of sceptical hypotheses. This is the very inference that, I have argued, an externalist should resist. The move to allowing the sceptical premise—and thus to denying the sceptic’s closure-type inference—is thus inadequately supported.

Furthermore, another advantage of the view expounded here is that it both legitimates our intuition that if we know anything then we must know the denials of sceptical hypotheses whilst being duly respectful of the competing intuition that there is *something* amiss with such knowledge. For indeed, there *is* something amiss with such knowledge. For one thing, it lacks the kind of ‘internal’ epistemic properties that are desirable in any instance of knowledge. For another, it seems extremely plausible to suppose that the absence of such properties means that such knowledge could never be properly claimed to be possessed. Nevertheless, the key thought that if I am able to know mundane things such as whether or not I have two hands then I must be able to know the denials of sceptical hypotheses, such as the BIV hypothesis, remains intact. We thus have a promising diagnosis of both the attraction of radical scepticism and of the revisionist epistemological proposals designed to meet scepticism, such as the argument for non-closure. In both cases, they over-react to a lack of internalist knowledge, falsely concluding on this basis that either closure must go or the appearance of widespread knowledge possession is unreal. An externalist epistemology, properly understood, gives us the means to resist both these moves.

The account of knowledge sketched here is thus preferable to the Dretskean account both in its general epistemological ramifications and in its application to the radical sceptical paradox.²⁸

7. CONCLUDING REMARKS

What we have thus seen is that, properly understood, the key externalist intuition enables us to meet the sceptical argument head-on without forcing us to endorse the usual revisionistic proposals associated with an externalist epistemology. Indeed, it would appear that where the main anti-sceptical theses go wrong is in making important concessions to epistemological internalism at critical junctures—despite the fact that these theories are explicitly externalist in motivation.

One theory of knowledge that I have not accorded a detailed discussion to here, however, has been that of the other main revisionist anti-sceptical theory, contextualism. Recall that this (generally externalist) account of knowledge attempts to meet the sceptical argument

with closure intact by indexing knowledge to different epistemic standards. In line with the proposal sketched here, it thus allows that we can know the denials of radical sceptical hypotheses—that is, so long as we are working within a context that employs a sufficiently low epistemic standard. It should be easy to see that the theory put forward here goes a long way to undermining the attraction of the contextualist thesis. After all, if one can make sense of the possession of knowledge of the denials of radical sceptical hypotheses within the type of externalist epistemology that contextualists standardly accept, then why adopt the concessive stance of only allowing such knowledge relative to low epistemic standards? Indeed, why adopt the revisionist paraphernalia of contextualism at all if one can meet the sceptical argument head-on without such additions to one’s account?

This is certainly not a knock-down argument against contextualism, of course, but it does highlight a further line of inquiry for anyone impressed by an externalist proposal configured along these lines. This is, however, a topic for another day. For now it is enough that we have shown that epistemologically externalist responses to scepticism, properly understood, are perfectly able to retain the very epistemic principle that the sceptic uses to motivate her argument.²⁹

NOTES

¹ Although it is standard to express the sceptical argument in terms of knowledge, one possible challenge to this formulation may come from those who, like Wright (1985; 1991), believe that one should run the sceptical argument in terms of a weaker epistemic operator. I ignore this proposal here on the grounds that it is at least doubtful that the epistemic role of this weaker operator would be sufficiently *epistemic* to meet a sceptical challenge. If, as is commonly believed, justification is not sufficient, with true belief, for knowledge, then it would be problematic to argue that one could meet the sceptical challenge with justification alone. Moreover, if the operator were ‘beefed-up’ so that it could fulfil this criterion, then the question would then be why it should be necessary to run the argument in terms of this operator rather than knowledge. After all, it is the *epistemic* status of our beliefs that the sceptic is objecting to, and any operator that satisfied this test would encapsulate this element of knowledge in full. What gain would there then be in being agnostic about the obtaining of the truth condition? For a fuller discussion of these points, see Pritchard (2001c). In any case, as Williamson (1995) has suggested, there are good reasons for thinking that knowledge is a concept that cannot be ‘carved-up’ into constituent elements in the way traditionally supposed.

