

Skepticism as a Theory of Knowledge

Skepticism about the external world may very well be correct, so the question is in order: what theory of knowledge flows from skepticism itself? The skeptic can give a relatively simple and intuitive account of knowledge by identifying it with indubitable certainty. Our everyday 'I know that p' claims, which typically are part of practical projects, deploy the idea of knowledge to make assertions closely related to, but weaker than, knowledge claims. Roughly, we are asserting that we know p 'for all practical purposes.' The truth of such claims is consistent with skepticism; various other vexing problems don't arise. In addition, even if no claim about the world outside my mind can be more probable than its negation, the project of pure scientific research remains well motivated.

Epistemology is largely a response to skepticism. A subtext of virtually every theory of knowledge has been to show how knowledge is possible or, at the least, to avoid an account that delivers us unto the skeptic. Yet skepticism remains robust after 2,500 years. The question is surely in order: what if skepticism is correct? How bad would that be? And what sort of epistemology flows from skepticism? A tortured epistemology has resulted from the conviction that a theory of knowledge must explain how knowledge is possible. If its consequences turn out not to be so terrible as we feared, and it yields an intuitive and useful account of knowledge, then skepticism itself becomes an attractive option in epistemology. This paper will support that alternative. In any case, it is unlikely that skepticism should have proved so durable if it reveals nothing of importance about knowledge; and we can hardly expect to learn what something teaches so long as we doggedly resist it. To learn what it teaches, therefore, I propose to capitulate to skepticism.

The skepticism that will concern me is three-fold: the Problem of the External World, the Problem of Induction, and the

Problem of the Past--we don't know there is a world outside our minds, we know nothing about the future, and we don't know there is a past. (The area of traditional epistemology these constitute might be called 'The Problem of the External World and Its Suburbs.') I focus on these problems because each of them is of particular concern to science, and because all are motivated by forceful arguments.¹ Skepticism remains unalarming, of course, if it merely insists that we don't know these things. We might gladly give up the word 'knowledge,' it is often noted, so long as we can keep 'probable belief.' Skepticism emerges in its full horror, and is far more interesting, when it insists (again, motivated by forceful arguments) that no belief belonging to these domains is warranted--none can be more probable than its negation. The skepticism to which I will capitulate, therefore, is Industrial Strength.

I

Once epistemology is relieved of the task of rescuing knowledge from skepticism, it becomes relatively easy to give an account of knowledge:² knowledge is indubitable certainty, as Descartes believed. That is, S knows that p if and only if S believes p truly and, given S's warrant, rational doubt that p is impossible.³ Our theory of knowledge will be internalistic, so that knowledge, if we had it, would be a rational antidote to doubt and wonder, one of the chief reasons we've always wanted knowledge. This theory doesn't entail we know nothing, however, but only that we know very little. Perhaps we know the sort of things Descartes thought we knew. My attentive belief that there

is a thought now is indubitable, for instance; for it must be true if I doubt it. And when I know there is a thought now, this makes it indubitable that something exists--if only the thought. I may know some purely logical truths as well, for example, that if something has hands, then it isn't a disembodied brain in a vat (a 'BIV,' for short). Nonetheless, to know that p I must be able to rule out every conceivable mistake. In the case of contingent beliefs, on which I will focus, this requires that E (the evidence on the basis of which I believe p) entails the non-actuality of every possible world in which not-p. Otherwise, for all I know, this is one of those worlds; so, for all I know, p is false. This argument is intuitive and forceful, and, once we no longer need to resist the skeptic, we have no reason not to accept it.

Skepticism is often rejected on the ground that it denies the truth of all our ordinary knowledge claims. Keith DeRose writes: "The bold skeptic thus implicates us...in systematic and widespread falsehood in our use, in speech and in thought, of our very common word 'know'."⁴ Worse, the skeptic denies them on account of ridiculous considerations. When I say 'I know there are three chairs at the table, I've counted them twice,' it is absurd to object: 'How do you know you're not a brain in a vat?' It is worth emphasizing, therefore, that it is not a consequence of an enlightened skepticism that such utterances are false, but only that they make, not claims to knowledge, but a closely related, weaker sort of statement. The skeptic can explain the utility of the idea of knowledge in our ordinary practices, even though we never have knowledge. Knowledge functions there as an

idealization. When involved in practical pursuits, we bracket skeptical possibilities; naturally we assume we aren't BIVs-- otherwise the practical problems we're addressing wouldn't arise. I use the sentence 'I know there are three chairs at the table' to make the statement that, supposing various odd possibilities conflicting with my theory of the world are ruled out, I know there are three chairs. My epistemic relation to that proposition would attain the ideal if I could foreclose those alternatives.

More precisely, the form of the proposition I assert by my everyday 'I know that p' utterance is: 'I believe p truly and, supposing not-p worlds that conflict with my theory of the world are already ruled out, then, given E, it is indubitably certain that p.' In the midst of practical pursuits, I routinely use sentences of the form 'I know that p' to assert, not that p is indubitably certain, but this closely related, weaker proposition. The literal meaning of the sentence 'I know that p,' namely, that I know that p, becomes a consequent in the compound proposition I actually assert. My knowing p would consist of the fact that I can rule out every not-p world; my everyday 'I know that p' claim (that is, the statement I make by uttering a sentence of the form 'I know that p' in an everyday situation) is made true largely by the fact that, supposing worlds that conflict with my theory of the world are ruled out, I can rule out every not-p world.⁵ So, roughly, the force of my utterance is: 'Supposing certain worlds are already ruled out, I know that p.' This centrally involves the idea of knowledge, but, as one can't rule out alternatives by supposing them ruled out, it

doesn't entail that I know that p.

