

How Hard Are the Sceptical Paradoxes?

ALEX BYRNE

Massachusetts Institute of Technology

The sceptic about the external world presents us with a paradox: an apparently acceptable argument for an apparently unacceptable conclusion—that we do not know anything about the external world. Some paradoxes, for instance the liar and the sorites, are very hard. The defense of a purported solution to either of these two inevitably deploys the latest in high-tech philosophical weaponry. On the other hand, some paradoxes are not at all hard, and may be resolved without much fuss. They do not contain profound lessons about the human condition. Where should we place the sceptical paradoxes?

Rather strikingly, there is no consensus on the matter. Everyone agrees that the liar and the sorites are deeply challenging, even—or perhaps especially—those who claim to have solved them. Many philosophers think that (at least some) sceptical paradoxes are up there with the liar and sorites. Some of these philosophers are optimists—they think that there is a way out, short of embracing scepticism. Some are pessimists, who see little hope. But they are all united in the conviction that one or another sceptical paradox is fearfully difficult.¹ On the other side, there are many who are decidedly underimpressed. To them, sceptical paradoxes do not “signal genuine collisions between features of our thinking which go deep”²—the sceptic is just another guy with a bad argument.³

This paper argues for the second “ho-hum” reaction. Two sceptical paradoxes are considered in some detail (sections 3 and 4): the first formulated with particular clarity by Ayer (1956, ch. 2); the second the centerpiece of the classic pessimist manifesto *The Significance of Philosophical Scepticism* (Stroud 1984a). It is argued—admittedly with the benefit of a century of hindsight—that each has a premise that is only superficially plausible. As far as I know, there is really only one other candidate for a sceptical paradox that is worth taking seriously. As it happens, I find it completely unparadoxical;

for reasons that will become apparent (section 2), I cannot hope to convince anyone who thinks otherwise. Although this paper has relatively sweeping ambitions, I stress that it makes little pretence to novelty; however, the case evidently needs to be prosecuted again. I also stress that *external world* scepticism is the sole topic: for all that is argued here, the problem of other minds (for example) might pose a genuinely difficult sceptical paradox. One bonus of our discussion of sceptical paradoxes is some insight (section 5) into Stroud's recent claim that the philosophical project of "understanding human knowledge in general" must end in failure.

First, though, some preliminaries.

1. Sceptical games

It is useful to think of ourselves and the sceptic as playing a game—a deadly serious one, with nothing less than our knowledge of the external world at stake. What are its rules?

Since we are concerned with the sort of sceptic who presents a grave threat to knowledge, the rules are simple.⁴ En route to arguing that we have no knowledge of the external world, the sceptic is only allowed to appeal to more-or-less uncontroversial or obvious premises (we may assume that the sceptic's premises entail her conclusion). If the sceptic succeeds in producing such an argument, she has constructed a genuinely difficult paradox, and may therefore be somberly declared the winner. (Naturally that doesn't mean that scepticism is *true*—it is a *paradox*, after all.) If, on the other hand, we manage to expose one of the sceptic's premises as false or highly dubious, then the sceptic loses.

Here there is a significant point of disanalogy with the sorites paradox. The conclusion of the sorites is—or is easily extended to—a contradiction (say, there is someone who is both bald and not-bald). The conclusion is therefore utterly unacceptable (even those who accept that liar sentences are true contradictions balk at extending this idea to the sorites). So the paradox-monger cannot fairly protest, against a solution to the sorites involving some complicated and controversial technical apparatus, that embracing the conclusion affords a simpler and less ad hoc solution. Contrast the sceptical paradox: its conclusion is not completely beyond the pale (Unger 1975).⁵ Indeed, as is familiar to anyone who has taught the *Meditations*, undergraduates often fall over themselves in their eagerness to proclaim that they know practically nothing at all. Consequently, denying one of the sceptic's premises by appeal to, for instance, a philosophical analysis of knowledge, is best avoided. Such a solution involves establishing the following conditional, and then affirming the antecedent: *If* [insert abstruse piece of philosophy here] is true, *then* a certain apparently innocuous sceptical premise is false. The choice between the sceptical conclusion and the abstruse piece of philosophy may not be so clear.

This game—which may be called *expose-the-sceptic*—should be distinguished from two others, both of which are easier for the sceptic to play. One is much like *expose-the-sceptic*, except that the sceptic is allowed to appeal to contentious or unobvious premises. So, for example, a sceptic playing this game might appeal to some version of the sense-datum theory, or the claim that knowledge implies certainty, or that reasons are either “deductive or defective”. If the sceptic succeeds in showing that scepticism follows, then this is undeniably a very useful result, for the sceptic’s argument serves as a *reductio* of the contentious premise; of course it is not a paradox. Naturally it is a vague matter where the first game ends and the second begins; but the distinction is serviceable enough.

In the other alternative game—*convince-the-sceptic*—we make the first move. Our objective is to try to convince the sceptic that we have knowledge of the external world, without “begging the question” against her. Specifically, we cannot appeal to any premise that is known on the basis of perceptual experience, because the legitimacy of such premises is precisely what the sceptic doubts. G. E. Moore is sometimes supposed to have been a brave but spectacularly incompetent player of this game, striking out on behalf of Commonsense.⁶

If we win at the first game of *expose-the-sceptic*, then our knowledge of the external world is secure. In other words: if there *isn't* an apparently acceptable argument for the sceptical conclusion, then the sceptic may safely be ignored. She may be admirably consistent in refusing any knowledge tainted with perceptual experience, but starvation diets are generally unwise, in philosophy as elsewhere. To redirect Searle’s invective against the functionalist, the sceptic doesn’t need an argument, she needs treatment.

This claim, that our (putative) external world knowledge is innocent until proven guilty, is occasionally denied, famously so by Kant, who complained that the existence of “things outside of us . . . must be accepted merely on *faith*”, unless we are able to counter the sceptic’s doubts with a “satisfactory proof” (1933, Bxl). For a contemporary—and less extreme—example, take Audi, who distinguishes between “rebutting” and “refuting” scepticism: the former involves “showing that [the sceptical argument] is not sound”; the latter involves showing that the sceptical conclusion is false (1993, 354). “Refuting” the sceptic is not quite as arduous as *convincing* the sceptic, but evidently a simple Moorean claim that one knows one has a hand is not good enough to count as an Audian refutation.⁷ With these distinctions, Audi explains the dialectical position as follows: “Rebuttal of a skeptical thesis would entitle us to withhold it, and a rebuttal based on successful criticism of a sufficiently wide range of skeptical arguments might warrant regarding skepticism as, say, probably false; refutation of a skeptical thesis would entitle us to deny it” (1993, 354). But this is exactly like saying that if Zeno’s paradoxes have only been *rebutted*, then we should regard the thesis that nothing moves as, at best, “probably false”. It is not a speculative

conjecture, a prejudice of our culture, or church dogma, that people frequently move from A to B. Absent considerations to the contrary, we really can be quite confident that people move. It is likewise not a speculative conjecture (etc.) that some people know a lot about the weather, or that they have hands. Absent considerations to the contrary, we really can be quite confident that people know such things. This is not a deep point about “Moorean facts”, “ordinary language”, “hinge propositions”, “epistemic conservatism”, or anything like that; it is a point of extreme shallowness, which is why it is clearly true.

Not only is it not *necessary* to convince the sceptic to win the first game, it isn’t even *sufficient*. Convincing the sceptic does not diagnose the mistake in the sceptical paradox. If the sceptic is convinced, then there is an impressive argument that the conclusion of the sceptical paradox is false. If anything, this makes the paradox more acute (Wright 1985, 440–1).

Only muddied waters can result from failing to separate this last game of convince-the-sceptic from the first game of expose-the-sceptic. For example, in *The Significance of Philosophical Scepticism*, Stroud starts by setting out a sceptical paradox. The sceptic is vividly presented as a wood-boring weevil aboard Neurath’s boat, destroying our knowledge from within: simply drawing on the epistemic standards employed by detectives, cooks, plumbers and garage mechanics, she leads us to admit that we know considerably less than we originally thought. Clearly we are playing the first game, and so one might anticipate that when Stroud goes on to consider some responses, they will be directed against the sceptic’s seemingly innocuous premises. However, for the most part that is not what happens: the responses Stroud discusses seem best classified as attempts to *convince* the sceptic—a Moorean proof; an Austinian reproach (“‘You cannot fool all of the people all of the time’ is ‘analytic’”); transcendental idealism (“inner experience... is possible only on the assumption of outer experience”); verificationism (the sceptical conclusion is meaningless); and naturalized epistemology (a Moorean proof with “scientific” premises).⁸ Stroud finds all these wanting, but the present point is that even if one succeeded, it is not at all clear how it would help resolve the initial sceptical paradox. Even if, say, transcendental idealism convinces us that the sceptical conclusion cannot possibly be accepted, the seductiveness of Stroud’s sceptical reasoning has not been explained away.⁹

Here is another illustration of the importance of separating convince-the-sceptic from expose-the-sceptic. Consider Strawson’s well-known neo-Humean proposal that “sceptical doubts... are simply to be neglected... They are to be neglected because they are *idle*; powerless against the force of nature, of our naturally implanted disposition to belief” (1983, 13). Strawson is sometimes interpreted as conceding that sceptical paradoxes are genuinely difficult, while recommending that we don’t waste time with them. This is a curious view for an erstwhile Waynflete Professor of Metaphysics, to say the least. Imagine a philosopher who counsels against research into vagueness,

because the conclusion of a sorites argument is “powerless against our naturally implanted disposition to belief”! (Cf. Sosa 1988, 163.) Barbers, painters and builders’ merchants need not bother themselves with the sorites paradox, but what *else* are philosophers supposed to do? However, once the first and last sceptical games are sharply distinguished, a much more attractive interpretation suggests itself. For Strawson nowhere claims that there is an apparently flawless argument *for* scepticism. Rather, his point seems to be that it is both futile and unnecessary to try to *convince* the sceptic—the sceptic’s challenge to produce a non-question-begging argument against her cannot and need not be met.¹⁰

2. The canonical sceptical argument

The sceptic argues that we have no knowledge of the external world; in other words, that we have no knowledge of *external world propositions*. (It will not be necessary to be more informative about propositions of this type: the intuitive conception will suffice.)