² In order to avoid unnecessary complications, the formulation of each epistemic principle in this paper will assume that the subject *does* have the requisite belief in the consequent proposition. Moreover, the reader should take every occurrence of the ‘K’ operator to be indexed to a particular agent and a particular time.

³ The exception will be those sceptical arguments, such as the Cartesian dreaming argument, that do not sever the ætiological link between the subject and her environment. Here the problem is that the entailment in question will not (usually) hold because dreaming does not preclude the subject from, for instance, sitting at a computer and typing a paper on scepticism. Accordingly, one cannot straightforwardly employ closure in such cases because one is unable to know that if one is, say, currently seated, then one is not dreaming. For more on dreaming scepticism, see Wright (1991) and Pritchard (2001e).

⁴ Of course, this is not all that there is to the Dretskean account. Dretske adds, for example, a reference to the “evidence-base” which produces the belief. Moreover, Nozick (1981) goes further to argue for an extra modal condition on knowledge. Nevertheless, these are all the details we need for the current discussion.

⁵ Stine (1976) is one commentator who has expressed the difference in this way. For some of the key recent treatments of the contextualist thesis, see Cohen (1999); Feldman (1999); Heller (1999b); Vogel (1999); and Pritchard (2001a). Heller (1999a) is one philosopher who has argued that we should deny closure *and* endorse contextualism. For a discussion of this proposal, see Pritchard (2000).

⁶ There has been some debate about whether closure really is as central to the radical sceptical debate as many suppose. See, for example, Yalçın (1992), Brueckner (1994), and Cohen (1998).

⁷ Note that sceptical possible worlds will not just be worlds in which a radical sceptical hypothesis is true, but also worlds in which such a hypothesis, though false, could very easily have been true. In a world in which human beings are regularly (and covertly) ‘harvested’ for their brains, an agent will not know very much, even if, as a matter of fact, she is not actually a BIV herself.

⁸ Of course, one way of making this element of the sceptical reasoning more explicit would be to express the argument in specifically modal terms. The conclusion of the sceptical argument would thus be something like the following: for any agent, time, or everyday proposition, it is *impossible* that the agent knows that proposition.

⁹ One finds such a view expressed in, for example, Chisholm (1989, chapter 1).

¹⁰ For the key discussions of the internalism/externalism distinction, see Kornblith (2001). For related discussions of how the internalist thesis is implicated in the sceptical paradox, see Williams (1991), Hill (1996), and Sosa (1999b).

¹¹ Both DeRose (1995) and Lewis (1996) offer externalist contextualist theses. One exception to this trend is Cohen (1999) who has argued for an internalist version of contextualism.

¹² For more on this, see Brandom (1998). Of course, this is not to deny that some externalists—such as, for example, Lewis (1996, 551)—do deny the relevance of an internal epistemic condition to knowledge out-right.

¹³ In a similar fashion, Sosa (1985) distinguishes between mere “animal” knowledge and “reflective” knowledge.

¹⁴ One should note, however, that if this rough characterisation of externalist knowledge is to be turned into a fully-fledged account then work needs to be done to insulate it from a number of straightforward counterexamples. For example, as it stands this account would treat cases where one reliably formed one’s true belief in a ‘lucky’ fashion—as when one forms a true belief regarding a burglar’s identity because one happens to look at his face just at the exact moment when his mask momentarily slips—as being cases where, contrary to intuition, one does *not* have knowledge (the example is adapted from one used by Nozick (1981)). After all, in the example just given there are a number of near-by possible worlds where the proposition at issue is true (that the burglar has such-and-such identity), but where the agent lacks the belief in this true proposition because he didn’t happen to look-up at the moment when the burglar’s mask slipped. There are a number of ways of getting around this problem (for example, one could, following Dretske (1971), restrict the range of relevant possible worlds to those where the “epistemic credential” that gave rise to the belief is the same as in the actual world), though I will spare the reader such an involved and technical discussion. After all, my primary aim in advancing this characterisation of externalism is to convince the reader of the plausibility of the proposal in general, not to present a detailed account of knowledge.