My 'theory of the world' ('WT,' for short) is an epistemic theory, an account of the mistakes I believe I can actually make in judging that p, given the way I believe things are in the particular circumstances in which I judge. A WT is an account of reality put to epistemic purposes. It changes when my beliefs shift so that I come to believe new mistakes can be made or old ones are impossible, and it determines the set of not-p worlds E must rule out for my everyday claim to be true. A WT is a rough and ready instrument conditioned by practical concerns; a functional WT can contain falsehoods, warrantless beliefs about the world, and the list of mistakes I believe I don't have to rule out may not be warranted by my beliefs about how things are, after all.⁶ Our principal requirement of WTs is that, when combined with perception, they enable us to negotiate this sharp-edged world well enough to satisfy our desires. How we acquire WTs, their precise contents, the role of perception, reason, and education in their development and correction, are legitimate topics for naturalized epistemology.⁷ What I want to maintain here is that we often deploy the ideal of knowledge against the background of these rough, practical theories in order to make useful epistemic assertions that are weaker than knowledge claims.⁸

When I now say 'I know there are three chairs at the table,' part of my WT is that I'm not hallucinating--I don't have to rule out that mistake--but also that there aren't in the vicinity holograms of chairs, papier maché facsimiles, and so on. Of course, when I make an everyday 'I know that p' assertion, my WT

is adjusted so that it excludes p; I don't know p 'for all practical purposes' if I assume p is true. E must do some work: the set of not-p worlds consistent with my WT must be non-empty.⁹ Consequently I don't know 'for all practical purposes' that I'm not a BIV, for once my WT is adjusted to exclude the belief that I'm not a BIV, my evidence doesn't rule out all the vat worlds that remain. In addition, such theories automatically incorporate new beliefs making them more demanding, so that E must rule out more worlds.

Further, we impose imperialistically our more demanding WT in evaluating other people's (and our own past) utterances. Learning that Henry has unwittingly driven into a district where many apparent barns are holograms, and that he says 'I know that's a barn' while luckily pointing at a real barn, I revise my WT accordingly. What I state by saying 'Henry knows he is looking at a barn' can be expressed 'Henry believes truly that he is looking at a barn (let this belief=p), and, assuming he isn't in a not-p world that conflicts with my WT, then, given his evidence, it is indubitably certain that he is looking at a barn,' which is plainly false. And the proposition Henry asserts by saying 'I did not know I was looking at a barn,' after he learns of the holograms, is true. It can be expressed: 'Either I did not believe truly that I was looking at a barn, or, setting aside worlds that conflict with my (present) WT, my evidence did not make it indubitably certain that I was looking at a barn.'

Nonetheless, as the indexical 'my theory of the world' latches onto whatever WT the speaker (or thinker) has presently, the proposition Henry actually asserted when he pointed to the

barn was true, in fact, though its truth is irrelevant to us now because it brackets too many worlds.¹⁰ Of course, if I believe the district in which Henry is driving contains holograms, and I hear him say the words 'I know I'm looking at a barn,' I can reply 'That's false' or 'What you say isn't true.' But this should be construed simply as my saying 'It's false that you know you're looking at a barn,' which is true. In fact, Henry and I disagree sharply: we have conflicting WTs. Henry believes his evidence rules out every mistake he can actually make in the perceptual situation in which he judges; I deny that. I respond to his utterance as if it asserts that belief. Roughly, Henry's true statement means: 'I believe that the mistakes consistent with my beliefs about how things are in this situation are the only ones I can actually make, and E rules them out.' As I believe another mistake can be made, one that E doesn't rule out, I respond: 'That's false.' (Compare: Henry says truly that he believes $2+2=5$. I respond: 'That isn't true.')

Our disagreement is expressed quite naturally by everyday 'I know that p' claims. I say: 'No, Henry, you don't know you're looking at a barn--this district is full of holograms that look just like barns from the road.' As soon as I persuade Henry about the holograms, he must say 'I did not know I was looking at a barn'--a judgement he can express by saying 'It's false that I knew I was looking at a barn' or simply 'What I said before wasn't true.'¹¹

Consequently, the truth of the proposition Henry actually asserts when he points to the barn doesn't require that he can rule out holograms. That is, he needn't be able to distinguish a case in which p is true from a case in which p is false in a way

it could easily have been false:¹² only not-p worlds consistent with his WT are relevant. To put the matter more obscurely, the truth of Henry's statement doesn't require that he not believe p in the not-p worlds closest to the actual world, only that he not believe p in the not-p worlds he believes are closest. Henry's WT can be construed as a theory about which not-p worlds are closest (namely, those consistent with his WT); and it, not the world outside his mind, determines which worlds E must rule out. So the presence or absence of holograms has no bearing on whether his claim's truth condition is satisfied. Our conviction that 'Henry does not know he is looking at a barn' is a consequence of the fact that the proposition we assert by saying this sentence is plainly true, since it enlists a WT consistent with holograms and Henry can't distinguish them from barns. In short, the truth of the proposition S asserts by saying 'I know that p' in everyday circumstances depends entirely on factors to which S has cognitive access (namely, S's belief that p, the set of not-p worlds consistent with S's WT, and the fact that E rules out that set)--with the exception of the truth of p.

