

SCEPTICISM, EPISTEMIC LUCK AND EPISTEMIC *ANGST*

Duncan Pritchard

ABSTRACT. A commonly expressed worry in the contemporary literature on the problem of epistemological scepticism is that there is something deeply intellectually unsatisfying about the dominant anti-sceptical theories. In this paper I outline the main approaches to scepticism and argue that they each fail to capture what is essential to the sceptical challenge because they fail to fully understand the role that the problem of epistemic luck plays in that challenge. I further argue that scepticism is best thought of not as a quandary directed at our possession of knowledge *simpliciter*, but rather as concerned with a specific kind of knowledge that is epistemically desirable. On this view, the source of scepticism lies in a peculiarly epistemic form of *angst*.

"It is always by favour of Nature that one knows something."
Wittgenstein [1969: §505]

1. THE SCEPTICAL PARADOX IN CONTEMPORARY DEBATE

The dominant conception of the problem of epistemological scepticism in the current literature understands it in terms of an argument from ignorance that is focussed upon the possibility of knowledge of a wide class of empirical propositions. Essentially, the sceptical claim on this view is that since we are unable to know that we are not the victims of any one of a number of radical sceptical hypotheses—such as the ‘BIV’ hypothesis that one is presently a brain-in-a-vat being ‘fed’ one’s experiences by neuroscientists—we are unable to have knowledge of any one of a wide class of ‘everyday’ empirical propositions that we ordinarily take ourselves to know, such as that we have two hands. Taking ‘SH’ to be a representative sceptical hypothesis, and ‘E’ to be a representative everyday empirical proposition, we thus get the following template formulation of the sceptical problem:

The Template Sceptical Argument

- (1) S does not know that \neg SH.
- (2) If S does not know that \neg SH, then S does not know E.
- (C) S does not know E.

As the sceptic points out, both of the premises of this argument are highly plausible. Accordingly, since this argument is valid and can be repeated with just about any everyday empirical proposition, so we get the paradoxical conclusion that we know next to nothing of what we think we know. Indeed, strictly speaking, the sceptical claim here is the stronger modalized conclusion that it is *impossible* for us to have this everyday knowledge since both of the premises are meant to be necessarily true (or at least, in the case of the first premise, necessarily true when it comes to creatures like us). One can find formulations of the sceptical problem in essentially these terms in numerous places in the contemporary

literature.¹

What motivates the first premise is the widespread intuition that sceptical hypotheses are, by their nature, error-possibilities that one is unable to know to be false. Part of the problem here is that one's evidence for thinking that they are false is, it seems, of its nature inadequate to the task. Since being a BIV being 'fed' one's experiences as if everything is normal is phenomenologically indistinguishable from everything being normal, it is hard to see how one could ever have sufficient empirical grounds to support one's knowledge that one is not the victim of such a scenario (we will return to consider this claim in more detail below).

The motivation for the second premise is more indirect, and typically concerns the so-called 'closure' principle for knowledge.² Roughly, this principle can be formulated as follows:

Closure for Knowledge

For all S, ϕ , ψ , if S knows ϕ , and S knows that ϕ entails ψ , then S knows ψ .

Such a principle seems, on the face of it at least, entirely uncontroversial. If I know that the murder was committed in the pantry, and that if the murder was committed in the pantry then it wasn't committed in the bedroom, then intuitively I must also know that it wasn't committed in the bedroom. Crucially, however, with this principle in play the second premise of the sceptical argument follows immediately, at least where the E-type proposition is known to be inconsistent with the sceptical hypothesis under consideration, which is what we would typically expect. Consider the following BIV version of the template sceptical argument, for example, where E is an everyday empirical proposition that is inconsistent with this sceptical hypothesis, such as that one has two hands:

The BIV Sceptical Argument

- (1) S does not know that \neg BIV.
- (2) If S does not know that \neg BIV, then S does not know E.
- (C) S does not know E.

Given closure for knowledge, if one is unable to know that \neg BIV then, granted that one has knowledge of the (obvious) entailment from E to \neg BIV, it follows that one does not have knowledge of E either, since if one did then this would mean that one would also know that one was not a BIV.

With the sceptical paradox understood in this way, three types of anti-sceptical response suggest themselves. The first—the '*Dretskean*' proposal—is simply to deny closure, and therefore deny the second premise of the sceptical argument. Such a move is typically associated with the work of Fred Dretske [1970] and Robert Nozick [1981]. The second—the '*contextualist*' proposal—is to argue for an attributer contextualist account of knowledge which maintains that the term 'knows' is context-sensitive in such a way as to enable one to claim that whilst in some epistemically demanding contexts of ascription the sceptical argument holds, in the sense that agents neither know the denials of sceptical

hypotheses nor the relevant everyday propositions, in other less epistemically demanding contexts it fails, in that agents know both the everyday propositions and the denials of sceptical hypotheses. One thus gets a victory over the sceptic (at least in normal contexts of ascription that are not epistemically demanding) without having to deny closure. Such an account is typically associated with the work of Keith DeRose [e.g., DeRose 1995] and David Lewis [1996].³ Finally, there is the so-called ‘*neo-Moorean*’ strategy of allowing closure whilst rejecting attributer contextualism which essentially consists of the claim (in opposition to the first premise of the sceptical argument) that agents *do* know the denials of sceptical hypotheses after all (at least insofar as they know anything much at all at any rate). Such a view is often attributed to Ernest Sosa [e.g., 1999], and I have defended a version of it myself [e.g., Pritchard 2002d].

Each of these anti-sceptical strategies has their problems and, as we will see, these difficulties ultimately relate to their commitment, usually explicit but sometimes implicit, to some version of epistemological externalism. Call the condition or conditions that turn true belief into knowledge ‘epistemic conditions’. What makes an epistemic condition an internal epistemic condition is that it incorporates some sort of ‘access’ requirement whereby the facts that determine that this condition has obtained are facts that the agent has special access to. Typically, this will mean that the agent is able to know these facts by reflection alone, where by ‘reflective knowledge’ one usually means *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.⁴ The justification condition is traditionally understood along internalist lines in that the kinds of facts that determine this condition—such as facts about what grounds the agent is able to offer in favour of her belief—are facts that the agent has (it is thought) special reflective access to. What demarcates an epistemologically internalist theory of knowledge from an externalist theory of knowledge is that the former, but not the latter, makes the meeting of an internal epistemic condition a *necessary* condition for knowledge.

In order to see this, consider the way in which the debate between internalists and externalists about knowledge tends to focus upon certain kinds of examples. Take, for example, the famous case of the ‘chicken sexer’. This, putatively, is an agent who, by being raised around chickens, acquires the highly reliable ability to tell male and female chicks apart. Crucially, however, the chicken sexer usually has false beliefs about how she is forming her beliefs, supposing that they are due to the fact that she is seeing or touching something distinctive even though, tests show, there is nothing distinctive for her to see or touch (she’s actually doing it via her sense of smell). What should we say about such an agent who reliably forms true beliefs in this fashion but who has false beliefs about how she is forming her beliefs and who also, let us stipulate for good measure, lacks grounds for thinking that she is reliable in this respect? If one’s intuition in this regard tends towards the ascription of knowledge on

the grounds that she really is exercising a reliable cognitive disposition, then one will be inclined to side with the epistemological externalist. If, on the other hand, one's intuition is that the agent does not have knowledge because she fails to have good reflectively accessible supporting reasons for her belief (and thus fails to meet a relevant internal epistemic condition), then one will tend to side with the epistemological internalist.⁵

The obvious way in which epistemological internalism about knowledge is implicated in the sceptical argument under consideration is with the first premise, since its clearest motivation comes from this quarter. Whatever else one might say about the epistemic status of one's beliefs in the denials of sceptical hypotheses, they certainly don't seem to meet internal epistemic conditions at all since in terms of what we can know by reflection alone we appear to have no good grounds at all for thinking that we are not BIVs and so forth.⁶ That is, more precisely, whilst we might *seem* to have good grounds for this belief (we can, for example, see and interact with what appears to be a genuine non-envatted environment), whether or not these grounds are indeed the good grounds that they appear to be depends upon facts obtaining which we are not able to reflectively know to be true. With this claim in place, and coupled with our obvious attachment to the denial of the sceptical conclusion—that we can know lots of everyday propositions—the tension with closure encapsulated in the second premise of the sceptical argument is clear.