Let p be some external world proposition, say that one has a hand, and let \spadesuit be some sceptical hypothesis inconsistent with p , say that one is a handless brain-in-a-vat having experiences qualitatively identical to those one is actually having. Then what is sometimes called the “canonical sceptical argument” (Brueckner 1994) is this:

- (1) If one knows p then one knows not- \spadesuit
- (2) One does not know not- \spadesuit

Therefore:

- (3) One does not know p

(1) follows from the assumption that one knows that p entails not- \spadesuit and the principle that knowledge is closed under known entailment.¹¹ It is consequently extremely plausible, although of course the closure principle has been denied (Dretske 1970; Nozick 1981, 204–11). If (2)—that one does not know that one is not a brain-in-a-vat—is likewise plausible, then the canonical sceptical argument presents a serious paradox.

Well, *is* (2) plausible? (Let us grant straightaway that \spadesuit is metaphysically possible.¹²) The two main lines of argument for (2) will be examined and rebutted in sections 3 and 4. In fact, on one natural way of formulating these arguments, they can be used to show (3) directly, without appeal to (1). But even if (1) is redundant, given the argument for (2), the canonical sceptical argument is the correct formulation of *one* sort of sceptical reasoning. For according to some philosophers, no *argument* for (2) is needed—or, if an argument *is* required, it merely consists in drawing attention to a sceptical

hypothesis. On this view, when one contemplates some sceptical hypothesis \spadesuit , without any further assumptions it becomes pretty much *obvious* that one does not know not- \spadesuit (see, e.g., Unger 1975, 13–28; Lewis 1996, 503).

Even if (2) does not need an argument, this does not mean that the resulting paradox cannot be solved—although it surely does mean that the paradox is very difficult. Perhaps some sophisticated philosophy will overcome initial resistance to closure, or to the sceptical conclusion; perhaps, following a recent trend, the “contextualist” solution can be adequately defended (for a review, see DeRose 1999). But many can’t bring themselves to accept any of these positions; for them, if mere contemplation of (2) is enough to make it plausible, the paradox is close to intractable.

Here we come to the part of this paper that is unquestionably disappointing. For I can do little more than record my conviction that (2) *does* need an argument.¹³ To the extent that I ever found it plausible, it is because of considerations that I subsequently learnt were unmasked as baseless—which exposés I will be reporting shortly. Those on the other side can hardly deny that I am in good company; alas, I cannot deny that they are either.

Further discussion of this issue is unlikely to resolve anything, and in any case a proper treatment would require another paper. So let us agree to disagree, and turn to the first of the two sceptical paradoxes that do not take (2) as an axiom.¹⁴

3. The argument from underdetermination

The sceptic begins the *argument from underdetermination* with a couple of apparent platitudes. First, one’s knowledge of the external world ultimately rests on perceptual evidence; that is, on appearances. Perhaps perceptual evidence is not *sufficient* for external world knowledge, but it is at least *necessary*. Second, there is a “logical gap” between appearance and reality: propositions about the external world are “underdetermined” by one’s perceptual evidence. An oar in water may *look* bent, but this does not entail that the oar *is* bent. The next step is to argue that the gap cannot be bridged: one’s perceptual evidence e is not strong enough to allow one to know p .

This step can be more or less direct. The sceptic might claim that e does not favor p —or favor p to a high enough degree—over a sceptical hypothesis \spadesuit that is compatible with e . Given this, premise (1) of the canonical sceptical argument—that if one knows p then one knows not- \spadesuit —is not needed to conclude that one does not know p (see Brueckner 1994).¹⁵ This would appear to be Russell’s strategy, who reports that once when asleep he

...did really have an experience intrinsically indistinguishable from that of seeing a ruined church when awake. It follows that the experience which I call “seeing a church” is not conclusive evidence that there is a church, since it may occur when there is no such external object as I suppose in my dream. It

may be said that, though when dreaming I *think* that I am awake, when I wake up I *know* that I am awake. But I do not see how we are to have any such certainty... (1948, 171)

Here Russell is claiming that he does not know that there is a church, because his evidence does not “conclusively” favor this proposition over the sceptical hypothesis that he is falsely dreaming that there is a church.¹⁶

Alternatively, perhaps the claim that *e* does not favor not-♠ (or favor not-♠ to a high enough degree) over ♠ is initially more plausible than the claim that *e* does not favor *p* over ♠, in which case premise (1) of the canonical sceptical argument is not redundant (see Cohen 1998).

For convenience the discussion will revolve around the first variant; the criticism to come will apply equally to the second variant.

The quotation from Russell does not really do the argument from underdetermination justice. Let us grant that one’s evidence for *p* is not *conclusive*—why does that show one does not know *p*? Here the clever sceptic will admit it doesn’t. Instead, she will point out that, given her starting point—a modest foundationalism and a “gap” between one’s perceptual evidence and the external world—our epistemic position is very much like that

... of someone waking up to find himself locked in a room full of television sets and trying to find out what is going on in the world outside. For all he can know, whatever is producing the patterns he can see on the screens in front of him might be something other than well-functioning cameras directed on to the passing show outside the room. The victim might switch on more of the sets in the room to try to get more information, and he might find that some of the sets show events exactly similar or coherently related to those already visible on the screens he can see. But all those pictures will be no help to him without some independent information, some knowledge which does not come to him from the pictures themselves, about how the pictures he does see before him are connected with what is going on outside the room. The problem of the external world is the problem of finding out, or knowing how we can find out, about the world around us if we were in that sort of predicament. It is perhaps enough simply to put the problem this way to convince us that it can never be given a satisfactory solution. (Stroud 1984a, 33)

We do not consider sceptical hypotheses in everyday life, but the sceptic charges that in this respect we are like detectives who have understandably overlooked a bizarre but powerful rival explanation of the evidence—say, as in Conan Doyle’s “The Lion’s Mane”, that the victim died in an accidental encounter with a large jellyfish. Once this alternative explanation is pointed out, the detectives will realize that by ordinary standards of evaluating hypotheses, they should withdraw their former claim that the victim was murdered. Even if the victim was murdered, the alternative hypothesis explains the evidence to a degree that precludes the detectives from knowing the truth.

More detail, obviously, is required before the argument from underdetermination begins to resemble a proof of scepticism. But this much seems fair: the literature strongly suggests that once one has argued oneself into the predicament illustrated by Stroud's example of the room full of television screens, it is no easy matter to argue oneself out.¹⁷

Let us look more closely at the two apparently platitudinous assumptions that introduced the argument from underdetermination. The first was that one's knowledge about the external world rests on perceptual evidence. And the second was that propositions about the external world are underdetermined by this evidence.

The first assumption does appear attractive, especially when it is emphasized that it need only commit us to a weak form of foundationalism: the foundations may be fallible, and coherence may also play an important role.¹⁸

What about the second? Take a particular case. An ordinary subject *S* is looking at a certain object *o*, a square tile on a contrasting background; his eyesight and the lighting conditions are normal. *S* thereby forms the belief that *that* (viz., *o*) is square. According to the second assumption, this external world proposition *p*, *that o is square*, is known—if at all—by *S* on the basis of underdetermining perceptual evidence *e*.

What is this perceptual evidence *e*? According to Russell, one's perceptual evidence is an *experience*. However, the sceptical case is considerably clearer if *e* is taken to be a *proposition*. Then the sceptic's claim that *e* does not "determine" *p* and is "compatible" with some sceptical hypothesis ♠ can be cashed out straightforwardly: *e* does not *entail* *p* and is *consistent* with ♠. And taking *e* to be a proposition is a thoroughly commendable policy in any case, since non-propositional conceptions of evidence are obscure.

(The rest of this section argues that the argument from underdetermination with a propositional *e* fails. In response, the sceptic might want to revive Russell's claim that *e* is an experience; this will be briefly examined in a note at the very end.)