An attractive response to skepticism, one often given by non-philosophers, is to allow that, while we may not know what the skeptic denies we do, still, we do know 'for all practical purposes.' Our everyday claims are not so ambitious as to be upset by skeptical arguments; understood properly, what we say is compatible with skepticism. As such utterances are typically part of activities that proceed on the basis of implicit practical assumptions, it is sensible to take them to be asserting that, supposing alternatives that conflict with the assumptions are

already ruled out, we know what we say we do. This response is exactly right, I maintain; and the claim that I know that p 'for all practical purposes' amounts to precisely what I have explicated above. There are not two kinds of knowledge, however, 'real' knowledge and knowledge 'for all practical purposes'; nor do the standards for knowledge shift from higher to lower in skeptical versus everyday contexts. Rather, there is knowledge, which is evaluated by the invariant standard of indubitable certainty, and assertions that enlist that ideal as a benchmark to express the strength of our evidence, various possibilities bracketed, concerning propositions we do not maintain we know.

These assertions use 'know' in the very same sense the skeptic does, namely 'indubitable certainty,' which is why the conclusions of skeptical arguments appear to deny our everyday claims. My assertion owes its different truth condition wholly to the fact that it involves an implicit conditional, maintaining that, if I assume certain alternatives are ruled out, then, given my evidence, I cannot rationally doubt that p--which is why skeptical arguments often seem to trade on irrelevancies. In fact, many such statements are true:¹³ right now, assuming the world is as I believe it to be in other respects, my evidence makes it indubitable that I'm looking at a room full of people. They are useful because, against our shared background of beliefs, our theory of the world, the information they express warrants confidence that p. Our epistemic situation, for all practical purposes, is as good as knowledge.

Nonetheless the truth of everyday 'I know that p' claims is compatible with Industrial Strength Skepticism. Suppose no

contingent belief my WT contains is more probable than its negation, except for some Cartesian certainties. If it's no more likely than not that I'm not a BIV, then my perceptual experiences make it no more likely than not that there are three chairs at the table. My everyday claim can still be true, however, for it can still be true that I believe truly that there are three chairs, and, bracketing worlds that conflict with my WT, E rules out every world in which my belief is false. Consequently the fact that I know that p 'for all practical purposes' doesn't entail that p is more probable than its negation.

II

Crucial to the skeptical epistemology I'm proposing is that often the statement we make by uttering a sentence of the form 'S knows that p' doesn't entail that anyone knows that p. This has substantial advantages. Consider the Principle of Closure for Knowledge (PCk):

If S knows that p and S knows that p entails q, then S knows that q.

(For example, if I know I have hands and I know that if I have hands, then I'm not a BIV, then I know I'm not a BIV.) As the skeptic can prove I don't know I'm not a BIV, and as I know that if I have hands, I'm not a BIV, it follows that I don't know I have hands. Naturally this seems intolerable to many philosophers who assume it follows that the claim I make by uttering the sentence 'I know I have hands' is false.

They are left in a dilemma. First, they can deny PCk. But it

is strongly counterintuitive to maintain that I know I have hands but not that I'm not a BIV. Of course, some philosophers embrace the denial of PCK (most notably Nozick, 1981).¹⁴ I'm not trying to present decisive objections against such (or any) epistemologies, but to motivate a real interest in alternatives. PCK appears to be a necessary truth, and it's hard to believe we will ultimately forsake the principle: 'Knowing an argument is valid, and its premisses are true, is sufficient warrant to know the conclusion.' In addition, recall that Robert Nozick identifies knowledge with true belief that 'tracks the truth': I wouldn't have believed p if it had been false, and I would have believed p if it had been true. (Consequently I know I have hands but not that I'm not a BIV, for my belief that I have hands, but not my belief that I'm not a BIV, tracks the truth.) He denies PCK because he holds that 'tracking' is both necessary and sufficient for knowledge; otherwise PCK could still be true.¹⁵ But such a theory invites a Gettier-style counterexample: Sally shows me her 'new Porsche.' I conclude she owns one, and I whimsically infer D: 'Sally owns a Porsche, or Venusians have bribed her to deceive me.'¹⁶ Unbeknownst to me, a Venusian research team has bribed Sally to deceive me, and lent her the Porsche. (Suppose, too, that these Venusians are, in fact, the only intelligent ETs, no one else wants to bribe Sally, she isn't the sort who plays practical jokes without a substantial bribe, she's broke, doesn't steal cars, and nobody is about to give her one.) Plainly I don't know D. I inferred it from the false disjunct, and I have no warrant for the true one; indeed, I attached that because I disbelieve it. Note, though, that I

wouldn't have believed D if it had been false, and I would have believed D if it had been true.¹⁷ This suggests that true belief that tracks the truth is insufficient for knowledge.¹⁸