The problem with the Dretskean proposal that we should reject closure lies in the diverse motivation that is cited in favour of making such a move. On the one hand, the natural informal motivation for this denial is the 'relevant alternatives' intuition that some error-possibilities, such as the sceptical error possibilities encapsulated in sceptical hypotheses, are just not relevant to our knowledge of everyday propositions. Of course, motivating such an anti-sceptical line within the internalist conception of knowledge is going to be highly suspect for the simple reason that to allow that we are in a position to make such convenient discriminations concerning which error-possibilities are relevant is going to look like a retreat to subjectivity. Intuitively, and taking it as given that there is no adequate *a priori* refutation of scepticism available, one has no reflectively accessible grounds for thinking that sceptical error-possibilities are irrelevant to one's everyday knowledge. For one thing, for all we can know by reflection alone the sceptical error possibility could be *true*, and in this case clearly it *would* be relevant to our knowledge (indeed, its truth would ensure that we know very little, if anything, of what we thought we knew). The only reflectively accessible grounds that one has for thinking that sceptical error-possibilities are irrelevant are pragmatic ones (such as that allowing their relevance would seem to undermine one's knowledge of everyday propositions). If one is to make any anti-sceptical headway with the relevant alternatives intuition one must thus develop the intuition along epistemically externalist lines which allow the discrimination to be made on external, and thus non-

reflectively accessible, grounds.

The way the Dretskean does this is by conceiving of relevance in modal terms via the principle of sensitivity. This demands that an agent's belief be such that in the nearest possible world or worlds in which what the agent believes is false the agent no longer forms that belief via the same method as that employed in the actual world. And since the ordering of the possible worlds is dictated by their similarity to the actual world, the relevant class of worlds will be determined by the facts concerning the nature of the actual world, facts which the agent has no special access to. All parties to the sceptical dispute, both sceptical and anti-sceptical, grant that agents lack knowledge if they form beliefs about the external world in sceptical worlds (or worlds near to sceptical worlds) where sceptical hypotheses are true. The issue thus rests upon what happens if agents form their external world beliefs in worlds that are like the world as we take it to be. The sceptic will contend that we still lack knowledge even under this supposition because we cannot, via the grounds that are reflectively accessible to us, rule-out sceptical hypotheses, but the Dretskean has a response to this move. On this understanding of the actual world and the resultant ordering of possible worlds, one's beliefs in everyday propositions can be sensitive, and so candidates for knowledge, even though one's beliefs in the denials of sceptical hypotheses are not sensitive and so not candidates for knowledge. This is because the worlds that are relevant to whether or not one's everyday beliefs are sensitive will be different from the worlds that are relevant to whether or not one's anti-sceptical beliefs are sensitive. In the former case, it will be relatively near-by possible worlds that are at issue, whilst in the latter case it will be, by hypothesis, far-off sceptical possible worlds that are at issue.

For example, if the actual world is pretty much as one takes it to be, then one's belief that one has two hands will be sensitive in that, in the nearest possible worlds where this is false (where I have had the misfortune to have lost my hands and so am looking at two stumps at the ends of my arms right now), I no longer believe that I have them. Crucially, however, even with the same supposition about the actual world in play, one's belief that one is not, say, a BIV is *insensitive*, in that in the nearest possible worlds where this is false (where one is a BIV), one will continue to believe that one isn't. The Dretskean thus seems to have identified an external criterion for relevance that supports an anti-sceptical argument to the effect that closure, and thus the second sceptical premise, should be rejected.

Matters are not quite as simple as they at first appear, however. The reason for this is that, as we saw above, the primary motivation for regarding us as lacking knowledge of the denials of sceptical hypotheses comes from an epistemologically internalist theory of knowledge, but if one is committed, with the Dretskean, to an externalist thesis, then it is far from clear why one should accept this particular argument for non-closure. This problem is exacerbated once one considers that on the most obvious reading of the relevant alternatives thesis, far-off possible worlds, such as sceptical possible

worlds, are irrelevant to *all* knowledge—this is just what it means for an error-possibility to be irrelevant. The Dretskean cannot make use of this particular externalist criterion of what counts as a relevant alternative, however, for the simple reason that on this view there is now no obvious reason why closure should fail at all. If the actual world is pretty much as we take it to be, then just as one's beliefs in everyday propositions will track the truth across near-by possible worlds, so will one's beliefs in the denials of sceptical hypotheses (we will return to this point in a moment). It seems then that the Dretskean is eschewing a natural epistemically externalist rendering of the relevant alternatives thesis in order to opt for a different externalist account which explains why closure fails, but now the only reason we have for thinking that closure fails is the epistemologically internalist intuition that we are unable to know the denials of sceptical hypotheses. As an epistemological externalist, however, the Dretskean ought to be suspicious of this intuition. The dilemma for the Dretskean is thus between, on the one hand, adopting a consistent epistemological externalism, in which case he needs to offer further reasons for thinking that we are unable to know the denials of sceptical hypotheses; or else, on the other hand, adopting a consistent epistemological internalism, in which case it isn't at all clear that he can help himself to the relevant alternatives intuition in order to explain the failure of closure in the first place.⁷

It is also the commitment to epistemological externalism that creates problems for the attributer contextualist response to scepticism. Clearly, attributer contextualists need to be epistemological externalists because they allow that agents are able to know the denials of sceptical hypotheses, albeit only in everyday contexts that employ low epistemic standards. Since one cannot make sense of such knowledge along internalist lines, epistemic externalism is the only option here.⁸ Granted this, however, the obvious question that arises is just why we should endorse attributer contextualism given that this move alone will suffice to meet the template sceptical argument outlined above by blocking the opening premise. That is, if one is willing to grant that one can indeed have knowledge of the denials of sceptical hypotheses, in *any* sense, then this alone will suffice for meeting the sceptical problem under consideration. There seems no need to saddle this seemingly counterintuitive move with the highly contentious attributer contextualist theoretical machinery.

Of course, the attributer contextualist will argue that we need the contextualist theoretical machinery in order to explain why we were ever taken in by the sceptical argument in the first place, but a more immediate explanation of this fact now presents itself—*viz.*, our implicit commitment to epistemological internalism about knowledge. That is, given this conception of the sceptical problem and given also the legitimacy of advancing an epistemological externalist theory in response to this problem in the manner of both the Dretskean and contextualist proposals, the more straightforward response to the sceptic that suggests itself is of a Moorean form that holds that insofar as one knows

everyday propositions then, given closure, one must also know the denials of sceptical hypotheses. This position is ‘Moorean’ because it mirrors, in broad outline if not in detail, the ‘common sense’ response to the sceptic offered by G. E. Moore [1925; 1939]. Where it differs from Moore is in offering (amongst other things) an account of knowledge that explains how we could come to know the denials of sceptical hypotheses.