¹⁵ There are a number of versions of externalism currently being discussed in the literature that bear more than a passing resemblance to the ‘simple externalism’ described here. For example, the currently popular movement of ‘virtue’ epistemology is in the spirit of this sort of construal of the externalist position. Indeed, in line with the proposal sketched here, one of the key advocates of a position of this sort, Sosa (1991; 1999a; 2000; cf. Greco 1994), has argued in a number of recent papers that it follows from such a virtue theoretic treatment of knowledge that one should endorse a similar counterfactual condition on knowledge which he calls ‘safety’. Furthermore, Sosa also argues (as I do below) that it follows from such an analysis of knowledge that one can know the sort of propositions that figure in the sceptical premise. He is thus in a position to retain closure (at least as regards ‘sceptical’ inferences) on externalist grounds. Where the view advocated here diverges from Sosa’s, however, is in terms of the diagnosis that it offers for the failure of arguments for non-closure in terms of an ambiguous reading of the knowledge operator (and thus of closure itself). Relatedly, Sosa is not in a position to explain why the proponents of the Dretskean view are able to make their position look so persuasive. In this sense, Sosa fails to fully develop the case for the externalist defence of closure. I critically discuss the view that Sosa offers in more detail in Pritchard (2002). For more discussion of Sosa’s view in this respect, see Williamson (2000, chapter 7).

Interesting, Vogel (2000) also discusses a version of externalism that is similar to that advanced here (what he calls “neighbourhood reliabilism”), and he too notes that such a view ought to be able to avoid some of the more problematic counterexamples to closure entailed by the Dretskean view. Unlike Sosa and myself, however, Vogel ultimately rejects both this form of externalism and its Dretskean counterpart on internalist grounds.

¹⁶ After all, no one actually *endorses* sceptical doubt—radical scepticism is an intellectual *challenge*, not a philosophical *position*. Besides, we have already granted that the subject believes both the antecedent proposition and the relevant entailment, so it is difficult to see how she could lack a belief in the consequent proposition.

¹⁷ That this is so, however, does not itself ensure that closure-type inferences hold in other ‘non-sceptical’ contexts. This is because although it is part of the nature of a sceptical hypothesis that it must be barred from figuring either in or near-to the domain of sensitivity if the agent is to have knowledge of the everyday antecedent proposition, this need not be the case in other examples. Consider, for instance, a parallel case where the consequent proposition is not the anti-sceptical proposition (not: BIV), but rather another everyday proposition that is entailed by P; namely, Q: “I am not wearing my red shoes”. As before, insofar as our protagonist really does know P, then she will have a belief as to whether P is true which matches the facts as to whether P is true across the range of possible worlds included in the domain of sensitivity. The orderings of possible worlds will thus be set as before with sceptical worlds well beyond this domain. Significantly, however, and in contrast to the example involving not: BIV as the consequent proposition, it is perfectly consistent with this ordering that there be not-Q worlds within the domain of sensitivity. As a result, we cannot ensure that closure holds purely on the basis that the agent has a belief in Q which cannot help but match the truth as to whether Q in all the possible Q worlds in the domain of sensitivity. Nevertheless, this is not to say that closure will fail in such cases. It is, after all, difficult to imagine how one could have a sensitive belief in P without a corresponding sensitive belief in Q. What, for instance, could get in the way of an agent’s sensitivity as regards Q which does not likewise affect her sensitivity to P? We do not need to get into such issues here, however, since it suffices for our purposes that closure works perfectly well in the cases that the sceptic uses. Given this characterisation of sensitivity, the sceptic is entirely within her rights to use closure for sceptical ends.

¹⁸ In Pritchard (2001b), I outline how one could construe the distinction made by Davies (1998) and Wright (2000) between closure and what they call ‘transmission’ in terms of this contrast between (CK) and (CKI).