Second, they can affirm PCk, admit the force of the skeptical argument, yet remain adamant that I know I have hands. But how to escape the conclusion that I do and do not know I'm not a BIV? The most promising strategy is to endorse the 'contextualist' theory that standards for evaluating knowledge claims vary with conversational context.¹⁹ On this account, what we are talking and thinking about when we make a knowledge attribution determines the 'standards for knowledge,' how much the putative knower must be able to rule out to know that p. As I know I have hands in the ordinary conversational context (call it 'C') relative to which a claim attributing that knowledge to me is evaluated, then, given PCk, I must also know in C that I'm not a BIV. As standards shift with conversational context, knowledge is relative: my true belief that I have hands is knowledge by the standards determined by ordinary conversational contexts, but it wouldn't have been knowledge if we'd just been talking about BIVs. As the standards determined by C don't require me to exclude exotic alternatives, I know, relative to those standards, that I'm not a BIV, because it is true of me, relative to those standards, that I know I have hands. Of course, the skeptic can prove in an epistemology seminar that I don't know I'm not a BIV, and therefore that I don't know I have hands; but as soon as he mentions BIVs, we're no longer in C.

Well, why not embrace contextualism? An immediate difficulty is that it seems a necessary truth that, where p is contingent, S

knows p only if S wouldn't believe p if it were false.²⁰ As it is true, even in C, that I would believe I'm not a BIV if I were one, it follows that my belief isn't knowledge in C, context be damned! Given PCK, it follows that I don't know in C that I have hands, either. In short, as knowledge must be 'sensitive' to the falsity of what is believed, contextualism is false.²¹ How might the contextualist respond? Keith DeRose agrees that 'we have a very general inclination to think that we don't know that P when we realize that our belief that P is insensitive...' Nonetheless he proposes the 'Rule of Sensitivity': 'When it is asserted that some subject S knows...some proposition P, the standards for knowledge...tend to be raised, if need be, to such a level as to requires S's belief in that particular P to be sensitive for it to count as knowledge.'²² So knowledge can be insensitive after all. Note that this response is implausible. If knowledge doesn't require sensitivity, why do we do that? DeRose is silent on the matter, and it's hardly easy to think of a plausible answer.

Suppose we accept the proposal anyway. As I know I have hands only if I know I'm not a BIV, it is crucial to the contextualist response to the skeptic that this belief is knowledge, according to ordinary, low standards for knowledge. Now if there is a candidate for a self-evident epistemological truth, I submit it is this: What can have no epistemic bearing on the question of whether or not p, should it arise, is not the knowledge that p. (Let me stipulate that the question of whether or not p 'arises' only when we don't dismiss it or treat it as a joke, but actually consider whether or not p.) The view is common among contextualists, however, that even considering whether or

not I'm a BIV raises the standards to a level at which I don't know I'm not one. David Lewis writes: 'No matter how far-fetched a certain possibility may be, no matter how properly we might have ignored it in some other context, if in this context we are not in fact ignoring it but attending to it, then for us now it is a relevant alternative.'²³ He dubs this the 'Rule of Attention.' Stewart Cohen observes that 'focusing on skeptical alternatives can lead us to consider them relevant.'²⁴ Keith DeRose, articulating a similar view, writes that the 'skeptic's mentioning the BIV hypothesis...makes that hypothesis relevant.'²⁵ He calls this the 'Rule of Relevance.' (DeRose appears to accept this principle, with the reservation that it doesn't satisfactorily explain why we don't know the relevant skeptical hypothesis is false.²⁶) Certainly it's hard to see how anyone who believes that conversational context (which includes internal dialogue) can affect standards for knowledge could plausibly deny the Rule of Attention--the stuff of which contextualism is made, plainly. Consequently my knowing that I'm not a BIV, by ordinary, low standards, can have no epistemic bearing on the question of whether or not I am one, should it arise--for if it arises, I no longer know I'm not a BIV.

Contextualism, therefore, centrally involves the claim that, according to ordinary standards for knowledge, for some p, I know p even though (a) I would believe p just the same if it were false, (b) anybody who says that I don't know p is right (due to the Rule of Sensitivity), and (c) my knowing p can have no bearing on the question of whether or not p is true, should it arise. This verges on incoherence. Epiphenomenal knowledge,

necessarily riding above the epistemic fray, which can have no bearing on the question of whether what is known is true, is 'knowledge' in name alone, whatever the standards. Just as the tiny, trunkless 'elephant' in the corner, who adores cheese but squeaks at cats, is more accurately called a 'mouse,' the insensitive, unclaimable 'knowledge' that I'm not a BIV, which must vanish whenever I wonder if I am one, is really an 'assumption.' As I know I have hands only if I know I'm not a BIV, if the latter 'knowledge' is nominal, so is the former. In short, contextualism, if it isn't simply false, is skepticism wearing a fig leaf.