We have already seen the basics of such an account in terms of a natural rendering of the relevant alternatives intuition which regards the relevant possible worlds as simply being the near-by possible worlds rather than the variable set of possible worlds dictated by the sensitivity condition (and much less the variable context-sensitive set of error-possibilities, and thus possible worlds, at issue in attributer contextualism). Typically, one finds this intuition encapsulated in terms of the ‘safety’ condition that (roughly) one’s belief should match the truth in most (if not all) of those possible worlds which are near-to the actual world in which one forms one’s belief in the same way as the actual world, such that where the agent believes the proposition in question, then it is true. On this conception of knowledge, the epistemic status of one’s beliefs in everyday propositions will tend to ebb and flow in concert with the epistemic status of one’s beliefs in the denials of sceptical hypotheses, just as one would expect (and as closure predicts). In particular, as long as the actual world is pretty much as one takes it to be, then in most (if not all) near-by possible worlds in which one believes that one has two hands by, say, looking at them, one will form a true belief as a result. Similarly, since on this supposition BIV-worlds are far-off worlds, in all near-by possible worlds in which one believes that one is not a BIV, one’s belief will also be true. Both one’s beliefs in everyday propositions and one’s beliefs in the denials of sceptical hypotheses can thus be safe, and hence insofar as one knows anything much at all, one also knows the denials of sceptical hypotheses. Accordingly, there is no anti-sceptical need, on this view, for the denial of closure, much less for the endorsement of an attributer contextualist thesis.⁹

All-in-all, then, *if* one understands the sceptical paradox in terms of the template sceptical argument outlined above, and *if* one thinks that epistemological externalism can be a legitimate component of an anti-sceptical thesis, then the neo-Moorean response seems to have the theoretical edge. Nevertheless, this does not mean that such a response to scepticism is altogether intellectually compelling, even if it does have advantages over its two key competitor anti-sceptical theories, and the crux of the matter is to understand why.

The problem seems to lie in the move to epistemological externalism, and it is this move that has prompted the so-called ‘metaepistemological’ scepticism that one finds in such writers as Barry Stroud. He writes:

[... S]uppose there are truths about the world and the human condition which link human perceptual states and cognitive mechanisms with further states of knowledge and reasonable belief, and which imply that human beings acquire their beliefs about the physical world through the operation of belief-forming mechanisms which are on the whole reliable in the sense of giving them mostly true beliefs. [...] If there are truths of this kind [...] that fact alone obviously will do us no good as theorists who want to understand human knowledge in this philosophical way. At the very least we must believe some such truths; their merely being true would not be enough to give us any illumination or satisfaction. But our merely happening to believe them would not be enough either. We seek understanding of certain aspects of the human condition, so we seek more than just a set of beliefs about it; we want to know or have good reasons for thinking that what we believe about it is true. [Stroud 1994: 297]¹⁰

That is, it does not seem to be enough to merely show that knowledge is possible if part of what this consists in includes the concession that whilst, if certain external (and anti-sceptical) conditions obtain, we have knowledge, we lack good reflectively accessible grounds for thinking that those conditions do obtain. This result seems to cast our epistemic situation as being one of mere bad faith. As Richard Fumerton puts the point:

[... E]xternalist analyses of knowledge [...] simply remove one *level* of the traditional problems of skepticism. When one reads the well-known externalists one is surely inclined to wonder why they are so sanguine about their supposition that our commonplace beliefs are, for the most part, [...] knowledge. [...] Perception, memory, and induction *may* be reliable processes (in Goldman's sense) and thus given his metaepistemological position we may [... *have knowledge of*] the beliefs they produce but, the sceptic can argue, we have no reason to believe that these processes are reliable and thus even if we accept reliabilism, we have no reason to think that the beliefs they produce [*constitute knowledge*]. [Fumerton 1990: 63]¹¹

That is, we proceed on the assumption of knowledge and with the assurance that such knowledge is at least *conditionally* possible, though since this knowledge is conditionally possible only given that which we have no good reason for thinking is true, we also proceed without the subjective assurance that we originally sought and on the basis of which the sceptic motivated her initial worry in the first place.

There is certainly something in this complaint, but I don't think that it can be the entire story as to where these anti-sceptical responses go wrong. After all, the conclusion that knowledge of much of what we think we know is possible *is* a legitimate response to the sceptical argument under consideration above. Moreover, the claim that Stroud and Fumerton are making cannot simply be that we should be internalists about knowledge and *therefore* seek an internalist resolution of the sceptical problem since there are independent grounds in favour of epistemological externalism which surely must suffice to ensure that, *prima facie* at least, externalism can be applicable to the sceptical problem. (Moreover, there are also *prima facie* grounds for thinking that no epistemically internalist resolution to the problem is going to be possible, which is partly what prompts the move to externalism in the first place). Accordingly, if we are to get a handle on the problem here, we are going to have to look a little deeper into what is motivating the kind of metaepistemological worry that Stroud and Fumerton are attempting to give expression to.

2. EPISTEMIC LUCK

One of the key intuitions that drives epistemological discussion is the thought that knowledge is in some sense an achievement on the part of the agent, something that the agent can take credit for. A related intuition in this respect is that whatever knowledge is, it cannot be due to luck since luck undermines genuine achievement. It is, for example, an intuition of this sort that guides discussion of the Gettier-style counterexamples to the classical tripartite account of knowledge. As Jonathan Dancy casually puts the point:

[... *J*]ustification and knowledge must somehow not depend on coincidence or luck. This was just the point of the Gettier counter-examples; nothing in the tripartite definition excluded knowledge by luck. [Dancy 1985: 134]

An agent who, for example, gains a true belief about what the time is by looking at a stopped clock, even if that belief is fully internalistically justified in the way that the classical tripartite account of knowledge would demand, still lacks knowledge because it is simply a matter of luck that her belief is true (that the stopped clock happened to be ‘telling’ the right time) and this undermines the sense in which the truth of her belief can be considered a genuine cognitive achievement on her part.¹²

One of the morals of the Gettier-type examples is that meeting merely internal epistemic conditions will not, with true belief, suffice for knowledge. The reason for this is that no matter how good these internal grounds might be, they will never suffice to ensure that one’s belief has not been ‘Gettiered’ in such a way as to make the truth of the belief in question due, to some substantive degree, to luck. In the case just cited, for example, no matter how good one’s reflectively accessible grounds might be for thinking that the clock is working, such grounds are consistent with the clock in fact having stopped and thus with one only luckily gaining a true belief about what the time is from this clock. An external epistemic condition is thus needed to ‘de-Gettier’ one’s true belief.

One can achieve this by adducing either of the sensitivity or safety principles introduced above. In the case of safety, this will meet the problem by ensuring that such cases do not count as knowledge because there will be a class of near-by possible worlds in which one forms the belief in the same way as in the actual world and yet forms a false belief as a result (e.g., the world in which the clock has still stopped, but the time is different). In the case of the sensitivity principle, the problem cases will be met because there will be a class of near-by possible worlds in which what is believed is false and yet one continues to form, on the same basis, the belief that one did in the actual world (as before, this class will include the world in which the clock has still stopped and yet the time is different). Indeed, sensitivity is clearly the stronger of the two principles in this respect, since whereas safety only deals with near-by possible worlds sensitivity will count far-off possible worlds as being relevant to

knowledge as well, at least in those cases where the negation of the proposition believed is only true in far-off possible worlds. Crucially, however, one only needs to consider near-by possible worlds in order to deal with the challenge posed by the Gettier-style examples, since they always concern local error-possibilities that have obtained but obtained in such a way as to not, as it happens, undermine the truth of the agent's belief or her meeting of the (internal) justification condition in the actual world. Accordingly, the world in which the agent forms her belief in the same way as in the actual world but forms a false belief as a result will, perforce, always be a near-by possible world. Given that this is the case, however, one only needs to adopt a safety-style principle to meet Gettier-type examples.