Does *S* believe *e*? Suppose *S* doesn't. Assume, with the sceptic, that there is no other proposition *e** that *S* believes and that is suitably strong evidence for *p*. Then *S* won't know *p*, even if *e* is excellent evidence for *p*. (That there is Trichinopoly cigar ash at the crime scene is excellent evidence that the murderer smokes such cigars; if Lestrade does not believe this evidence, then—absent other evidence—he does not know that the murderer smokes Trichinopoly cigars.) So, if *S* doesn't believe *e*, then the sceptic has a very quick argument for *S*'s not knowing *p*. Too quick, in fact. *S* is an ordinary person in an everyday situation. A convincing case for scepticism will not assume from the outset that gross epistemic failings, such as ignoring essential evidence, are widespread. The sceptic will concede that ordinary people do their best to attain knowledge of the external world, and hence that *S* believes *e*.

Does *S* know *e*? If *S* doesn't know *e* then *S* doesn't know *p*. (Lestrade justifiably believes that there there is Trichinopoly cigar ash at the crime

scene; his belief is true, but by accident, and so isn't knowledge; if Lestrade concludes that the murderer smokes such cigars, this is not knowledge either.) Again, this route to scepticism is too fast: if a *premise* of the sceptical argument is that we lack knowledge of "internal world" propositions like *e*, it can hardly be compelling. Accordingly, the sceptic will concede that *S* knows *e*.¹⁹

The known proposition *e* has yet to be identified. The candidates may be divided into two classes. The first—class I—consists of propositions about *S*'s sense-data, ideas, impressions, phantasms or other queer entities allegedly "given" in experience. The second—class II—consists of propositions about how things look or (visually) appear to *S* (cf. the first paragraph of this section).²⁰

It is quite doubtful that (trivial exceptions aside) any propositions in class I are true, *a fortiori* known; they may accordingly be dismissed. This would have sounded dogmatic as recently as the first half of the twentieth century: it is only in the last fifty years or so that the deep flaws in what used to be called the "representative theory of perception" have become gradually visible. Admittedly, not everyone agrees that the theory rests on a soggy bog of error: in one form or another, it still has its defenders.²¹ However, it is unnecessary here to rehash the argument: because we are playing the first sceptical game, the sceptic must steer clear of philosophical controversy.

That leaves the members of class II: propositions about how things look or appear to *S*—in other words, certain propositions about *S*'s mental states. But because the representative theory of perception is off-limits, there is very little motivation for thinking that one's knowledge of the external world rests on a foundation of knowledge about one's own psychology. For consider the most natural way of trying to motivate it. Perception is a causal process: when one sees a hand, there is a causal chain beginning at the hand and ending in the mind or brain, via the eye. At the distal end of the chain, there is the hand; at the proximal end, there is a certain mental event or some other kind of mental entity, which we may neutrally call a "percept". So perception is importantly like—to borrow Johnston's helpful analogy—Morse code transmission: "both involve interpretable effects at the end of an information-bearing process or signal" (1992, 257).²² In the case of Morse code, one's "perceptual evidence" is that a certain series of dots and dashes has been received; in the case of perception, one's perceptual evidence is that certain percepts are present to the mind. On the Morse code model, there is an object of awareness common to both veridical perception and hallucination: when one sees a tomato there is a "tomato percept" (mental dots and dashes) of which one is aware; likewise when one hallucinates a tomato. However, the claim that there is a common object is the central component of the representative theory of perception, traditionally taken to be established by the notorious "argument from illusion".²³ Once the representative theory of perception is undermined, so is the Morse code model, and with it the main motivation for thinking that one's perceptual evidence exclusively concerns something mental.

The considerably more attractive replacement for the Morse code model does not (or need not) deny that perception is a causal process; the difference is in the nature of the effect. On the Morse code model, when someone looks at his hand, the psychological effect is a “percept”, something that the subject can (it is supposed) unproblematically become aware of because it is in his own mind. The percept itself is not supposed to be an awareness *of* anything: it is the subject’s awareness *of* the percept that explains how, at one remove, he can become aware of his hand. On the replacement “intentional” picture, the effect itself is an *awareness* or an *appearing*—specifically an awareness of the hand, or the hand’s appearing to be a certain shape, or something along similar lines. On some views, the same effect is present in both the veridical and hallucinatory cases, on other views the effects are different. But whatever the details, the subject is not (or need not) be aware *of* the effect in order to be aware of his hand.²⁴

So, once the representative theory and the Morse code model have been left behind, there should be no temptation to seek the elusive proposition *e* in class II. Still, since *e* is in class II if it’s anywhere, it is instructive to examine this option further.

Propositions about how things look or appear to *S* can be divided into two types. The first—type II_E—comprises external world propositions, because they entail the existence of *o*: that *o* looks square to *S*, that it appears to *S* that *o* is square, etc. Hence, propositions of type II_E, despite not entailing *p*, and perhaps being known by *S*, are quite unsuitable candidates to be *e*. For *e* is not supposed to be an external world proposition.

The second—type II_I—comprises those propositions about how things look or appear to *S* that are not external world propositions (or so we may suppose): that it appears to *S* that (some *x*) *x* is square, that it appears to *S* that the *F* is square (for various fillings for ‘*F*’, e.g. ‘tile’, ‘pink thing’), etc. If *e* is to be found in class II, it must be of type II_I.

Moreover—and this is a point in *favor* of the sceptic—if *e* is a type II_I proposition, then no fancy argument is required to show that *S* cannot know *p* on the basis of *e*. Grant that *S* knows that it appears to him that the *F* is square; without the background knowledge that *o* is the *F*, this seems no evidence at all that a *particular object*, *o*, is square.

However (setting aside the more fundamental problem of motivation), it is not plausible that *e* is a type II_I proposition. First, *e* is *true*. And it is somewhat contrived to describe how things look or appear to *S* in existentially quantified terms: that (some *x*) *x* is square, that the pink tile is square, etc. It cannot simply be assumed that this sort of description is correct—an argument is required. Second, *S* *believes* *e*, and it is quite unobvious why *S*, if he is to know *p* via his senses, must have *any* beliefs about how things appear, let alone believe one of the specific propositions under consideration.²⁵ Suppose *S* is a conceptually challenged animal who cannot entertain

these comparatively sophisticated thoughts about appearances; does this fact alone imply that *S* cannot use his eyes to come to know that *o* is square?

We have set the representative theory and the Morse code model aside; but there remains one—albeit superficial—reason for thinking that perceptual evidence concerns the way things look or appear, namely that it is natural to respond to a question like ‘How do you know that this is square?’, with something along the lines of: ‘I see that it is’, ‘It looks square’, ‘It appears to me that it’s square’, etc.

But this consideration does not amount to much. First, note that the answers mentioned are all of (the prohibited) type II_E: the previous three quoted sentences all express *external world* propositions. It would not be at all natural to reply to the question ‘How do you know that this is square?’ by asserting a proposition from the second category. Second, “How-do-you-know” questions are not always properly answered by *asserting* one’s evidence: often one answers by explaining how one *acquired* one’s evidence: ‘I asked him’, ‘I was across the street when it happened’, ‘I used litmus paper’, ‘I’m a doctor’, etc. (See Austin 1946, 79.) Sometimes this is more useful to the questioner than an assertion of one’s evidence would be—it may, for example, permit the questioner to check one’s sources independently.

It would be a mistake to say that propositions about the way things look or appear *never* form an important part of one’s evidence for believing how things are. They frequently do. Searching for one’s car at night in a sodium-illuminated parking lot, a crucial piece of evidence might be that a certain car looks black, for one knows that this is how one’s blue car looks under these conditions. If a line of type looks blurry, that is evidence that one has presbyopia. Examples concerning the other senses are equally easy to find. But the point remains that a large portion of one’s perceptual evidence is not about mental states at all.

The sceptic’s first assumption was that one’s knowledge of the external world rests on perceptual evidence. If this is correct, then this perceptual evidence includes external world propositions, which is inconsistent with the sceptic’s second assumption—that there is a “logical gap” between one’s evidence and external world propositions. The argument from underdetermination is therefore unsound. (This is actually stronger than necessary. For our purposes, it is enough to establish the absence of any good reason to believe that the argument is sound.)

All this is not to say, of course, that one’s perceptual evidence *determines* all the propositions one knows on the basis of it. One may know that it will snow soon, that the cat is shut in the cupboard, that the murderer smokes Trichinopoly cigars, and so on, on the basis of non-entailing perceptual evidence. Given this fact, could the argument from underdetermination be revived in a less ambitious form? The conclusion would be, not that perception gives us *no* knowledge of the external world, but that it only gives us

knowledge of *perceptual evidence*. There are various sceptical hypotheses that are arguably consistent with our perceptual evidence and inconsistent with many ordinary beliefs—for instance the hypothesis that objects exist only when observed, or that objects have no occluded parts. The sceptic might try to argue that some of these hypotheses are good enough potential explainers of the perceptual evidence to undermine many of our everyday claims to knowledge.