The skeptical epistemology I'm proposing, on the other hand, affirms PCk (unlike the conditional theory), insists that knowledge must be sensitive (unlike contextualism), yet remains adamant that the claim I make by uttering the sentence 'I know I have hands' can be true. When I now say to you, paraphrasing G.E. Moore, 'How absurd it would be to deny that I know I have hands; it would be like suggesting that I don't know I am standing up and talking!' I'm using the concept of knowledge as an idealization, to point out to you that, supposing alternatives incompatible with my standing here talking to you have been ruled out, I've provided evidence that makes it indubitably certain that I have hands.²⁷ Considering the implications of the claim I make by saying 'I know I have hands' doesn't prevent me from making the same claim truthfully again. As it doesn't entail that I know I have hands, combining what I assert with PCk doesn't entail that I know I'm not a BIV, which is why raising the vat possibility is otiose.²⁸ I know 'for all practical purposes' that

I'm wearing trousers, even in the epistemology seminar.

Of course, it follows that I don't know I have hands--how could I? I would have to know I'm not a BIV, and how could anybody know that? But why on earth would I, in the midst of practical pursuits like setting the table, ever need to make a claim so strong as to entail I know I'm not a BIV? It is misguided to force all this ambitious epistemic content into these innocent pragmatic utterances, then tie ourselves in knots trying to explain how, so interpreted, they could possibly be true. Better to have the humility to accept that we know very little, and the charity to interpret them so that they are useful and true. Then we can move on to other things.

III

But why not say that what I'm calling "knowledge 'for all practical purposes'" is knowledge proper, so our everyday claims are knowledge claims after all? Whatever role it may play in framing such assertions, the philosopher's notion of 'indubitable certainty' has little purchase in the rough and tumble circumstances in which we successfully use the word 'know.' It is far more reasonable, surely, to identify knowledge with whatever satisfies the truth conditions of the utterances in which we typically use that word! But now we must face a question: what do we do about the skeptical arguments? As the skeptic can prove I don't know I'm not a BIV, it follows, given PCK, that I don't know I have hands. So if we insist that everyday claims are knowledge claims after all, they're false. If that is intolerable, we must meet somehow the skeptical challenge. We must deny PCK, or relativize 'knowledge,' or insist there is a

fallacy in the skeptic's argument even though nobody can find it, or.... The alternative to perpetuating this ancient strife, therefore, is to allow that everyday 'I know that p' claims, while they centrally involve the idea of knowledge, aren't claims to knowledge.

But isn't it implausible to maintain that, when people say things like 'I know that p,' they are hardly ever claiming to know that p? Why isn't this just as wrongheaded as insisting that, when they say 'Grass is green,' they aren't asserting that grass is green? Surely the gap between what we say and what our words would most naturally be construed to mean is too large! This objection would have more force if skeptical arguments didn't work. Of course, on grounds of simplicity we should construe everyday 'I know that p' utterances to assert the proposition that the speaker knows that p. On the other hand, we want to interpret common and useful discourse so that it isn't always false. If the simple interpretation engenders awful problems, a non-standard but closely related reading that explains the utterances' truth and utility becomes plausible.²⁹ As the skeptic can prove we don't know we're not BIVs, some apparently obvious principle must be abandoned or we can't preserve the truth of our everyday claims. Perhaps we'll deny PCK, or insist that knowledge need have no epistemic bearing on the question of whether what is known is true. Now it's plausible to suppose that everyday 'I know that p' claims are made against the background of practical assumptions that affect their truth conditions; but then it's hardly implausible to suppose the propositions we assert advert to those assumptions, and involve

the supposition that they are satisfied--though we don't spell all that out in our utterances, naturally. If it doesn't assert the false proposition that I know that p, my everyday claim asserts the closely related true one that, roughly, supposing the world is one in which my present practical pursuits make sense, I know that p. As I observed earlier, this is often the response of non-philosophers confronted with skeptical arguments--no surprise, as it seems pretty intuitive under the circumstances.

But doesn't this miss the thrust of the objection? It's not just that, all things being equal, we should construe an utterance 'X is Y' to assert that X is Y. Things are quite unequal, obviously, if the skeptic can prove we don't know we're not BIVs. The objection's real force flows from the fact that, when we utter the words 'I know that p,' we naturally take ourselves to be making knowledge claims. My supposedly enlightened version of skepticism denies that we typically assert what we universally think we do assert! However this turns out to be a consequence of my theory. For I do know something when I know that p 'for all practical purposes.' Let S be the set of not-p worlds consistent with my WT. My everyday claim entails that it is indubitably certain that I'm not in a member of S.³⁰ That is, I know that I'm not in a not-p world consistent with my WT. Now if we're proceeding on the assumption that our shared WT is true, then, for all practical purposes, this knowledge is the knowledge that p. Hence my everyday claim has exactly the force of that knowledge claim; for all practical purposes, I've made it. Consequently when we ordinarily say things like 'I know that p,' we are claiming to know that p, as far as we are concerned.

Naturally we take ourselves to be making those knowledge claims. Generally, if a state of affairs is a consequence of a well motivated theory T--exactly what one would expect if T were true--then the fact that it obtains isn't a forceful objection to T. Hence the fact that, when we say things like 'I know that p,' we naturally take ourselves to be making knowledge claims, doesn't constitute a forceful objection to my skeptical epistemology, which cogently 'explains the appearances.' So an enlightened skepticism gets us everyday 'I know that p' claims that have the force of knowledge claims, and are true, plus an end to the war with skepticism.