That one needs a safety-type principle in order to deal with Gettier-type examples does not itself prejudge the issue of whether we should be externalists or internalists about knowledge. As we noted above, the distinction rests upon whether one regards the meeting of an internal epistemic condition as being necessary for knowledge and so one could consistently advocate, say, an internalist safety-based theory of knowledge simply by demanding that agents need to meet a further internal epistemic condition over and above the safety condition. The issue is thus whether we should add the extra demand that the agent must meet an internal epistemic condition and this takes us back to such cases as the chicken sexer example. Clearly, the beliefs formed by this chicken sexer are safe in the required sense since in all near-by possible worlds where she forms her belief in the same manner as in the actual world her belief will track the truth in the required way, but the question will be whether this alone is enough for knowledge.¹³

I think that the way to see our way through this debate, and in the process throw some light on the issue raised above about scepticism, is to try to understand just why the safety principle is able to directly meet the challenge posed by the Gettier-style example. The answer lies, I claim, in how the safety principle rules-out the specific kind of luck that is at issue in Gettier-style examples.

What makes an event lucky? There is surprisingly little in the philosophical literature devoted to this question, and what work has been done on this issue has tended to be merely suggestive, but we can, I think, capture the heart of the notion by understanding luck in modal terms roughly as follows:

Luck

An event is lucky if it obtains in the actual world but does not obtain in most of the nearest possible worlds where the relevant initial conditions for that event are the same as in the actual world.

As it stands, this is somewhat vague and is incomplete in at least one respect. It is vague primarily because of the use of the word 'relevant' to identify the initial conditions in question, and it is incomplete because only events that are significant in some way to agents are ever considered lucky (or unlucky) and this characterisation of the notion makes no mention of the significance of the event. These two issues set aside, however, there are good grounds for thinking that such a characterisation of

luck is on the right tracks. Consider the paradigm case of a lucky event—the lottery win. Clearly we have here an event that obtains in the actual world but which does not obtain in most of the nearest possible worlds to the actual world where, intuitively, the relevant initial conditions are the same. Although one might win a free and fair lottery with long odds in the actual world, in most near-by possible worlds where, for example, one continued to buy the lottery ticket that one did and the lottery remained free and fair, one would not be standing here now clutching the winning lottery ticket.¹⁴

This account of luck also seems to capture what is at issue in the Gettier-type cases, where the luck concerns the truth of the belief in question (and where the relevant initial conditions are here largely determined by the way in which the agent formed her belief in the actual world). So whilst, in the actual world, the agent ends up with a true belief by looking at a stopped clock, this belief is only luckily true because in most near-by possible worlds where the agent forms her belief in the same way as in the actual world (by looking at this clock), her belief will be false.

The type of luck at issue in the Gettier-type cases is what I will call *veritic luck*, and this can be characterised as follows:

Veritic Luck

For all S, the truth of S's belief in a contingent proposition, φ , is veritically lucky if, and only if, S's belief that φ is true in the actual world, but false in most near-by possible worlds in which the belief is formed in the same manner as in the actual world.

Given this understanding of what is at issue in Gettier-type cases, it should come as little surprise to find that safety-based theories of knowledge are able to directly deal with the problem since what safety ensures is that one's true belief is not veritically lucky. If in order to meet this principle one must only believe that proposition in near-by possible worlds where the belief is true, then it is little wonder that the belief in question will not be veritically lucky, and thus immune to a Gettier-style scenario. This is because there won't be the logical space within which the agent has a true belief in the actual world but not in most near-by possible worlds where she forms her belief in the same way as in the actual world.

Where Gettier-style cases are our focus, it can thus seem as if the goal of epistemology is merely to ensure that our beliefs are not veritically lucky. Notice, however, that veritic luck and the safety principle that eliminates it both understand the possible worlds in question in terms of an ordering that does not essentially depend upon what the agent has reflectively accessible grounds for believing this ordering to be but only on the nature of the actual world (call this the 'objective' ordering). As a result, a mere safety-based account which only eliminated veritic luck would lead to epistemological externalism by counting such agents as our chicken sexer as knowers. After all, their beliefs meet the safety principle and thus are not susceptible to veritic luck (and thus cannot be

‘Gettiered’). Nevertheless, one might worry that there is still a sense remaining in which such beliefs are lucky.

Indeed, one finds commentators remarking as regards examples just like the chicken sexer case that the true belief in question *is* lucky. Here, for example, is Linda Zagzebski:

The value of the truth obtained by a reliable process in the absence of any conscious awareness of a connection between the behaviour of the agent and the truth he thereby acquires is no better than the value of the lucky guess. [Zagzebski 1996: 304]

And elsewhere she has argued that what distinguishes externalists from internalists about knowledge is the fact that the former are unconcerned about luck:

The dispute between externalists and internalists looms large mostly because of ambivalence over the place of luck in normative theory. Theorists who resist the idea that knowledge [...] is vulnerable to luck are pulled in the direction of internalism [...] Externalists are more sanguine about luck. [...] There is lots of room for luck in externalist theories since the conditions that make it the case that the knower is in a state of knowledge are independent of her conscious access. [Zagzebski 1996: 39]

This cannot be altogether correct, however, since, as we have just seen, even epistemological externalists are keen to eliminate the veritic luck at issue in the Gettier-type cases. So given that it is not the kind of veritic luck that is at issue in the Gettier-type cases, what type of luck is it that Zagzebski has in mind here?

My claim is that it is a variety of luck where the ordering of the possible worlds is determined not in the usual ‘objective’ way in terms of the nature of the actual world, but rather in a different fashion which orders the worlds in terms of what the agent can know by reflection alone (call this the ‘reflective’ ordering). As in the characterisation of epistemic internalism that was offered above, reflective knowledge here includes *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. The idea is that what fixes the orderings of worlds on this conception are not the relevant facts about the actual world, but rather what grounds the agent has for believing the relevant facts about the actual world to be, where one’s grounds are here understood in the usual (internalist) way as being reflectively accessible to one. So whilst the fact of the matter as to whether or not, for example, one’s belief was formed reliably will be crucially relevant to determining the objective ordering, it won’t be relevant to the reflective ordering. Instead, what will be important in this regard will only be what grounds the agent has reflectively available to her for believing that she is reliable. Where the agent’s beliefs are inadequately supported by good reflectively accessible grounds, what counts as a near-by world on the objective ordering could be very different to what counts as a near-by world on the reflective ordering, and thus an agent’s belief could count as lucky by the lights of the reflective ordering even though it doesn’t count as lucky by the lights of the objective ordering (we will consider an example in a

moment).¹⁵

I will call this notion of luck, *reflective luck*:

Reflective Luck

For all S, the truth of S's belief in a contingent proposition, φ , is reflectively lucky if, and only if, S's belief that φ is true in the actual world, but, *given only what S is able to know by reflection alone*, false in most near-by possible worlds in which the belief is formed in the same manner as in the actual world.

Furthermore, I contend that it is specifically this sort of luck that is at issue in the sceptical argument, not its veritic counterpart.