Of course, even if this more restricted sort of sceptical argument posed a difficult paradox, it would not affect the conclusion of this paper. It is worth noting, though, that the argument has little intuitive force. This sort of restricted scepticism is importantly different from the more global sort, because external world propositions are among the things a restricted sceptical hypothesis is supposed to explain. And once the familiar external world has a toehold among the *explananda*, it is hard to stop more of it creeping in to the *explanans*. The theory that some mischievous demon destroys everything whenever our backs are turned would seem to fail on any ordinary way of comparing hypotheses: in particular, it gratuitously imposes a fantastically baroque causal structure on material objects.²⁶

It might be objected—as it sometimes is in the case of “inference to the best explanation” replies to the traditional sceptic—that this is not *our* reason for believing that tomatoes exist when we aren’t looking at them. Imagine that Lestrade believes, on the basis of spurious reasoning from his evidence, that the undergardener committed the murder. He does not even consider any rival hypotheses, let alone try to rule them out. As it happens, if any of these rivals were considered by competent detectives, they would be rejected, and further, Lestrade is right. Yet his true belief does not amount to knowledge. In the case of the proposition that tomatoes exist unperceived, why aren’t we—a few philosophers excepted—in Lestrade’s position?

But this is to forget that we are playing expose-the-sceptic, and so the sceptic bears the onus. The sceptic has not given us any *argument* to suppose that we are in Lestrade’s position—that our beliefs not entailed by our perceptual evidence are at best lucky guesses. Even though we do not arrive at these beliefs on the basis of an explicit inference to the best explanation, that is no reason to suppose that they do not amount to knowledge. If we are appropriately sensitive to our evidence, that may be enough.²⁷

4. The argument from the unproved assumption

Sometimes the argument from underdetermination is put in terms of an “unproved assumption” (M. McGinn 1989, 6), or the need for “independent information” (see the quotation from Stroud in the previous section²⁸). The sceptic claims that one is able to cross between appearance and reality only by making a certain assumption (typically, that one’s perceptual faculties are reliable, or something to this effect). But how, the sceptic asks, can this

assumption be known? It can't be known on the basis of perception, because knowing *anything* on the basis of perception requires one to appeal to the assumption as a *premise*. Neither is the assumption knowable a priori. Hence the assumption cannot be known, and scepticism follows. (On this way of running the underdetermination argument, it is of course structurally parallel to Hume's sceptical argument about induction.²⁹)

The previous section argued that there is no gap to bridge, and so the argument from underdetermination doesn't even leave the starting gate. The sceptic propounding the *argument from the unproved assumption* can be thought of as rejecting the argument from underdetermination, but still insisting that an assumption is required. That one's (ostensible) perceptual evidence consists of external world propositions does nothing by itself to rule out the possibility that one's perceptual faculties are unreliable, etc. And, the sceptic claims, if one knows that one has a hand, this knowledge rests in part on the premise that one's perceptual faculties are reliable, and/or premises of a similar sort. According to the sceptic, and in Pryor's terminology, one needs to know that one's perceptual faculties are reliable *antecedently* to knowing that one has a hand (2000, 524–5); in Klein's terminology, knowing that one's perceptual faculties are reliable is a *prerequisite* for knowing that one has a hand (1981, 97).³⁰

An analogy will help illustrate this type of reasoning. Suppose a certain book contains various claims for which we have no independent evidence. The provenance of the book is doubtful—for all we can tell from the cover, it might be a work of fiction. However, our ancient tradition has been to take the book at its word. One day some free thinker inconveniently points out that this practice is underwritten by the tacit assumption that the book is not a work of fiction. If the book truly reports that it is in fact a work of responsible journalism, could we thereby come to know this assumption? Surely not: before we can trust the testimony of the book on *any* topic, we need *independent evidence*—that is, evidence that is *not* based on the testimony of the book—to believe that the book is not a work of fiction. If no such evidence is forthcoming, then the book does not give us knowledge.

Some notation will be useful. Let 'a' be a sentence expressing an unproved assumption, 'p' be a sentence expressing some external world proposition, for instance that one has a hand, and interpret ' $K\alpha > K\beta$ ' as: one knows that α antecedently to knowing that β .³¹ (Note that ' $K\alpha > K\beta$ ' implies both ' $K\alpha$ ' and ' $K\beta$ '.) Then the argument from the unproved assumption can be put as follows:

UA

(U) $K\alpha \rightarrow (Ka > K\alpha)$

(U) is a schema, instances of which are obtained by replacing ' α ' with sentences expressing external world propositions. Replacing ' α ' once

with ‘**p**’, and then replacing it with ‘**a**’ (which itself expresses an external world proposition, for instance that one’s perceptual faculties are reliable), yields:

$$(1) \mathbf{Kp} \rightarrow (\mathbf{Ka} > \mathbf{Kp})$$

And:

$$(2) \mathbf{Ka} \rightarrow (\mathbf{Ka} > \mathbf{Ka})$$

Since the “antecedently-to” relation is (we may assume) asymmetric:

$$(3) \neg(\mathbf{Ka} > \mathbf{Ka})$$

It follows from the consequent of (1) that **Ka**, and from (2) and (3) it follows that $\neg\mathbf{Ka}$. Hence the sceptical conclusion:

$$(4) \neg\mathbf{Kp}$$

In fact, a famous version of (1) is one of the two jointly inconsistent halves of the “Cartesian Circle”: If I know that whatever I perceive clearly and distinctly is true then I *antecedently* know that God exists and is not a deceiver.³²

Now consider the following sceptical argument due to Wright (with “sufficient reason” in place of knowledge):

I cannot acquire sufficient reason to believe that I am not dreaming... by any empirical procedure. *For before carrying out an empirical procedure can give me reason to believe something*, I need to have sufficient reason to believe that it has been properly carried out, *a fortiori* that it has been carried out at all. And I can have sufficient reason to believe that only if I have sufficient reason to believe that I did not dream its execution. So empirically based reason to believe that I am not dreaming is excluded. Since the proposition seems quite unsuitable to be reasonably believed by me a priori, I cannot, the sceptic will contend, acquire sufficient reason to believe it at all. (1985, 433, footnote omitted, first emphasis mine.)

(It seems clear that Wright’s ‘before’ in the second sentence should be construed as ‘antecedently to’.) Running this for knowledge instead of sufficient reason gives us a fragment of UA, the derivation of ‘ $\neg\mathbf{Ka}$ ’ from (2) and (3). The assumption one needs to know antecedently is that one is not “dreaming”, which amounts to knowing antecedently that one’s perceptual faculties are reliable. To see this, note that the hypothesis that one *is* (“philosophically”) dreaming is equivalent to the conjunction of (i) a certain maximally specific proposition detailing the character of one’s actual experience, and (ii) the proposition that one’s perceptual faculties are unreliable. (One may “philosophically” dream that one is speaking in the House of Lords, and one may actually be speaking there. If so, this will be a fluke: “philosophical” dreaming can at best be veridical hallucination, not

veridical perception.) Since (i) is *true*, in order to know antecedently that one is not dreaming one must know antecedently that (ii) is *false*—that one’s perceptual faculties are reliable.

Wright himself conjoins the conclusion of the above argument—that one does not have sufficient reason to believe that one is not dreaming—with other premises to derive the sceptical conclusion that one does not have sufficient reason to believe that one has a hand, etc. But this would appear to be unnecessary. If Wright is correct in claiming that “*before* carrying out an empirical procedure can give me reason to believe something, I need to have sufficient reason to believe . . . that I did not dream its execution”, then the argument in the quotation can be straightforwardly adapted to show that one does not have sufficient reason to believe one has a hand, along the lines of UA.

Wright does not explain why one must have sufficient reason to believe that one is not dreaming antecedently to having sufficient reason to believe anything about the external world, but the quoted passage does contain a footnote referring to Stroud, and indeed on an attractive interpretation the main line of sceptical reasoning in *The Significance of Philosophical Scepticism* is a version of UA.³³ (This interpretation is henceforth assumed; another interpretation, on which the argument is considerably less interesting, is discussed in note 37 below.) So let us turn to Stroud.

Stroud’s discussion is in terms of knowledge, rather than sufficient reason, but the unproved assumption is the same, that one is not dreaming. Letting ‘*d*’ express the proposition that one *is* dreaming, the schema in the Stroudian version of UA is therefore:

$$(U^*) \quad K\alpha \rightarrow (K\neg d \supset K\alpha)$$

Following Stroud, call the claim that every “external world” instance of (U*) is true, *Descartes’ Condition*. We may concede that once we accept Descartes’ Condition, the game is over. But why should we accept it? According to Stroud, it appears to be an instance of a more general principle that we accept in everyday life:

Given our original favourable response to Descartes’s reasoning, then, it can scarcely be denied that what I have called his assumption or condition seems perfectly natural to insist on. Perhaps it seems like nothing more than an insistence of a familiar commonplace about knowledge. We are all aware that, even in the most ordinary circumstances when nothing very important turns on the outcome, *we cannot know a particular thing unless we have ruled out certain possibilities that we recognize are incompatible with our knowing that thing.* (1984a, 24, my emphasis).

Note that the italicized claim should be read as: we cannot know a particular thing unless we have *antecedently* ruled out possibilities that we recognize are incompatible with our knowing that thing.

And later Stroud mentions another possibility for the general principle of which Descartes' Condition is an instance:

Reflecting even on the uncontroversial everyday examples alone can easily lead us to suppose that [the general principle] is something like this: if somebody knows something, *p*, *he must know the falsity of all those things incompatible with his knowing that p* (or perhaps all those things he knows to be incompatible with his knowing that *p*). (29–30, my italics)

Again, for explicitness the word 'antecedently' should be inserted before 'know' in the italicized part.