In sum, the skeptical epistemology this paper defends provides a relatively simple and intuitive account of knowledge as indubitable certainty. It explains how that concept is deployed in our ordinary practices, why these are useful, and how what we say when we use the word 'know' can be true. It puts to rest one of the oldest problems in philosophy. In addition, vexing and apparently intractable difficulties in contemporary epistemology don't arise (e.g. Gettier problems and the lottery paradox).³¹ As I argued above, the obvious objections fail straightforwardly. There are no interesting uncontroversial theories of knowledge, plainly. I submit that this theory is sufficiently interesting to be worthy of controversy--especially if we take skeptical arguments seriously. Of course, none of this will comfort the scientist and philosopher; for, if no contingent belief (save some Cartesian certainties) is ever more probable than its negation, the project of science appears to become irrational. I turn now to the question: if Industrial Strength

Skepticism is true, what is lost?

IV

Consider the claim that the universe started five minutes ago, complete with libraries full of books about World War I, automobiles with 100,000 miles on the odometers, people with apparent memories of lives they never lived, and so on. There is a possible world w that satisfies that description; but then, as we have no immediate access to the past, all our evidence that there is a past (our present 'memories,' 'worn out' shoes, grey hairs, and so on) exists in w , too. As our evidence cannot rule out the possibility that this is w , for all we know the universe started five minutes ago. Indeed, as our 'evidence' is the same whether or not there is a past, it makes one claim no more probable than the other.

Now consider the same claim as a scientific hypothesis. Why was the universe as it was five minutes ago? Why these 'fossils' and not others? Why these 'memories' and not different ones? Why us and not other people? Why any people? Such questions would seem to be unanswerable: that's just how things began. But the hypothesis that there is a past which caused our memories, and which we can reconstruct from fossils, ruins, geology, and so on, enables us to explain all of this. Further, it enables us to explain why the universe was as it was six minutes ago, seven minutes ago, and so on. In short, the hypothesis of the past 'projects an explainable universe,' one in which we find more explanation: it maximizes explanation.

Note, too, that the Problem of the Past, the Problem of

Induction, and the Problem of the External World have the same logical structure. In each case there is an assumption that makes evidence possible--an Evidence Creating Assumption ('ECA'). The assumption that I perceive an external world makes my sensory experiences evidence about what is happening in it. The assumption that the future will resemble the past makes them evidence about what will happen in the future. The assumption that there is a past makes memories, photos, fossils, evidence about what that past was like. What skeptical arguments show is that ECAs cannot, without circularity, be justified by the evidence they create: if we aren't justified in accepting the ECAs to begin with, there is no evidence at all. Consequently they cannot be justified by evidence, unless it is evidence that doesn't depend on an ECA; but, if skepticism is true, there is none. Therefore ECAs are no more probable than their negations, and the particular claims warranted by the evidence ECAs create are also no more probable than their negations. Note, however, that, in each case, the ECA helps to project an explainable universe. The assumption that I perceive an external world enables me to explain my sensory experiences, then the surface features of the objects I perceive, then the features that explain the surface features. ECAs are instances of a grander, overarching assumption: that the universe is explainable. What entitles us to make it? After all, it is no more probable than not that the universe is explainable.

'It is because of wondering that men began to philosophize and do so now' Aristotle writes. 'First, they wondered at the difficulties close at hand; then, advancing little by little,

they discussed difficulties also about greater matters, for example, about the changing attributes of the Moon and of the Sun and of the stars, and about the generation of the universe.'³² Consider the fascination of the pre-Socratics with cosmology and cosmogony. Their first concern wasn't prediction or control; they weren't pragmatists. Above all, they were motivated by raw curiosity about the way things really are: they wanted to explain the universe. Now if that is our project, we must suppose our perceptions are more or less accurate representations of the surface features of their objects. Otherwise we can explain nothing. Also, we must suppose a past that our memories recall. If this is my first glimpse of a brand new universe, and all my 'memories' are illusions, I lack the information I need to mount an explanation, not to mention a fix on what needs explaining. The very act of looking for explanations supposes there may be regularities that we can find. Therefore, if we are entitled to try to explain the universe, this entitles us to view it as we must if explanations are to be found--to suppose there is a past, that the future will resemble it, and that we're not BIVs.

Of course, the universe already seems to be one in which these assumptions are true; indeed, if none of them seemed true we could have no projects. As I will argue below, the project of explanation flows from, and is a response to, the way things seem. But skepticism has no quarrel with our being motivated by appearances, as long as we don't suppose we have any reason to think they represent realities. And if a project is caused by, or is a natural response to, the way things seem, so be it, so long as it presupposes nothing that skepticism denies.

Well, then, what entitles us to look for explanations? On the face of things, we are entitled to pursue any worthwhile project to which we can contribute; that is a good and reasonable way to spend our time, surely. Of course, it will be objected immediately that skepticism renders the scientific project worthless. For, if skepticism is true, we can never know we have found an explanation; indeed, it can never be more likely than not that we've found one. Surely it is irrational to look for what we can never know we've found! A project is senseless if we can't tell the difference between success and failure. Given skepticism, therefore, the project of looking for explanations is a waste of time.