Consider again the chicken sexer that we considered above who, whilst forming beliefs in a safe manner, nevertheless is doing so without having any significant amount of reflective knowledge in support of her beliefs. For example, she does not have any good reflectively available reasons for thinking that she has this ability, much less does she have any reasons available to her which would indicate how such an ability would work if she did have it. One reason why we might consider her beliefs to be lucky is that whilst the truth of her beliefs is not a matter of veritic luck, that this is so is dependent upon the actual world being a certain way and yet our chicken sexer lacks any good reflectively accessible grounds for thinking that the world is this way. After all, although, as it happens, she really does have a reliable belief-forming ability in this regard, given that she lacks any good reflectively accessible reasons for thinking that she has this ability, it is just a matter of luck, *from her point of view*—i.e., in terms of what she can know by reflection alone at that moment in time—that she is reliable in this regard. She could, we might say, have just so easily been wrong and thus be forming her beliefs in an entirely unreliable manner.

Of course, when we say that the chicken sexer 'could just so easily have been wrong', we do not have in mind here veritic luck, since given that she really does have the ability in question then she will not only get to the truth via this method in the actual world but also in most near-by possible worlds as well. Instead, the luck at issue here is *reflective* and is captured by our characterisation of reflective luck above. Although on the standard objective ordering of possible worlds—the one at issue in our characterisation of veritic luck—the chicken sexer's belief is not lucky, the kind of luck at issue here does not concern this objective ordering but rather the reflective ordering, one that only takes into account what the agent can at that moment know by reflection alone—i.e., the reasons in support of her belief that are reflectively available to her.

In order to see this contrast in more detail, compare our 'unenlightened' chicken sexer with an 'enlightened' counterpart who not only has safe true beliefs about the subject matter in question, but also has good reflectively accessible grounds in favour of her beliefs, such as good reasons for thinking that she has this reliable ability and that it works in a certain kind of way. Neither of these counterparts

is forming beliefs in a veritically lucky fashion, but the unenlightened chicken sexer does seem to be forming beliefs in a reflectively lucky fashion in a way that the enlightened chicken sexer isn't—at least if we set aside the problem posed by scepticism for a moment (I will return to this important qualification below).

Take the enlightened chicken sexer first. On a reflective ordering of the possible worlds which only takes into account what she knows by reflection alone, the near-by possible worlds will tend to be the same as on the objective ordering since she has good reflectively accessible grounds for thinking that the actual world is the way that it is in the relevant regards, and it is the nature of the actual world that determines the objective ordering. For example, she has good reflectively accessible reasons available to her for thinking that she has the ability in question, and with this in mind the reflective ordering of possible worlds will tend to place those worlds in which she lacks this cognitive ability further away from the actual world than those worlds in which she retains this ability. As a result, her beliefs will tend to track the truth not only in near-by possible worlds on the objective ordering, but also in near-by possible worlds on the reflective ordering as well, and hence her belief will neither be veritically lucky nor reflectively lucky.

In contrast, in the case of the unenlightened chicken sexer just about any possible world can count as a near-by possible world on the reflective ordering since she knows next to nothing by reflection alone that would distinguish between possible candidate worlds in this respect. Compare, for example, the possible world in which she retains her ability and continues to form true beliefs in this regard and one in which she doesn't have this ability at all and is in fact forming beliefs in a fairly random way that bears no relation whatsoever to the sex of the chicks before her. Since she has next to no relevant grounds that are reflectively accessible to her, both of these worlds will count as near-by on the reflective ordering, and thus there will be a large number of near-by possible worlds where she lacks the ability in question and so is forming a false belief as a result. Hence, whilst her beliefs in this respect are not veritically lucky, they are reflectively lucky.

It thus seems that where externalists and internalists about knowledge diverge is not (*contra* Zagzebski) over whether or not they think luck is incompatible with knowledge (since they both think that luck in some sense is incompatible with knowledge), but rather over the specific sort of luck that they think is incompatible with knowledge. Whilst epistemological externalists focus only on veritic luck, epistemological internalists are also concerned with the elimination of reflective luck.

What might seem to play into the hands of the epistemological internalist here is that there is manifestly something epistemically desirable about having beliefs that are neither reflectively nor veritically lucky, in that we would clearly prefer to be in the epistemic position of the enlightened chicken sexer rather than her unenlightened counterpart. Part of the reason for this is that whilst the true

beliefs of both chicken sexers is of credit to them in the weak sense of being the result of one of their cognitive traits, only the true beliefs of the enlightened chicken sexer is of credit to them in the fuller sense of being a true belief that the *agent* can take credit for. That this is so is highlighted by the fact that we would think it somewhat inappropriate for the unenlightened chicken sexer to claim knowledge of what she believes for the simple reason that, from her point of view, she has no reason for thinking she has knowledge at all. This is not the case for her enlightened counterpart.¹⁶

That granted, however, this doesn't altogether decide the issue of externalism *versus* internalism about knowledge for the simple reason that there is, *prima facie* at any rate, room for what we might term an *epistemic pluralism* here. After all, the externalist can perfectly well concede that reflective luck in one's beliefs is epistemically undesirable whilst consistently maintaining that nevertheless one's beliefs only have to be immune to veritic luck in order to count as knowledge. One could thus regard the elimination of reflective luck as being a desirable epistemic condition to meet even if meeting that condition is not essential for knowledge.

The issue of externalism *versus* internalism in the theory of knowledge aside, however, it does seem that it is reflective luck that is at issue in the sceptical debate and that the externalist responses to scepticism considered above are intellectually dissatisfying precisely because they are only concerned with the more elementary veritic luck. That is, externalist responses to scepticism merely identify a clear sense in which knowledge is possible given that certain conditions obtain which the agent is unable to know by reflection alone to have obtained. In doing so, they establish that the truth of our beliefs is not necessarily due to luck in the specific veritic sense of that term. The sceptical worry, however, is not that one's beliefs are luckily true in a veritic sense, but rather the concern that they are, if true at all, luckily true in a reflective sense. One could therefore construe the sceptical counter-response to anti-sceptical treatments of the closure-based sceptical argument as being that all it ensures is that one's beliefs are, at best, no worse-off than the unenlightened chicken sexer's beliefs, when what we want from a response to scepticism is cognitive responsibility in our beliefs in something like the manner of the enlightened chicken sexer.

Indeed, this seemed to be what was underlying Stroud's 'metaepistemological' complaint that we looked at earlier on. Recall that what he wanted was not merely an account of how, if certain external conditions obtain, we have knowledge, but what he termed a 'philosophical understanding' of the epistemic aspect of the human condition, which demanded "good reasons for thinking that what we believe [*about that epistemic aspect of the human condition*] is true". Accordingly, externalism as an account of knowledge that is applicable to the sceptical problem seems to be ruled-out *tout court* because it isn't even attempting to offer such good, reflectively accessible, reasons.

On this construal of scepticism, the problem isn't merely to show that knowledge is possible in

the minimal sense of showing that we might be able to have non-reflectively lucky true beliefs, but the more substantive challenge to show that we can have non-reflectively lucky true beliefs, and this latter challenge is not touched on at all by externalist responses to scepticism. And notice that this isn't to ascribe a commitment to the internalist thesis regarding knowledge to the sceptic since the sceptic can also allow an epistemic pluralism in just the sense mentioned above. That is, the sceptic can consistently grant that there is a minimal 'brute' form of knowledge that merely demands that agents meet external epistemic conditions and thus which involves a non-reflectively lucky true belief. Nevertheless, she could further claim that what is lacking is the more interesting type of knowledge that is only possessed by those agents whose beliefs are not reflectively lucky. If all our knowledge were of the brute variety possessed by the unenlightened chicken sexer then this would have radical implications for the sense in which we are able to take cognitive responsibility for our beliefs. It would, for example, seem to undermine the possibility of our ever coherently claiming to have knowledge. This could be scepticism enough without the further claim that this kind of knowledge is the only kind of knowledge there is, and thus that knowledge *simpliciter* is impossible.¹⁷

3. EPISTEMIC *ANGST*

Does this mean then that the way to deal with the sceptic is to adopt an internalist theory of knowledge that specifically aims to eliminate the problem of reflective luck? Unfortunately, matters are not nearly so simple. The reason for this is that the internalist move to eliminate reflective luck only works in chicken-sexer-type cases provided we set to one side, as we did earlier, the sceptical problem and, in particular, the sceptical error-possibilities that this problem makes use of. For consider what would happen if we allowed sceptical possible worlds to count in the ordering. Regardless of the epistemic pedigree of the enlightened chicken sexer's reflective knowledge, this would still be consistent with a large class of sceptical worlds counting as near-by worlds. Accordingly, although we might regard her as epistemically better-off than her unenlightened counterpart, she would nevertheless be forming beliefs which were to some significant degree reflectively lucky. Wherever the answer to this form of scepticism lies, then, it does not obviously lie with epistemological internalism about knowledge any more than it does with the alternative externalist theory.