¹⁹ That said, one can detect a train of argument against the closure principle—or, at least, the sister principle (CKI)—in earlier work by Austin (1961) and Wittgenstein (1969).

²⁰ Dretske’s point is actually directed at epistemic operators *in general*, but I will be focussing on his remarks as they relate to knowledge in particular. For a list of the epistemic operators that he has in mind, see Dretske (1970, 1009).

²¹ Dretske (1970, 1009) refers to epistemic operators as being “semi-penetrating”.

²² The point is often expressed in terms of ‘relevant alternatives’ (e.g. Dretske 1970), where an ‘alternative’ to P is a proposition which is inconsistent with the truth of P (and thus a proposition which will be false if P is true). If Q is a ‘relevant’ alternative to P then one must be able to eliminate the possibility of Q in order to know P. Accordingly, for Dretske, not all alternatives to P are relevant alternatives - one can know that the animals in the pen are zebras without having to rule-out the irrelevant alternative that they are cleverly disguised mules. For discussion, see Stine (1976); Yourgrau (1983); Heller (1989); Cohen (1991; cf. Dretske 1991); and DeRose (1995; 1996).

²³ Dretske is actually noticeably coy about the epistemic status of the subject’s belief in the entailment. If Dretske is to be seen as making a substantive point, however, then it ought to be possible to run his argument, as I have done here, with knowledge of the entailment taken for granted.

²⁴ Of course, the ‘assumption’ in question here need only be tacit—it is no part of Dretske’s argument that such assumptions are explicitly made. Indeed what Dretske actually notes is that we tacitly assume some “general uniformities” which are given expression to by remarks such as “Why should the zoo authorities do that?”. It is not the denial of the sceptical error-possibility that is explicitly assumed then, so much as some general conception of one’s epistemic position which entails the truth of this tacit assumption. I am grateful to a referee from *Theoria* for pressing me on this point.

²⁵ Part of the problem here is that Dretske tends to run discussions on closure together with his conception of “contrast-consequences”. This latter phenomenon concerns how the same assertion in different conversational contexts can carry different contrast-consequences. For instance, in a conversational context in which the participants take for granted that there is no deception going on, then it may be legitimate to *claim* knowledge of (A), but not of (B), because a claim to know (A) will not carry the conversational implicature that one can rule-out the contrast-consequent embedded in (B). As far as closure goes, however, this is entirely besides the point because the issue here ought not to be dependent upon pragmatic features of the situation but on the hard-and-fast epistemic standards in operation. Indeed, as both Yourgrau (1983) and Sanford (1991) have pointed out (with Dretske’s agreement (Dretske 1991)), such examples do not seem to be cases where one knows the (A)-type proposition but not the (B)-type proposition, but rather cases in which the agent does not know the (A)-type proposition but may nevertheless legitimately claim to know the (A)-type proposition because of the restricted parameters of that conversational context (in effect, the assertion has the form, “Given that certain (B)-type conditions obtain, I know the (A)-type proposition”). For more discussion on this point, see Dretske (1968; 1970); Yourgrau (1983); Sanford (1991; cf. Dretske 1991); and Rieber (1998).

²⁶ For an account of the propriety of knowledge claims that would support this line, see Wittgenstein (1969) and Pritchard (2001d).

²⁷ Of course, Dretske has a relevant alternatives model prior to 1971, and this can be used to account for the failure of closure. Crucially, however, the basic relevant alternatives idea can be cashed-out in a number of ways, some of them consistent with closure (for more on this point, see, for example, Stine (1976)). Accordingly, it is not the relevant alternatives idea itself that accounts for the failure of closure on the Dretskean view, but rather the specific modal interpretation that he offers of this idea in 1971. I am grateful to a referee from *Theoria* for pressing me on this point.

²⁸ For another influential critique of Dretske's argument for non-Closure—and one which is, in the main, compatible with the line taken here—see Williams (1991, chapter 8).