But I believe it is reasonable to look for explanations, even though we can never know we've found one. The project of pure science stays meaningful, because a) there is something that counts as success and something that counts as failure, and b) the success state is definitely worth having. Consider three worlds:

1. In the Best Case World, we have a completed science, the universe is as we hope it is, and our science explains it.

2. In the Demon World, we are brains in vats, the universe started five seconds ago, and there are no regularities.

Nonetheless we have the same science we have in the Best Case World and, for the moment, the universe appears just as the Best Case World does. Given skepticism, it is no more probable that we are in the Best Case World than that we are in the Demon World; equal evidence for both hypotheses.

3. In the Mush World we inquire for centuries and find no natural

laws. Predictions always appear to go wrong, hypotheses are always disconfirmed. There are no candidate explanations. Science goes nowhere.³³

We succeed in both the Demon World and the Best Case World. In both we are entitled to say: 'There is a possible world we have explained, and if we are lucky, we are in it. And there is no reason to think we are not in it.' That is, if we are lucky, we have explained our universe. This is a kind of success. In the Mush World we cannot say this, which is a kind of failure. The former condition is success compared to the latter. In the Demon World, we are unlucky. In the Best Case World, we are lucky. We have no evidence which world this is, one way or the other. But at least we haven't come up empty (or even half empty), as we do in the Mush World, without even candidate explanations. In both the Demon World and the Best Case World, we know we are not in the Mush World. So there is something that counts as failure of the scientific project, and something that counts as success. I submit that it is enough to make the scientific project worthwhile. We need no more.

The wonder in which philosophy begins doesn't arise in a swirling chaos. Wonder is the response of curiosity to the appearance of order. I experience myself as an embodied being in a persisting world containing mountains and seas and trees and animals other than myself. Winters follow summers, cats have kittens, and we can navigate reliably by the stars. That is how things seem. What is going on in this apparently orderly universe, what is the underlying logos that explains what we see, and how did things get this way? This isn't epistemic wonder: am

I a BIV? Is there a past? Even if we sometimes wonder about skeptical possibilities, that isn't the wonder in which philosophy begins. Scientific wonder is about the orderly universe in which we believe ourselves embedded, and only peripherally about our ability to know about it. We can say, in both the Best Case World and the Demon World, that if the universe is anything like the one we want to explain, we've explained it. Wonder is satisfied if this is the universe that wonder presupposes. That, too, is a kind of success, one which makes the scientific project worthwhile, for our curiosity is satisfied.

Can wonder itself be called into question? If we have no more reason to think we're in the Best Case World than in the Demon World, then wonder is misguided, surely. Can it be rational to be so deeply curious about the underlying nature of a universe we have no reason to believe even exists? Well, why not? We can't create ourselves from nothing, after all. Sooner or later we must take ourselves as we are, and we find ourselves curious. Also, we don't experience ourselves as BIVs, but as embodied creatures going about an orderly world. As I have no more reason to think I'm in the Demon World than in the Best Case World, it certainly seems reasonable to want to see how far explanation goes on the supposition that I'm in the latter. I may very well be explaining how things are.

To sum up: We are entitled to project an explainable universe because we are entitled to try to explain the universe. We are entitled to try to do that because something counts as failing and something counts as succeeding, namely, being

entitled to say: 'If we're lucky, we've done it.' This project is worthwhile because it satisfies our curiosity within the assumptions that our curiosity presupposes. In addition, if a situation bears two interpretations, neither of which is more warranted than the other, and we can interpret it so that we have what we want, then, *ceteris paribus*, we should take it that way. That is because, *ceteris paribus*, it is good to have what we want. Of course, interpreting things so they constitute what I want may be bad if this will make me credulous, but it is unlikely that taking myself not to be a BIV will lead me inexorably to the New Age. So it is reasonable to project a universe in which explanation is maximized. As we are entitled to proceed on the assumption that the ECAs mentioned above are true, we are entitled to talk about evidence and probability. But at bottom, excepting some Cartesian certainties, no belief is more probable than its negation, and we know very little (though we can say we 'know' various things in the idealized way outlined in section I). That's life.

Skeptics are seldom content to maintain that skepticism is correct, however. In addition, we must be persuaded that it is good for us. The Pyrrhonists taught that it leads to tranquillity; Hume thought it makes for tolerance. If I may put in my two cents, skepticism frees us from grasping after impossible assurances we don't need. We are falling through epistemic space, without the twin parachutes of knowledge and probability to slow our descent, but that's no reason to stop doing science. The answer to skepticism isn't epistemology, but courage.³⁴

References

Apostle, H. G. 1979: Aristotle's Metaphysics, (The Peripatetic Press, Grinnell Iowa).

Brower, B. forthcoming: 'Contextualism' in E. Craig, ed. Routledge Encyclopedia of Philosophy.

Cohen, S. 1988: 'How to be a Fallibilist,' Philosophical Perspectives, 2, Epistemology.

DeRose, K. 1992: 'Contextualism and Knowledge Attributions,' Philosophy and Phenomenological Research, Vol. LII, No. 4, pp. 913-929.

DeRose, K. 1995: 'Solving the Skeptical Problem,' The Philosophical Review, Vol. 104, No. 1, pp. 1-52.