We can see this point more clearly by considering how in presenting reflectively accessible grounds in favour of a belief one must, perforce, assume that one is in the market for offering grounds in the first place, and this means that one must ungroundedly assume that one is not the victim of a radical sceptical hypothesis. Consider, for example, the grounds that one might offer in favour of one's

belief that one's car is parked outside, such as that one appears to see one's car parked there and one has no reason for thinking that this car on the driveway is not one's own. Clearly, such grounds are only suited to the purpose of providing support for the belief in question provided one already discounts the radical sceptical error-possibility that, for example, one is a BIV, since the truth of this error-possibility would undermine each and every ground that is being cited here. The problem is, however, one has no good reflectively accessible grounds for thinking that one is not a BIV. So insofar as one is in a position to offer reflectively accessible grounds one must assume that one is not the victim of a radical sceptical hypothesis, and yet one lacks adequate reflectively accessible grounds for thinking that this assumption is true.

It is this feature of the structure of reasons that has led some commentators to mistakenly reject the principle of closure, in that our reflectively accessible reasons for believing everyday propositions (e.g., that one's car is parked outside on one's driveway in England) will not necessarily transmit across a known entailment to be reflectively accessible reasons for believing the entailed proposition where that entailed proposition concerns the denial of a radical sceptical hypothesis (e.g., that one is not presently a BIV on Alpha Centauri). With epistemological externalism in play, however, there is no reason why a lack of reflectively accessible grounds in support of one's beliefs should prejudice the issue of whether or not one has knowledge, and thus the move to deny closure is premature. What does need to be denied, however, is the claim that our reflectively accessible reasons go, as it were, 'right down' to be reflectively accessible reasons for believing everything that is presupposed in one's knowledge. I can have a safe and to some degree internalistically justified belief in an everyday proposition (and thus know it in some internalistic sense), know that this entails the denial of a sceptical hypothesis, and yet whilst I might know this sceptical hypothesis to be false in some brute externalist sense (in that I have a safe belief in this proposition), I do not know it in an internalistic sense that demands sufficient reflectively accessible grounds in favour of the belief in question.¹⁸

The sceptical argument that is emerging is thus very different from the one that we started with in that it is not directed (at least not primarily) at the possibility of knowledge *simpliciter* at all, but rather at the possibility of knowledge of a certain sort—knowledge which one can take full cognitive responsibility for and which thus consists in the complete elimination of both veritic and reflective luck. What the sceptic highlights to us, however, is that one can only completely eliminate reflective luck provided that one groundlessly disregards sceptical error-possibilities, and thus that there is an sense in which our knowledge is ineliminably subject to reflective luck. In this attenuated sense, then, the sceptical challenge *does* go through in that we need to accept the ultimately groundless nature of our believing.¹⁹

I think that it is an implicit awareness that this is the case which at least partly explains why,

when it comes to discussions of scepticism, commentators depart from the standard line that knowledge is incompatible with luck and start to allow at least a certain degree of luck to be compatible with knowledge possession. Zagzebski writes, for example, that:

Epistemic luck permeates the human condition whether for good or ill. [Zagzebski 1999: 109]

And this claim is echoed by Michael Williams in the following passage:

Knowledge and justification always involve [...] an element of *epistemic luck*. A belief whose truth is *wholly* accidental cannot count as knowledge. But [...] getting things right is never wholly nonaccidental either. [Williams 1999: 59]

Williams and Zagzebski are wrong, however, to think that (some degree of) epistemic luck *simpliciter* is consistent with knowledge since veritic luck clearly isn't ever compatible with knowledge possession. Nevertheless, what they say is at least half-right, in the sense that it is only possible to have knowledge of a certain sort provided one allows reflective luck to infect one's beliefs.

So what then is moral of scepticism? It is, I think, that we are unavoidably subject to what I will call an *epistemic angst*, where *angst* is here understood as a general fear about the nature of our epistemic position which is not due to any *specific* empirical challenge to our putative knowledge. Instead, it is caused by the discovery, in the context of reflection, that the ultimate scope of our cognitive responsibility is severely restricted. More specifically, what we discover is that the kind of cognitive responsibility that we standardly attribute to ourselves is only possible given the correctness of a backdrop of anti-sceptical assumptions the truth of which we cannot be cognitively responsible for. Crucially, however, no plausible theory of knowledge (even an internalist theory of knowledge) can adequately allay the problem of epistemic *angst*, and this is the source of scepticism. Scepticism is therefore an *existential* problem, not in the sense that anyone actually endorses the sceptical conclusion, but in the sense that the source of scepticism lies in an essential feature of the human condition. Ultimately, what reasons we have for our everyday beliefs will be of a pragmatic, rather than an epistemic nature, in that it is only by setting aside certain kinds of error-possibilities that one can coherently engage in the practice of offering grounds in the first place, but this is, crucially, to supply a very different kind of buttress to our beliefs to that which we originally sought.²⁰

University of Stirling

REFERENCES

- Black, T. 2002. A Moorean Response to Brain-In-A-Vat Scepticism, *Australasian Journal of Philosophy* 80: 148-63.
- Brandom, R. 1998. Insights and Blindspots of Reliabilism, *The Monist* 81: 371-92.
- Cohen, S. 2000. Contextualism and Skepticism, *Philosophical Issues* 10: 94-107.
- Dancy, J. 1985. *Introduction to Contemporary Epistemology*, Oxford: Blackwell.
- DeRose, K. 1995. Solving the Skeptical Problem, *Philosophical Review* 104: 1-52.
- DeRose, K. 2000. How Can We Know that We're Not Brains in Vats?, *The Southern Journal of Philosophy* 38: 121-48.
- Dretske, F. 1970. Epistemic Operators, *Journal of Philosophy* 67: 1007-23.
- Foley, R. 1987. *A Theory of Epistemic Rationality*, Cambridge, MA: Harvard University Press.
- Fumerton, R. 1990. Metaepistemology and Skepticism, in *Doubting: Contemporary Perspectives on Skepticism*, eds. M. D. Roth & G. Ross, Dordrecht: Kluwer: 57-68.
- Fumerton, R. 1995. *Metaepistemology and Skepticism*, Lanham, MD: Rowman & Littlefield.
- Kornblith, H., ed. 2001. *Epistemology: Internalism and Externalism*, Oxford: Blackwell.
- Lewis, D. 1996. Elusive Knowledge, *Australasian Journal of Philosophy* 74: 549-67.
- McDowell, J. 1994. *Mind and World*, Cambridge, MA: Harvard University Press.
- Moore, G. E. 1925. A Defence of Common Sense, in *Contemporary British Philosophy*, 2nd series, ed. J. H. Muirhead, London: Allen and Unwin.
- Moore, G. E. 1939. Proof of an External World, *Proceedings of the British Academy* 25: 273-300.
- Neta, R. 2002. S Knows that P, *Noûs* 36: 663-81.
- Neta, R. 2003. Contextualism and the Problem of the External World, *Philosophy and Phenomenological Research* 63: 1-31.
- Nozick, R. 1981. *Philosophical Explanations*, Oxford: Oxford University Press.
- Pritchard, D. H. 2001. Radical Scepticism, Epistemological Externalism, and "Hinge" Propositions, in *Wittgenstein-Jahrbuch 2001/2002*, ed. D. Salehi, Berlin: Peter Lang: 97-122.
- Pritchard, D. H. 2002a. McKinsey Paradoxes, Radical Scepticism, and the Transmission of Knowledge across Known Entailments, *Synthese* 130: 279-302.
- Pritchard, D. H. 2002b. Radical Scepticism, Epistemological Externalism and Closure, *Theoria* 69: 129-61.
- Pritchard, D. H. 2002c. Recent Work on Radical Skepticism, *American Philosophical Quarterly* 39: 215-57.
- Pritchard, D. H. 2002d. Resurrecting the Moorean Response to the Sceptic, *International Journal of Philosophical Studies* 10: 283-307.