The first quotation (and the parenthetical part of the second) suggests the following schematic principle (where 'α' and 'β' may be replaced with any sentences expressing external world propositions):

$$(A) [K\alpha \ \& \ K(K\alpha \rightarrow \beta)] \rightarrow (K\beta > K\alpha)$$

And the second quotation suggests an alternative:

$$(B) [K\alpha \ \& \ (K\alpha \rightarrow \beta)] \rightarrow (K\beta > K\alpha)$$

Note that, because 'K(Kα → β)' implies 'Kα → β', the antecedent of (A) is equivalent to the antecedent of (B) plus the conjunct 'K(Kα → β)', and so (by the principle of "strengthening the antecedent") (B) implies (A). We may therefore exclusively consider (A), since the weaker the principle, the stronger the sceptic's case.

Assuming (reasonably enough, at least for knowers like ourselves) that one *does* know that if one knows some external world proposition *p*, then one is not dreaming, then an instance of (A), replacing 'β' with '¬d', gives us Descartes' Condition. (Note that sceptical consequences can be drawn from (A) more quickly: replacing 'α' and 'β' both with '**p**' and assuming that $K(K\mathbf{p} \rightarrow \mathbf{p})$, gives us that $\neg K\mathbf{p}$, because $\neg(K\mathbf{p} > K\mathbf{p})$.)

Sceptical consequences aside, (A) is not very appealing. Replacing 'β' with '**Kp**', and assuming that $K(K\mathbf{p} \rightarrow K\mathbf{p})$, we get the "KK" principle: $K\alpha \rightarrow KK\alpha$, against which there are compelling arguments.³⁴ And what's even worse, (A) implies that if one knows *p*, one knows that one knows *p* *antecedently* to knowing *p*! That just cannot be right.³⁵

Of course, this just mires us deeper in paradox if (A) simply records "a general procedure we recognize and insist on in making and assessing knowledge-claims in everyday and scientific life" (30–1). So, is that true?

One of Stroud's main pieces of evidence for something like (A) is his "Cleveland" example: "As a member of a jury I might find out that I have been ruling out one suspect in my mind because he was a thousand miles away, in Cleveland, at the time of the crime" (3). He continues:

If I come to suspect that all the witnesses have conspired and made up a story about the man's being in Cleveland that night, for example, and their testimony is all I have got to go on in believing that he was in Cleveland, I might find that I no longer know whether he was there or not until I have some reason to rule out my suspicion. If their testimony was all invented I would not know that the man was in Cleveland... [even if] he was there all the time. (25)

If this story is elaborated in a natural way, it is quite plausible that Stroud does not know that the man was in Cleveland, even if the witnesses have testified truthfully. In order for Stroud to know that the man was in Cleveland, he needs some independent reason for thinking that the witnesses have not conspired. However, the fact that the witnesses are testifying at a *trial* is not incidental, but a crucial piece of detail. It is easy to imagine that Stroud has good reason for his suspicion that a conspiracy is taking place, and it is this that underwrites the intuition that he does not know that the man was in Cleveland. If we change the example to an ordinary case of testimony, with no suspicious circumstances, then the negative verdict is considerably less compelling. If in some everyday setting Stroud asks a reliable student where Mr. N.N. is, and believes the truthful answer that Mr. N.N. is in Cleveland, then it is hardly intuitively obvious that Stroud's belief counts as knowledge only if Stroud *knows*, let alone knows *antecedently*, that the student is not guessing, etc. (Note that the student's guessing is not incompatible with Mr. N.N.'s *being* in Cleveland, only incompatible with Stroud's *knowing* that he is.) Again, by ordinary standards Stroud is not prevented from knowing that the date of the Battle of Fredericksburg is 1862, even if he has no memory of the history lesson in which he learnt it, and was certainly in no position at the time to know antecedently that his high school teacher wasn't dissembling. It simply isn't true that in making and assessing knowledge-claims in everyday life, we hold knowers to Stroud's principle (A).

It might be replied that the example of the student's testimony can be turned to the sceptic's advantage. She may claim that Stroud's knowledge that Mr. N.N. is in Cleveland is underwritten, not by the assumption that *the student's* testimony is reliable, but by the assumption that testimony is *in general* reliable. Surely *all* our knowledge based on testimony is underwritten by the assumption that our informants are generally reliable, or something to that effect. And if so, then presumably some similar assumption underwrites our knowledge based on "the testimony of the senses".

This reply is best thought of as attempting to make Descartes' Condition plausible, without deriving it from some more general principle like (A). However, it is not very persuasive. As has been noted many times, the claim that one must know that one's informants are generally reliable *antecedently* to knowing their testimony has little support from our everyday practice. By ordinary standards, a child who is in no position to know about the

general reliability of informants at all, is not thereby prevented from knowing that it's snowing on the basis of testimony. This is *not* to propound a controversial philosophical theory of testimony (say, that testimony always supplies defeasible reason for belief). The point is merely defensive: the analogy with testimony does not help establish that one must know that one's perceptual faculties are reliable antecedently to knowing anything on the basis of perception.^{36,37}

5. Stroud on “understanding human knowledge in general”

The dominant theme in Stroud's most recent writings on scepticism is that the philosophical project of “understanding human knowledge in general” is doomed to disappointment. Sceptical arguments have a role to play in establishing this conclusion, but Stroud does not rely on the assumption that these arguments are *sound*, just that they are valid—with possibly dubious premises. In the light of the foregoing, it will be instructive to examine Stroud's case for the distressing conclusion that “there is and could be no satisfactory answer” to the “distinctively philosophical problem of knowledge” (2000, xiv).

What is this “distinctively philosophical problem of knowledge”?

We want to understand how any knowledge at all is possible—how anything we currently accept amounts to knowledge. Or, less ambitiously, we want to understand with complete generality how we come to know anything at all in a certain specific domain. (Stroud 1989, 101)

And the “specific domain” of special relevance is of course the “external world”, or “material bodies around us” (101). Given that the understanding we seek is of *all* knowledge in a certain domain,

it will seem that any other knowledge that might be relevant could not be allowed to amount to already knowing something in the domain in question. Knowledge of anything at all in that domain is what we want to explain, and if we simply assume from the outset that the person has already got some of that knowledge we will not be explaining all of it. (103)

One strategy, then, is to explain how we have knowledge of the target domain “on the basis of another, prior kind of knowledge that does not imply or presuppose any of the knowledge we are trying to explain” (104). So, in the case of the external world, we might

single out epistemically prior ‘sensations’ or ‘sense data’ or ‘experiences’ or whatever it might be, and then ask how on that basis alone, which is the only basis we have, we can know anything of the objects around us. (105)

This is the first horn of a dilemma. And it is at this point that sceptical arguments are brought in to show

that if we really were in the position the traditional account in terms of epistemic priority describes us as being in, scepticism would be correct. (106)

The first horn, then, leads to scepticism. And here there is a complication in Stroud's presentation of the argument. According to Stroud, the reasoning that extracts scepticism from the traditional account turns on the "internalist" claim that if someone knows p on the basis of e , then this person must know that e is good evidence for p . Stroud then considers the "externalist" reply, according to which "[a] person who knows something does not himself have to know that what he has got in his prior evidential basis amounts to knowledge in the domain in question" (112). However, let us ignore this bit of the dialectic, and concede that we cannot explain our knowledge of the external world "on the basis of another, prior kind of knowledge", internalism or no.

The problem with the second horn of the dilemma might now seem immediate. If the explanation of our knowledge of the external world is not in terms of some *other* kind of knowledge, then the explanation will have to assume "that we have got some knowledge in the domain in question", in which case "we have lost the prospect of explaining and therefore understanding all of our knowledge with complete generality" (121).

Not so fast, though. Suppose that we have an account of knowing propositions about the external world that gives at least a reductively specified sufficient condition: a person S knows p if $C(S, p)$ (for instance—to take an example Stroud mentions—a "reliabilist" theory along the lines of Dretske 1981). Since we are following the second horn, this reductive account will not avert to another kind of knowledge (of "sensations", for instance). The *explanandum* is that S knows p . The *explanans* is that $C(S, p)$, which hypothesis is couched in terms that can be understood without presupposing that S himself knows p , or indeed has any knowledge of the external world at all. That is, we *aren't* assuming "from the outset that the person has already got some of that knowledge". So what's the problem?

It turns out that there *isn't* one, provided S is someone else: "It is not that there is any difficulty in understanding other people's knowledge in those terms. It is only with self-understanding that the unsatisfactoriness or loss of complete generality makes itself felt" (120).

Thus the second horn only gets uncomfortable when $S = \text{me}$. Stroud admits that it is "not easy to describe the deficiency in a few words" (1994, 148); on his official explanation of the deficiency, if I hold a theory of the sort mentioned two paragraphs back, then I can "say at most: 'If the theory I hold is true, I do know or have good reason to believe that I know

or have good reason to believe it, and I do understand how I know the things I do” (151).