²⁹ Acknowledgements plus thanks to the two referees from *Theoria*.

REFERENCES

- Austin, J. L. (1961). 'Other Minds', reprinted in his *Philosophical Papers*, (eds.) J. O. Urmson & G. J. Warnock, 44-84, Clarendon Press, Oxford, England.
- Brandom, R. (1998). 'Insights and Blindspots of Reliabilism', *The Monist* 81, 371-92.
- Brueckner, A. (1994). 'The Structure of the Skeptical Argument', *Philosophy and Phenomenological Research* 54, 827-35.
- Chisholm, R. M. (1989). *Theory of Knowledge*, (3rd Edition), Prentice Hall, Englewood Cliffs, New Jersey.
- Cohen, S. (1986). 'Knowledge and Context', *Journal of Philosophy* 83, 574-83.
- Cohen, S. (1991). 'Skepticism, Relevance, and Relativity', in *Dretske and his Critics*, (ed.) B. McLaughlin, 17-37, Basil Blackwell, Oxford, England.
- Cohen, S. (1998). 'Two Kinds of Skeptical Argument', *Philosophy and Phenomenological Research* 58, 143-59.
- Cohen, S. (1999). 'Contextualism, Scepticism, and the Structure of Reasons', *Philosophical Perspectives* 13, 57-90.
- Davies, M. (1998). 'Externalism, Architecturalism, and Epistemic Warrant', in *Knowing Our Own Minds: Essays on Self-Knowledge*, (eds.) C. Macdonald, B. Smith & C. J. G. Wright, 321-61, Oxford University Press, Oxford, England.
- DeRose, K. (1995). 'Solving the Sceptical Puzzle', *Philosophical Review* 104, 1-52.
- DeRose, K. (1996). 'Relevant Alternatives and the Content of Knowledge Attributions', *Philosophy and Phenomenological Research* 56, 193-7.
- Dretske, F. (1968). *Seeing and Knowing*, Routledge & Kegan Paul, London, England.
- Dretske, F. (1970). 'Epistemic Operators', *Journal of Philosophy* 67, 1007-23.
- Dretske, F. (1971). 'Conclusive Reasons', *Australasian Journal of Philosophy* 49, 1-22.
- Dretske, F. (1991). 'Knowledge: Sanford and Cohen', in *Dretske and his Critics*, (ed.) B. McLaughlin, 185-96, Basil Blackwell, Oxford, England.
- Feldman, R. (1999). 'Contextualism and Scepticism', *Philosophical Perspectives* 13, 91-114.
- Goldman, A. (1986). *Epistemology and Cognition*, Harvard University Press, Cambridge, Massachusetts.
- Greco, J. (1994). 'Virtue Epistemology and the Relevant Sense of "Relevant Possibility"', *The Southern Journal of Philosophy* 32, 61-77.
- Heller, M. (1989). 'Relevant Alternatives', *Philosophical Studies* 55, 23-40.
- Heller, M. (1999a). 'Relevant Alternatives and Closure', *Australasian Journal of Philosophy* 77, 196-208.
- Heller, M. (1999b). 'The Proper Role for Contextualism in an Anti-Luck Epistemology', *Philosophical Perspectives* 13, 115-29.
- Hill, C. (1996). 'Process Reliabilism and Cartesian Scepticism', *Philosophy and Phenomenological Research* 56, 567-81.