- Pritchard, D. H. 2002e. Two Forms of Epistemological Contextualism, *Grazer Philosophische Studien* 64: 19-55.
- Pritchard, D. H. 2003a. McDowell on Reasons, Externalism and Scepticism, *European Journal of Philosophy* 11: 273-94.
- Pritchard, D. H. 2003b. Virtue Epistemology and Epistemic Luck, *Metaphilosophy* 34: 106-30.
- Pritchard, D. H. 2004a. Epistemic Luck, *Journal of Philosophical Research* 29: 193-222.
- Pritchard, D. H. 2004b. Neo-Mooreanism, Contextualism, and the Evidential Basis of Scepticism, in *Contextualism*, ed. Danilo Suster, New York: Rutgers University Press.
- Pritchard, D. H. 2004c. The Structure of Sceptical Arguments, *The Philosophical Quarterly* 54.
- Pritchard, D. H. 2005. Wittgenstein's *On Certainty* and Contemporary Anti-Scepticism, in *Investigating On Certainty: Essays on Wittgenstein's Last Work*, eds. D. Moyal-Sharrock & W. H. Brenner, London: Palgrave Macmillan.
- Pritchard, D. H. *forthcominga*. *Epistemic Luck*, Oxford: Oxford University Press.
- Pritchard, D. H. *forthcomingb*. Externalism, Skepticism and Luck, in *Internalism and Externalism in Semantics and Epistemology*, ed. S. Goldberg, Oxford: Oxford University Press.
- Pritchard, D. H., & Smith, M. 2004. The Psychology and Philosophy of Luck, *New Ideas in Psychology* 22: 1-28.
- Pryor, J. 2000. The Skeptic and the Dogmatist, *Noûs* 34: 517-49.
- Pryor, J. 2001. Highlights of Recent Epistemology, *The British Journal for the Philosophy of Science* 52: 95-124.
- Putnam, H. 1981. *Reason, Truth and History*, Cambridge: Cambridge University Press.
- Russell, B. 1948. *Human Knowledge: Its Scope and its Limits*, London: Allen & Unwin.
- Sosa, E. 1994. Philosophical Scepticism and Epistemic Circularity, *Proceedings of the Aristotelian Society*, supp. vol., 68: 263-90.
- Sosa, E. 1999. How to Defeat Opposition to Moore, *Philosophical Perspectives* 13: 141-54.
- Stroud, B. 1994. Scepticism, 'Externalism', and the Goal of Epistemology, *Proceedings of the Aristotelian Society*, supp. vol., 68: 290-307.
- Unger, P. 1975. *Ignorance—A Case for Scepticism*, Oxford: Clarendon Press.
- Vogel, J. 1999. The New Relevant Alternatives Theory, *Philosophical Perspectives* 13: 155-80.
- Williams, M. 1991. *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*, Oxford: Blackwell.
- Williams, M. 1999. Skepticism, in *Epistemology*, eds. J. Greco & E. Sosa, Oxford: Blackwell.
- Williams, M. 2003. Wittgenstein's Refutation of Idealism, in *Wittgenstein and Scepticism*, ed. D. McManus, London: Routledge.

- Williamson, T. 2000. Scepticism, Semantic Externalism and Keith's Mom, *The Southern Journal of Philosophy* 38, 148-57.
- Wittgenstein, L. 1969. *On Certainty*, eds. G. E. M. Anscombe & G. H. von Wright, tr. D. Paul & G. E. M. Anscombe, Oxford: Blackwell.
- Wright, C. 2003. Wittgensteinian Certainties, in *Wittgenstein and Scepticism*, ed. D. McManus, London: Routledge.
- Zagzebski, L. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge; Cambridge University Press.
- Zagzebski, L. 1999. What is Knowledge?, in *Epistemology*, eds. J. Greco & E. Sosa, Oxford: Blackwell: 92-116.

NOTES

- ¹ One finds explicit formulations of the sceptical paradox along these lines in, for example, DeRose [1995], Sosa [1999], Vogel [1990], and Pritchard [2002d], and even where the sceptical paradox is not explicitly formulated in this way, it is usually clear that such a formulation is being presupposed. For more on how the sceptical problem is understood in the contemporary debate, see Pritchard [2002c].
- ² This is not the only way that one can argue for the second premise since one could, for example, defend it via an appeal to infallibilism. Nevertheless, this premise is standardly motivated via an appeal to closure and, in any case, closure is a logically weaker principle than any corresponding infallibilism principle, making it both more plausible and more dialectically interesting (in that any argument for the rejection of closure will, *ipso facto*, be an argument for the rejection of infallibilism also). The *locus classicus* for discussions of infallibilism is, of course, Unger [1975].
- ³ See also Cohen [e.g., 2000]. A very different form of contextualist anti-sceptical proposal—one that I have not the space to discuss here—is the inferential contextualism advanced by Williams [see, especially, 1991]. For discussion of this type of contextualist proposal in the light of its attributer contextualist counterpart, see Pritchard [2002c, §§5-7; 2002e].
- ⁴ In this characterisation of an internal epistemic condition I follow the account given by Pryor [2001] in his extremely useful survey of recent trends in epistemology.
- ⁵ For more discussion of the chicken sexer example, see Foley [1987: 168-9], Lewis [1996], Zagzebski [1996: §2.1 & §4.1] and Brandom [1998]. For more on the epistemological externalism/internalism distinction in general, see the papers collected in the excellent anthology edited by Kornblith [2001].
- ⁶ Note that it is important in this respect that we have defined the BIV hypothesis as we have, since on some construals of this hypothesis there may well be *a priori* (and thus reflectively accessible) grounds available for thinking that it is false, grounds which derive from a certain understanding of content externalism. The *locus classicus* in this respect is, of course, Putnam [1991: chapter 1]. In general, it is typically taken for granted that whilst one might be able to dismiss certain kinds of sceptical hypotheses on content externalist grounds, one will not be able to dismiss all of them, and thus that the sceptical problem cannot be entirely resolved on content externalist grounds alone. One exception to this is McDowell [e.g., 1994], who has argued in several places that scepticism is the result of a content internalist way of thinking about the relationship between thought and world. I discuss the McDowellian approach to scepticism in detail in Pritchard [2003a; *forthcomingb*].
- A related claim that also finds its most influential expression in the work of McDowell is the idea that we can have reflective access to *factive reasons*—i.e., reasons which entail truths about the world, such as that one *sees that* the world is such-and-such a way (which entails that the world is that way). On this view, one is able to have reflective knowledge of empirical facts about the world, such as that one has two hands (because one has reflective knowledge of the factive reason that one sees that one has two hands), and so one can, it seems, argue on this basis to anti-sceptical conclusions which are entailed by the empirical claim in question (such as that there is an external world). A similar view has been put forward more recently by Pryor [2000]. This proposal is highly controversial, and so I don't intend to dwell on it here (I discuss it at length in Pritchard [2003a]). A 'contextualised' version of the 'factive reasons' thesis has also recently been propounded by Neta [2002; 2003], and I discuss this variant on the thesis at length in Pritchard [2004b].
- ⁷ For further discussion on the relationship between the Dretskean anti-sceptical proposal and its equivocal attachment to epistemological externalism, see Pritchard [2002b]. In any case, as both Williams [1991, chapter 8] and Black [2002] point out, properly understood the sensitivity principle does not entail that we are unable to know the denials of sceptical hypotheses anyway, and so the issue of whether sensitivity entails the rejection of closure is moot.