“I think”, Stroud continues,

that . . . we can see a way in which the satisfaction the theorist seeks in understanding his knowledge still eludes him. Given that all of his knowledge of the world is in question, he will still find himself able to say only ‘I might understand my knowledge, I might not. Whether I do or not all depends on how things in fact are in the world I think I’ve got knowledge of.’ (151)

The point seems to be this. According to my theory, if $C(S, p)$ then S knows p . Hence, if $C(\text{Me}, \text{that I have a hand})$, then I know that I have a hand. The antecedent of this conditional expresses some complicated contingent condition, perhaps involving reliable connections between my head and my hand. If this condition does *not* obtain, then I may yet know that I have a hand—my theory only gives a *sufficient* condition for knowledge. However, since I have no *other* account of my knowledge, if it’s false that $C(\text{Me}, \text{that I have a hand})$, then I don’t “understand” my knowledge that I have a hand. So, do I or do I not understand my knowledge? It all depends on whether $C(\text{Me}, \text{that I have a hand})$. If I can assure myself that this condition obtains, then I can say that I really do understand my knowledge that I have a hand. And here the fact that I am the *subject* of my own investigation is relevant. For how can I assure myself that $C(\text{Me}, \text{that I have a hand})$? Evidently empirical investigation is the only way. Very well, I do some of that, and come to know that this condition obtains. Have I then arrived at the understanding I seek? No, because I want to understand *all* my knowledge of the external world. When I assert that $C(\text{Me}, \text{that I have a hand})$, I realize that I am presupposing or taking for granted that I *know* this external world proposition: if I didn’t know it, I would not be entitled to assert it.³⁸ So I can only assure myself I understand my knowledge that I have a hand if I help myself to another piece of *my* knowledge of the external world. And that’s exactly what I *don’t* want: I want to understand my knowledge of the external world without helping myself to any part of it. (Although Stroud chose the first-person singular to illustrate the difficulty, the problem could be put equally well in the first-person plural: *we* want to understand *our* knowledge without helping ourselves to any part of it.)

So, the demand to understand our knowledge of the external world presents us with a dilemma. One horn leads to scepticism; on the other horn, although we do have some knowledge of the external world, we can’t achieve the proper philosophical understanding of it. Perhaps, Stroud conjectures, the proper moral is that the “epistemological enterprise is not fully valid, or perhaps not even fully coherent”, in which case there is something wrong with the demand that leads to the dilemma. Disappointment all round.

Notice that, on this interpretation of Stroud, he is *not* saying that I am not entitled to assert that I know that I have hand—that I can only make the *conditional* assertion that *if C*(Me, that I have a hand), *then* I know that I have a hand. Admittedly, he does appear to make this claim in a number of places (see the quotations above from Stroud 1994, 151; see also 1989, 118–9). But appearances must be misleading, otherwise Stroud would have made a serious mistake (as is pointed out in Sosa 1997, 283); and in any case, the interpretation just offered is independently plausible.

Now all this should seem familiar. We have met Stroud’s project of “understanding human knowledge in general” before—in fact, it is nothing less than the task of *convincing the sceptic*. For recall Stroud’s first strategy, and suppose that it works: we can show how we have knowledge of the external world on the basis of some other kind of knowledge—of “sensations”, for example. Then we would, presumably, have a sound argument with the conclusion that we have knowledge of the external world, none of whose premises are external world propositions. This is, of course, exactly what is required to convince the sceptic. And what was supposed to be wrong with the second strategy? Precisely that our attempt to “understand” our knowledge of the external world appeals to an *external world premise*—something that the sceptic will dispute as begging the question against her.

If the project of “understanding human knowledge in general”, the “distinctively philosophical problem of knowledge”, boils down to the task of convincing the sceptic, then one might wonder why the project’s failure is such a big deal. Let us grant that we can’t beat the sceptic with both hands tied: is this really much cause for “dissatisfaction”?

The odd thing is that there is certainly a case to be made on Stroud’s behalf. Grant that there is no cogent sceptical *argument*. Still, there *does* seem to be something worrying about sceptical scenarios. In certain philosophical moods, sceptical doubts are hard to fight off—an argument that will convince the sceptic seems to be the only cure. As Guyer puts it on behalf of Kant: “an answer to skepticism must lie in what reflection on his own thoughts can reveal to the individual thinker even on the supposition that nothing but his own consciousness exists”.³⁹ However, if the drift of this paper is correct, the urge to convince the sceptic itself poses something of a paradox. One *does* know that one has hands and that one isn’t dreaming—further, one knows that one knows these things. The sceptic can offer nothing to threaten these obvious facts. And yet, somehow, one does feel threatened—even after cleansing oneself of confusion. This is one paradox that this paper does not attempt to solve.⁴⁰

Notes

¹ See Cavell 1979; Clarke 1972; Fumerton 1995; Greco 2000; C. McGinn 1989, 113–4; M. McGinn 1989; Nagel 1986, 67–74; Stroud 1984a; Williams 1996a; and Wright 1985. These

philosophers differ widely in their views, but they all profess some sympathy with claims like the following two: “*At no point does the [sceptical] argument require us to accept anything that is either obviously false or even open to doubt*” (M. McGinn 1989, 7, her emphasis); “skeptical arguments are powerful in the following sense: it is not easy to see where they go wrong, and rejecting them requires one to adopt substantive and controversial theses about the nature of knowledge and evidence” (Greco 2000, 2). Fumerton, C. McGinn, Stroud, and Nagel are to varying degrees in the pessimistic camp. Cavell, Clarke, Greco, M. McGinn, Williams, and Wright offer diagnoses of the sceptic’s mistake, which typically turns out to be extremely subtle. Williams, for instance, thinks that the sceptic mistakenly assumes “epistemological realism”: “our beliefs [fall] into a hierarchy of *natural epistemological kinds* ordered by natural relations of epistemic priority” (1996b, 397; see also 1996a, ch. 3). To take one more example, Wright thinks the sceptic mistakenly assumes that “hinge propositions” (e.g. that there is an external world) need justification at all; drawing on Wittgenstein’s *On Certainty*, Wright argues that such propositions are neither justified nor unjustified, because they are not “fact-stating”.

Both C. McGinn and Wright modified their earlier positions (McGinn 1993, Wright 1991), but remained impressed with the power of sceptical reasoning. In his 1989 McGinn’s view was that scepticism may well be correct; in his 1993 he conjectures that “scepticism is false but unknowably so” (117). In his 1985 Wright gave an argument for the sceptical conclusion that one doesn’t know one isn’t dreaming (433; see also section 4 below), and apparently later downgraded his assessment of the argument’s force (1991, 91–3).

Although (as will be argued) the works cited above overreact to the sceptical paradoxes, it needs emphasizing that in doing so they have produced many philosophically important spin-offs, as the space program produced the non-stick saucepan.

²Wright 1985, 429.

³Thus Lycan: “It is hard to find a sound argument taking us from any common theory of epistemic justification, or for that matter any plausible set of premises, to a seriously skeptical conclusion” (1988, 197).

⁴Sometimes the sceptical argument is explicitly directed against justification, not knowledge (e.g. Pryor 2000, Wright 1985). If knowledge entails justification, then any sceptical argument against justification is *a fortiori* an argument against knowledge. It will be assumed here that knowledge entails justification (for a recent denial see Lewis 1996). Accordingly, if the conclusion of this paper is correct, then sceptical arguments against justification are also flawed.

⁵See also note 1 above.

⁶Convince-the-sceptic comes in generic and specific variants: convince the sceptic that we have some knowledge or other of the external world (generic); convince the sceptic that we know *p*, for various external world fillings for ‘*p*’ (specific). We can afford to slur over the distinction here.

⁷An Audian refutation is something like Pryor’s “modest anti-skeptical project”: “to establish to *our* satisfaction that we can justifiably believe and know such things as that there is a hand” (2000, 517)—insisting that one does know such things is apparently insufficient, but we are allowed to appeal to assumptions that the sceptic won’t accept. (Pryor also mentions an “ambitious anti-skeptical project”, which corresponds to our last game of convince-the-sceptic; the prospects for success here, he sensibly says, “seem somewhat dim” (517).)

⁸The parenthetical quotations are from Austin 1946, footnote to 113, and Kant 1933, B275 (quoted by Stroud, 51, 114, respectively).

⁹Stroud himself is not under any confusion about the distinction between exposing the mistake in a sceptical paradox and showing that its conclusion is false (see 1984a, 139–140). This makes the structure of his book all the more puzzling.

¹⁰For examples of the curious interpretation of Strawson, see Sosa 1988, 160–3; Williams 1992, ch. 1; Wright 1985, 429, fn. 2, and 1990, 87, fn. 2. I am not claiming that any of these philosophers have arrived at their interpretation by overlooking the distinction just mentioned. (And it must be admitted that some passages in Strawson’s book do support the curious

interpretation.) All I am saying is that overlooking the distinction makes the curious interpretation seem more plausible than it ought to be. A recent exchange between Sosa and Strawson provided an opportunity for this exegetical matter to be cleared up—unfortunately, as far as I can see, it wasn't. Sosa complains, against what he takes to be Strawson's view, that "if we do see the skeptic as a purveyor of paradox, then we might surely take his arguments seriously without needing to produce new arguments in favor of our favored paradox-enmeshed propositions. Instead we might take his arguments seriously by examining them for defects" (1998, 366). In his reply, Strawson says that "[I]t is not merely a matter of dismissing the demand for a justification of one's belief in a proposition on the ground that one can't help believing it. That would be weak indeed. The position is, rather, that the demand for justification is really senseless" (1998, 370). This doesn't address Sosa's point: his sceptic does not insist that we bear the burden of proof, but rather produces an *argument* that we don't know anything.