- Klein, P. (1987). 'On Behalf of the Skeptic', in *The Possibility of Knowledge*, S. Luper-Foy (ed.), 267-81, Rowman & Littlefield, Totowa, New Jersey.
- Kornblith, H. (2001). *Epistemology: Internalism and Externalism*, (ed.) Basil Blackwell, Oxford, England.
- Lewis, D. (1996). 'Elusive Knowledge', *Australasian Journal of Philosophy* 74, 549-67.
- Nozick, R. (1981). *Philosophical Explanations*, Oxford University Press, Oxford, England.
- Pollock, J. (1986). *Contemporary Theories of Knowledge*, Rowman & Littlefield, Totowa, New Jersey.
- Pritchard, D. H. (2000). 'Closure and Context', *Australasian Journal of Philosophy* 78, 275-80.
- Pritchard, D. H. (2001a). 'Contextualism, Scepticism, and the Problem of Epistemic Descent', *Dialectica* 55, 327-49.
- Pritchard, D. H. (2001b). 'McKinsey Paradoxes, Radical Scepticism, and the Transmission of Knowledge across Known Entailments', *Synthese* 129, 1-24.
- Pritchard, D. H. (2001c). 'Meta-Epistemological Constraints on Anti-Sceptical Theories', *Facta Philosophica* 3, 101-26.
- Pritchard, D. H. (2001d). 'Radical Scepticism, Epistemological Externalism, and "Hinge" Propositions', *Wittgenstein-Studien*, 81-105
- Pritchard, D. H. (2001e). 'Scepticism and Dreaming', *Philosophia* 28, 373-90.
- Pritchard, D. (2002). 'Resurrecting the Moorean Response to Scepticism', *International Journal of Philosophical Studies* 11, 1-25.
- Pryor, J. (2001). 'Epistemology c. 1988-2000', *The British Journal for the Philosophy of Science* 52, 95-124.
- Putnam, H. (1981). *Reason, Truth and History*, Cambridge University Press, Cambridge, England.
- Rieber, S. (1998). 'Skepticism and Contrastive Explanation', *Noûs* 32, 189-204.
- Sanford, D. H. (1991). 'Proper Knowledge', *Dretske and his Critics*, (ed) B. P. McLaughlin, 38-51, Basil Blackwell, Oxford, England.
- Sosa, E. (1985). 'Knowledge and Intellectual Virtue', *The Monist* 68, 224-45.
- Sosa, E. (1991). 'Intellectual Virtue in Perspective', in his *Knowledge in Perspective: Selected Essays in Epistemology*, 270-93, Cambridge University Press, Cambridge, England.
- Sosa, E. (1999a). 'How to Defeat Opposition to Moore', *Philosophical Perspectives* 13, 141-54.
- Sosa, E. (1999b). 'Skepticism and the Internal/External Divide', in *Epistemology*, J. Greco & E. Sosa (eds.), 145-157, Basil Blackwell, Oxford, England.
- Sosa, E. (2000). 'Skepticism and Contextualism', *Philosophical Issues* 10, 1-18.
- Stine, G. C. (1976). 'Scepticism, Relevant Alternatives, and Deductive Closure', *Philosophical Studies* 29, 249-61.
- Vogel, J. (1999). 'The New Relevant Alternatives Theory', *Philosophical Perspectives* 13, 155-80.
- Vogel, J. (2000). 'Reliabilism Levelled', *Journal of Philosophy* 97, 102-24.
- Williams, M. (1991). *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*, Basil Blackwell, Oxford, England.
- Williamson, T. (1995). 'Is Knowing a State of Mind?', *Mind* 104, 533-65.
- Williamson, T. (2000). *Knowledge and Its Limits*, Oxford University Press, Oxford, England.
- Wittgenstein, L. (1969). *On Certainty*, (eds.) G. E. M. Anscombe & G. H. von Wright, (tr.) Denis Paul & G. E. M. Anscombe, Basil Blackwell, Oxford, England.
- Wright, C. (1985). 'Facts and Certainty', *Proceedings of the British Academy* 71, 429-72.
- Wright, C. (1991). 'Scepticism and Dreaming: Imploding the Demon', *Mind* 100, 87-115.
- Wright, C. (2000). 'Cogency and Question-Begging: Some Reflections on McKinsey's

- Paradox and Putnam's Proof', *Philosophical Issues* 10, 140-63.
- Yalçın, Ü. (1992). 'Skeptical Arguments from Underdetermination', *Philosophical Studies* 68, 1-34.
- Yourgrau, P. (1983). 'Knowledge and Relevant Alternatives', *Synthese* 55, 55-70.