⁸ Not all attributer contextualists have conceded this implication of their view. Some of them have claimed, for example, that the knowledge in question here is in some sense *a priori*, either in that it concerns some sort of presupposition of rational inquiry, or because of semantic externalist considerations. I do not have the space to go in to these issues in depth here (though needless to say I am unsympathetic to this general line of argument), but for the relevant literature in this respect, see Cohen [2000] and DeRose [2000; cf. Williamson 2000].

⁹ Notice that I am only claiming here that the safety-based view poses no problem for the sceptical and anti-sceptical employment of the closure principle, not that such views can retain closure in general, which would require further argument.

¹⁰ Stroud was explicitly responding here to Sosa [1994].

¹¹ Fumerton develops this line of argument at length in Fumerton [1995].

¹² The stopped clock Gettier-style example is due to Russell [1948: 170-1], though he didn't himself realise that it was an example of this sort. I discuss the relationship between epistemic luck and cognitive achievement in more detail in Pritchard [2003b; *forthcominga*: chapter six].

¹³ Note that her beliefs will be sensitive as well, since in the nearest possible world in which what she believes is false (because, for example, the chicks in question have been swapped around) she will not believe this proposition via the same method as in the actual world.

¹⁴ A critical survey of the philosophical literature on luck along with a detailed defence of this characterisation of luck in the light of that survey is offered in Pritchard & Smith [2004] and Pritchard [2004a; *forthcominga*: chapter five].

¹⁵ It will also be possible for an agent's belief to count as lucky by the lights of the objective ordering whilst not being lucky by the lights of the reflective ordering because she has good reasons for taking the actual world to be a certain way and yet, because of (for instance) the truth of a sceptical hypothesis, the actual world is in fact radically different from what the agent takes it to be. It is the possibility of 'reflective' luck even in the absence of 'objective' (i.e., veritic) luck that is our focus at present though.

¹⁶ I expand on this dual notion of epistemic credit in more detail in Pritchard [2003b; *forthcominga*: chapter six].

¹⁷ One of the main advantages of understanding the commonly stated worry about externalist responses to scepticism (such as the metaepistemological challenge we noted above) in terms of reflective luck is that it does not simply treat such challenges as being implicit appeals to an internalist epistemology. One might think, for example, that the metaepistemological challenge was simply a way of putting the internalist point that if we really do have knowledge then we must also know that we know. Accordingly, since the denial of this thesis (the so-called 'Iterativity' thesis) is associated with epistemological externalism, whilst its advocacy is associated with epistemological internalism, the complaint just seems to be another way of expressing one's prior commitment to epistemic internalism, thereby depriving it of any real interest to those who do not share this commitment. In contrast, if one understands the challenge in terms of the familiar (and universal) intuition that knowledge excludes luck, then one can characterise this problem in a non-partisan way in that whilst one *might* ally the complaint to an appeal to epistemic internalism, this would not be essential to stating the challenge. Instead, one could grant that knowledge is sometimes possessed even in the absence of adequate supporting grounds—such 'knowledge' is, after all, 'non-lucky' in one (veritic) sense of that phrase, even if it does not accord with the Iterativity principle—without thereby conceding that the externalist response to scepticism is entirely intellectually satisfying on the grounds that it fails to capture a further sense in which we desire our knowledge to be non-lucky, one that makes reference to reflective luck.

¹⁸ I develop this point about the failure of what I call 'internalist closure'—as opposed to the failure of closure *simpliciter*—in Pritchard [2002a; 2002b]. It should be noted that the picture is in fact even more complicated in that one lacks reflectively accessible grounds for believing lots of everyday propositions as well insofar as those propositions are treated as being as certain as anything in that context. Wittgenstein [1969] famously referred to these propositions as "hinge" propositions, offering the 'Moorean' proposition 'I have two hands' as an example. Since nothing is more certain than this in normal circumstances, it follows that nothing could coherently count as a supporting ground for belief in this proposition (since this would mean that it was more certain than one's belief in this proposition) and, relatedly, nothing could act as a ground for doubt in this proposition either (since the ground for doubt would, perforce, be more open to question than the belief itself). In this sense, then, hinge propositions are (like the denials of radical sceptical hypotheses) both groundless *and* not coherently open to doubt. That these everyday propositions play this role is because doubt of such a proposition as this in normal circumstances would be tantamount to doubt of the denial of a radical sceptical hypothesis. If, in normal circumstances, it is questionable whether or not I have two hands, then everything I believe is also questionable. An apparently local doubt about an everyday proposition is thus in fact a radical doubt about whether my reasons 'hook-up' with the world at all. I explore the relevance of Wittgenstein's remarks in this regard to this debate in Pritchard [2001; 2003a; 2005]. See also Williams [2003] and Wright [e.g., 2003].

¹⁹ How should one formulate such a sceptical argument? One could simply re-phrase the template argument set-out above so that it is made explicit that it is a certain kind of knowledge that is its focus. I argue in Pritchard [2004c], however, that the sceptical argument is best understood as turning on an 'underdetermination' principle rather than on a closure principle, where this demands that one's reflectively accessible grounds for one's beliefs should prefer those beliefs over known to be incompatible alternatives, at least insofar as one seeks the kind of (internalist) knowledge that requires reflectively accessible grounds.

²⁰ Earlier versions of this paper have been presented at a number of venues, including at *The Scots Philosophical Club* Annual Meeting (University of Stirling, 2002), at a conference on ‘Skepticism’ (Amsterdam Free University, 2002), at a *Society for Skeptical Studies* session at the Eastern APA (San Francisco, 2003), and at a NAMICONA conference on ‘Theories of Knowledge’ (Kolding, Denmark, 2003). I am grateful to the audiences on these occasions, especially Lars Binderup, Martijn Blaauw, Tim Black, Berit Brogaard, Sandy Goldberg, Richard Greene, Patrick Greenough, Lars-Bo Gundersen, Jesper Kallestrup, Klemens Kappel, Michael Williams, and René van Woudenberg. Thanks also go to the following people who read and commented on earlier versions of this material: Guy Axtell, Michael Brady, Rowan Cruft, Antony Duff, Neil Gascoigne, Ram Neta, Phyllis McKay, and Alan Millar. I am also extremely grateful to the editor of the *Australasian Journal of Philosophy*, Maurice Goldsmith, for his subtle commentary on this paper, and also to no less than *three* anonymous referees who commented on this article and enabled it to become a much more polished piece of work as a result. Finally, I would like to thank *The Leverhulme Trust* for the award of a two-year Special Research Fellowship which has enabled me to conduct work in this area.