¹¹In fact, the closure principle as stated is not quite correct, because one may know *p*, know that *p* entails *q*, and yet fail to put these two pieces of knowledge together and infer *q*. A better version can be obtained by replacing 'one knows/does not know' in the canonical sceptical argument and the closure principle with 'one is in a position to know/not in a position to know'. Similar repairs to various principles discussed in section 4 below can be left tacit.

¹²For a recent attack on this premise, see Levin 2000.

¹³Here I am joining Feldman (1999), and Greco (2000, 52, fn. 16).

¹⁴There are two other sceptical arguments that need at least a brief mention. First, the ancient argument from an infinite regress of reasons (see Williams 1996a, 60–68; 1999): if *p* is justified, then it is justified on the basis of a distinct justified proposition *q* (else *p* would be a mere assumption); repeating the argument for *q*, and assuming that the chain of justification cannot go on for ever or loop back in a circle, no proposition is justified, and hence no proposition is known. This is an argument for *total* scepticism, not just *external world* scepticism, and so is not discussed here.

Second, an argument for external world scepticism that appeals to the third of Nozick's proposed four individually necessary and jointly sufficient conditions for *S* to know α : if α were false, *S* wouldn't believe α (1981, 172). By this condition, *S* does not know not- \spadesuit ; since *p* (e.g., that *S* has a hand) entails not- \spadesuit , by closure, *S* does not know *p*. (Of course, Nozick himself resists the final step of this argument.) Nozick revises this condition a few pages later in the face of a counterexample (179), and on closer examination (provided by, among others, Vogel (1987) and Williamson (2000, ch. 7)) there are serious objections to any revision that yields sceptical consequences.

¹⁵The reformulated sceptical argument, not relying on the premise that if one knows *p* then one knows not- \spadesuit , is:

- (1) If one's evidence does not favor *p* over \spadesuit , then one does not know *p*.
- (2) One's evidence does not favor *p* over \spadesuit .

Therefore:

- (3) One does not know *p*.

¹⁶According to Russell, knowledge "is a matter of degree" (1948, 427). Russell appears to be only arguing in the quoted passage that he does not know "to some high degree" that there is a church.

¹⁷For some of these attempts, see Jackson 1977, 141–53; Russell 1912, ch. 2; Slote 1970; Vogel 1990a. For criticism, see Enç 1990; Williams 1977, ch. 4; 1999, 48–9; Wright 1985, 446–9. Cf. Kant 1933, A367.

¹⁸For reviews of foundationalism's recent resurgence, see Bonjour 1999 and Pryor 2001. Modest versions of foundationalism are endorsed by Alston (1989), Audi (1993), Bonjour (1999), and Pryor (2000).

¹⁹The sceptic is not assuming that *S*'s *total* evidence for *p* is *e*, only that *S* would not know *p* without *e* as evidence. Connectedly, the sceptic is not assuming the falsity of Williamson's thesis that all and only knowledge is evidence (see Williamson 2000, chs. 7, 8).

²⁰ Often the nature of propositions like *e* is given rather cursory treatment, as if the details are inconsequential. For example, according to Fumerton, *e* is a proposition “about the phenomenological character of [*S*’s] subjective and fleeting sensations” (1995, 32); according to Williams, *e* is a proposition about “sensory appearances” (1996a, 73); according to Wright, *e* is “some proposition describing in appropriate detail [*S*’s] total field of experience” (1985, 435). Stroud is particularly cagey about *e*: he tells us that “[w]hat we gain through the senses is on Descartes’s view only information that is compatible with our dreaming things about the world around us and not knowing anything about that world” (1984a, 12), but this “information” is not explicitly characterized further (see also Stroud 1984b, 5–6).

²¹ Jackson 1977 is a classic recent defense of the representative theory of perception, although Jackson later changed his mind. One of the prominent current defenders is Robinson (1994).

²² In the paper cited Johnston endorses this comparison; in later papers (e.g., 1998) he repudiates it.

²³ Another component is that objects like tomatoes are perceived “indirectly” or “mediately”. Two excellent discussions of this obscure claim are in Dretske 1969, 62–75; and Jackson 1977, ch. 1.

²⁴ For some examples of recent accounts of perception that broadly endorse this basic idea, but which radically differ in other ways, see Burge 1986, Dretske 1995, Johnston 1998, McDowell 1994, Peacocke 1992, and Searle 1983. Peacocke is an interesting case: he holds that there are sense-data-like “regions of the visual field” (see Peacocke 1983), and to this extent endorses the representative theory of perception. However, he certainly does not think that one’s perceptual evidence concerns these regions, and so repudiates the Morse code model.

²⁵ See Pollock 1971, 301; Plantinga 1993, 93; Greco 2000, 95–6; Pryor 2000, 536. Cf. Jackson 1977, 152.

²⁶ Another line of argument for “semi-skepticism” (Vogel’s terminology) is discussed in his 1990b.

²⁷ As mentioned earlier, Russell takes *S*’s perceptual evidence *e* to be an experience, not a proposition. In the light of the failure of the argument from underdetermination with a propositional *e*, it is worth briefly revisiting Russell’s suggestion.

Since speaking of non-propositional “evidence” can only cloud the issue (see Williamson 2000, ch. 9), Russell’s suggestion is best recast in some alternative terminology. Let us speak instead of “grounds”. A *ground* of one’s belief, let us say, is something that stands in a justification-conferring relation to it. (Note that one need not be *aware* of a belief’s grounds.) Thus the experience as of a hand might be the ground of the belief that one has a hand when the experience bears a certain kind of causal relation to the belief. For the sake of the argument let us concede that our external world knowledge requires experiential grounds: in particular, if *S*’s belief *p* does not have experience *E* as a ground, then *S* does not know *p*. Can the sceptic respond to the above criticisms by revising the argument from underdetermination in terms of grounds?

It is hard to see how. Certainly, *p* is “underdetermined” by *E* in this sense: the proposition that *E* occurs does not entail *p* (further, *p* is not entailed by the stronger proposition that *E* is a ground of *S*’s belief *p*). However, the sceptic will not suppose that the proposition that *E* occurs is *S*’s *evidence* for *p*—that would be to revert to the original formulation of the argument from underdetermination, with all its problems. Hence, the sceptic cannot fairly compare our predicament with those of the detectives who have overlooked a rival explanation of the (propositional) evidence; she thus loses the main way of motivating the sceptical conclusion. Of course, one might wonder about the nature of the justification-conferring relation between experience and belief, or question whether a particular candidate relation could do the trick. To the extent that this suggests the materials for a sceptical paradox immune to the objections raised so far, it is the one examined in the following section.

²⁸ See also Stroud 1989, 106ff.

²⁹ As Ayer notes (1956, ch. 2).

³⁰This extends Klein's terminology slightly, which officially applies to justification rather than knowledge.

³¹Cf. the connective '⇒' in Sosa 1988, 158–9.

³²See Van Cleve 1979, 242. The other half is the converse: If I know that God exists and is not a deceiver then I *antecedently* know that whatever I perceive clearly and distinctly is true.

³³See Pryor 2000, 529–30, to which I am much indebted. This interpretation is also briefly given in Sosa 1988, 158–9, and is strongly hinted at in Wright 1985 (see 433 and the footnote to 430).

³⁴See Williamson 2000, ch. 5; Stroud himself seems undecided whether the KK principle is correct (1989).

³⁵Stroud recognizes that (A) is problematic, as the above quotation continues: "I will not speculate further on the qualifications or emendations required to make the principle less implausible". See also note 36 below.

³⁶On testimony, see Coady 1992. Stroud himself emphasizes these sorts of points about testimony and Descartes' Condition when he discusses Austin (1984a, ch.2). He even admits that "[i]t seems clear that we simply do not insist on fulfilling Descartes's condition in order to know things in real life" (50). Stroud does not see this as casting doubt on the sceptical argument presented in the first chapter; this is presumably because chapter 2 marks an important shift in the dialectical structure of the book. Chapter 2 is not devoted to *exposing* the sceptic, by questioning whether there is any support for Descartes' Condition. Rather, its main purpose is to rebut the following attempt to *convince* the sceptic: "anyone who raises the [dreaming] possibility in normal circumstances [is] necessarily violating or rejecting the everyday meaning of the word 'know'" (57).

³⁷There is an alternative (and, to my mind, less plausible) interpretation of Stroud's sceptical argument that makes no use of knowing one thing antecedently to knowing another. However, it does make for a completely straightforward construal of various passages, and with the notation explained above may be set out as follows:

Principle (A') (cf. the previous quotations from Stroud): if one knows that α and knows that one's knowing that α is incompatible with it's being false that β , then one knows that β :

$$(A') \quad [K\alpha \ \& \ K(K\alpha \rightarrow \beta)] \rightarrow K\beta$$

Hence:

- (1) $[Kp \ \& \ K(Kp \rightarrow \neg d)] \rightarrow K\neg d$
- (2) $\neg K\neg d$ [Premise]
- (3) $\neg[Kp \ \& \ K(Kp \rightarrow \neg d)]$ [From(1), (2)]
- (4) $K(Kp \rightarrow \neg d)$ [Premise]
- (5) $\neg Kp$ [From (3), (4)]

(1)–(5) is the "Complex Argument" of Sosa 1988. This interpretation of Stroud is due to Wright (1990, 90–2). Setting aside the obvious problems with (A'), the argument needs the disputable premise (2) (see the discussion of the canonical sceptical argument in section 2 above).

³⁸In support of this step, see Williamson 2000, ch. 11.

³⁹Guyer 1980, 161; this remark is quoted (approvingly) by Levin (2000, 424).

⁴⁰Thanks to Juan Comesaña, Richard Feldman, Michael Glanzberg, Ned Hall, and Ralph Wedgwood. I am especially grateful to Jim Pryor for many helpful conversations.

References

- Alston, William. 1989. *Epistemic Justification*. Ithaca: Cornell University Press.
 Audi, Robert. 1993. *The Structure of Justification*. Cambridge: Cambridge University Press.

- Austin, J. L. 1946. "Other minds." *Proceedings of the Aristotelian Society* Supp. Vol. 20. Reference to the reprinting in Austin, *Philosophical Papers*, Oxford: Oxford University Press, 1979.
- Ayer, A. J. 1956. *The Problem of Knowledge*. Harmondsworth, England: Penguin.
- Bonjour, Laurence. 1999. "The dialectic of foundationalism and coherentism." Greco and Sosa 1999.
- Brueckner, Anthony. 1994. "The structure of the skeptical argument." *Philosophy and Phenomenological Research* 54: 827–35.
- Burge, Tyler. 1986. "Cartesian error and the objectivity of perception." *Subject, Thought, and Context*, eds. P. Pettit and J. McDowell. Oxford: Oxford University Press.
- Cavell, Stanley. 1979. *The Claim of Reason*. Oxford: Oxford University Press.
- Clarke, Thompson. 1972. "The legacy of skepticism." *Journal of Philosophy* 69: 754–69.
- Coady, C. A. J. 1992. *Testimony*. Oxford: Oxford University Press.
- Cohen, Stewart. 1998. "Two kinds of skeptical argument." *Philosophy and Phenomenological Research* 58: 143–59.
- DeRose, Keith. 1999. "Contextualism: an explanation and defense." Greco and Sosa 1999.
- Dretske, Fred. 1969. *Seeing and Knowing*. London: Routledge & Kegan Paul.
- Dretske, Fred. 1970. "Epistemic operators." *Journal of Philosophy* 67: 1001–23.
- Dretske, Fred. 1981. *Knowledge and the Flow of Information*. Oxford: Blackwell.
- Dretske, Fred. 1995. *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- Eng, Berent. 1990. "Is realism really the best hypothesis?" *Journal of Philosophy* 87: 667–8.
- Feldman, Richard. 1999. "Contextualism and skepticism." *Philosophical Perspectives* 13: 91–114.
- Fumerton, Richard. 1995. *Metaepistemology and Skepticism*. Lanham, MD: Rowman & Littlefield.
- Greco, John. 2000. *Putting Skeptics in Their Place: the nature of skeptical arguments and their role in philosophical inquiry*. Cambridge: Cambridge University Press.
- Greco, John, and Ernest Sosa, ed. 1999. *The Blackwell Guide to Epistemology*. Oxford: Blackwell.
- Guyer, Paul. 1980. Review of D. Heinrich, *Identität und Objektivität*. *Journal of Philosophy* 76: 151–67.
- Jackson, Frank. 1977. *Perception*. Cambridge: Cambridge University Press.
- Johnston, Mark. 1992. "How to speak of the colors." *Philosophical Studies* 68: 221–63.
- Johnston, Mark. 1998. "Are manifest qualities response-dependent?" *The Monist* 81: 3–43.
- Kant, Immanuel. 1933. *Critique of Pure Reason*. Trans. N. Kemp Smith. London: Macmillan.
- Klein, Peter. 1981. *Certainty: A Refutation of Scepticism*. Minneapolis: University of Minnesota Press.
- Levin, Michael. 2000. "Demons, possibility and evidence." *Noûs* 34: 422–40.
- Lewis, David. 1996. "Elusive knowledge." *Australasian Journal of Philosophy* 74: 549–57. Reference to the reprinting in Sosa and Kim 2000.
- Lycan, W. G. 1988. *Judgement and Justification*. Cambridge: Cambridge University Press.
- McDowell, John. 1994. *Mind and World*. Cambridge, MA: Harvard University Press.
- McGinn, Colin. 1989. *Mental Content*. Oxford: Basil Blackwell.
- McGinn, Colin. 1993. *Problems in Philosophy*. Oxford: Blackwell.
- McGinn, Marie. 1989. *Sense and Certainty*. Oxford: Basil Blackwell.
- Nagel, Thomas. 1986. *The View From Nowhere*. Oxford: Oxford University Press.
- Nozick, Robert. 1981. *Philosophical Explanations*. Oxford: Oxford University Press.
- Peacocke, Christopher. 1983. *Sense and Content*. Oxford: Oxford University Press.
- Peacocke, Christopher. 1992. *A Study of Concepts*. Cambridge, MA: MIT Press.
- Plantinga, Alvin. 1993. *Warrant and Proper Function*. Oxford: Oxford University Press.
- Pollock, John. 1971. "Perceptual knowledge." *Philosophical Review* 80: 287–319.
- Pryor, James. 2000. "The skeptic and the dogmatist." *Noûs* 34: 517–49.
- Pryor, James. 2001. "Highlights of recent epistemology." *British Journal for the Philosophy of Science* 52: 95–124.

- Robinson, Howard. 1994. *Perception*. London: Routledge.
- Russell, Bertrand. 1912. *The Problems of Philosophy*. New York: Henry Holt.
- Russell, Bertrand. 1948. *Human Knowledge: Its Scope and Limits*. New York: Simon and Schuster.
- Searle, John. 1983. *Intentionality*. Cambridge: Cambridge University Press.
- Slote, Michael. 1970. *Reason and Scepticism*. London: George Allen & Unwin.
- Sosa, Ernest. 1988. "Beyond scepticism, to the best of our knowledge." *Mind* 97: 153–88.
- Sosa, Ernest. 1997. "Reflective knowledge in the best circles." *Journal of Philosophy* 94: 410–30. Reference to the reprinting in Sosa and Kim 2000.
- Sosa, Ernest. 1998. "P. F. Strawson's epistemological naturalism." *The Philosophy of P. F. Strawson*, ed. L. E. Hahn. Chicago: Open Court.
- Sosa, Ernest, and Jaegwon Kim, ed. 2000. *Epistemology: An Anthology*. Oxford: Blackwell.
- Strawson, P. F. 1985. *Scepticism and Naturalism: Some Varieties*. New York: Columbia University Press.
- Strawson, P. F. 1998. "Reply to Ernest Sosa." *The Philosophy of P. F. Strawson*, ed. L. E. Hahn. Chicago: Open Court.
- Stroud, Barry. 1984a. *The Significance of Philosophical Scepticism*. Oxford: Oxford University Press.
- Stroud, Barry. 1984b. "Scepticism and the possibility of knowledge." *Journal of Philosophy*: 545–51. Reference to the reprinting in Stroud 2000.
- Stroud, Barry. 1989. "Understanding human knowledge in general." *Knowledge and Scepticism*, ed. M. Clay and K. Lehrer, Boulder: Westview Press. Reference to the reprinting in Stroud 2000.
- Stroud, Barry. 1994. "Scepticism, 'externalism', and the goal of epistemology." *Proceedings of the Aristotelian Society* Supp. Vol. 68: 291–307. Reference to the reprinting in Stroud 2000.
- Stroud, Barry. 2000. *Understanding Human Knowledge*. Oxford: Oxford University Press.
- Unger, Peter. 1975. *Ignorance*. Oxford: Oxford University Press.
- Van Cleve, James. 1979. "Foundationalism, epistemic principles, and the Cartesian circle." *Philosophical Review* 88: 55–91. Reference to the reprinting in Sosa and Kim 2000.
- Vogel, Jonathan. 1987. "Tracking, closure, and inductive knowledge." *The Possibility of Knowledge*, ed. S. Luper-Foy, Totowa, NJ: Rowman & Littlefield.
- Vogel, Jonathan. 1990a. "Cartesian skepticism and inference to the best explanation." *Journal of Philosophy* 87: 658–66.
- Vogel, Jonathan. 1990b. "Are there counterexamples to the closure principle?" *Doubting*, ed. D. Roth and G. Ross. Dordrecht: Kluwer.
- Williams, Michael. 1977. *Groundless Belief*. Oxford: Basil Blackwell.
- Williams, Michael. 1996a. *Unnatural Doubts*. Princeton: Princeton University Press.
- Williams, Michael. 1996b. "Understanding human knowledge philosophically." *Philosophy and Phenomenological Research* 56: 359–78.
- Williams, Michael. 1999. "Scepticism." Greco and Sosa 1999.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wright, Crispin. 1985. "Facts and certainty." *Proceedings of the British Academy* 71: 429–72.
- Wright, Crispin. 1991. "Scepticism and dreaming: imploding the demon." *Mind* 100: 87–116.