DeRose, K. 1996: 'Relevant Alternatives and the Content of Knowledge Attributions,' Philosophy and Phenomenological Research, Vol. LVI, No. 1, pp. 193-202.

Dretske, F. 1970: 'Epistemic Operators,' Journal of Philosophy, pp. 1007-1023.

Goldman, A. 1976: 'Discrimination and Perceptual Knowledge,'
Journal of Philosophy, pp. 771-791.

Kornblith, H. 1993: Inductive Inference and Its Natural Ground:
An Essay in Naturalistic Epistemology (Bradford: The MIT
Press).

Lewis, D. 1996: 'Elusive Knowledge,' Australasian Journal of
Philosophy, Vol. 74. No. 4, pp. 549-567.

Moore, G. E. 1959: Philosophical Papers (London: Allen and
Unwin).

Nozick, R. 1981: Philosophical Explanations, (Cambridge: Harvard
University Press).

Stone, J. 1984: 'Dreaming & Certainty,' Philosophical Studies, 45
(1984), pp. 353-368.

Footnotes

1. I set aside the dreaming argument, which I believe has been
refuted (see Stone, 1984).

2. The emphasis here is on 'relatively.' Much needs to be said to
make clear how this account (or any other) will work for
necessary truths, for instance. This paper will focus on

contingent truths. Plainly the burden of giving an adequate account of knowledge becomes lighter if we needn't have any knowledge.

3. The condition that p is true is redundant, for the condition that my warrant renders rational doubt impossible is meant to entail that, given my warrant, p must be true. I include it because I assert p in claiming to know it--my audience isn't meant to infer that p merely from the strength of my warrant.

4. DeRose 1995, p. 42.

5. Of course, in both cases it is also required that p is true and I believe it. My evidence E 'rules out' all the members of a set of worlds R just in case there is no member of R in which E exists (here I follow Lewis, 1996). Also, E 'rules out' alternative A just in case there is no A-world in which E exists; so seeing that the car is in the driveway doesn't 'rule out' the alternative that my wife is at work, as I will use these words.

6. WTs aren't 'creatures of probability.' When I judge that p in a particular situation T, my WT provides a menu of mistakes that I must rule out in T in judging that p; other mistakes I believe cannot be made. As I believe that newspapers don't make mistakes in announcing lottery tickets, I believe I cannot be mistaken in judging that I've lost the lottery because the newspaper has announced the wrong winning number (see note 8). But it's certainly possible that I believe a mistake cannot be made even though my other beliefs suggest that it can. In this connection, note that much research on human inferential tendencies is pessimistic. Reviewing it, Hilary Kornblith (who is himself optimistic) writes:

The catalog of inferential error which we are naturally inclined to commit seemed limitless: we routinely violate the probabilistic law of large numbers by confidently making judgments about populations on the basis of extremely small samples...;...we have an unseemly attachment to our beliefs, holding on to them even when our evidence has been completely undermined ...;...our degree of confidence in our own judgments far outstrips our objective reliability.... The list goes on and on. (Kornblith 1993, p. 83)

7. When it comes to determining the contents of an individual's WT, I recommend asking her: 'Do you think this is a mistake you can actually make in this situation?' (e.g. 'Do you think that judging you see a chair because you see a hologram that looks like a chair is a mistake you can actually make in this situation?').

8. This provides a simple solution to a version of the 'lottery paradox.' Stewart Cohen presents it (roughly) in this way:

Suppose S holds a ticket in a fair lottery with n tickets, where the probability $n-1/n$ of S losing is very high. Does S know that his ticket will lose? Although...S has good reasons to believe he will lose, it does not seem right to say that S knows he will lose. ...Now...suppose S reads in the paper that another ticket has won. [W]e are inclined to say that S does know that he loses. (Cohen 1988, pp. 92-93)

The paradox is that we're inclined to attribute knowledge to S in the second case but not in the first, even though, by increasing the number of tickets in the lottery in the first case, we can make it more likely that S loses in the first case than in the second: maybe the paper made a mistake. Here's the solution. Suppose I'm S. I don't know in either case that I will lose, but my everyday claim 'I know I will lose' is false in the first case, true in the second. In the first, I assert that I believe truly that I will lose and, assuming this isn't a 'win-world' that conflicts with my WT, then, given E (the odds against my winning), it is indubitable that I will lose. But that's false: there are win-worlds consistent with my WT that E doesn't rule

out. In the second case, I confess it: I believe newspapers don't make mistakes about winning lottery tickets! So judging I've lost because the newspaper has announced the wrong number isn't a mistake I believe I can actually make. Newspapers make some mistakes, of course, but this one is especially easy to avoid and they are powerfully motivated not to make it. No newspaper has ever made it, to my knowledge, and I believe I would have heard if it had happened. Most people think as I do. Maybe I'm wrong, perhaps my belief that I can't make this mistake is unwarranted by my other beliefs (see note 6), but it's part of my WT. Therefore, setting aside worlds that conflict with my WT, my evidence (seeing '22635' in the paper and looking at '935701' on my ticket) rules out every win-world consistent with my WT.

9. An apparent consequence is that, where p is a necessary truth, I can't know p 'for all practical purposes.' For example, I don't know that $2+2=4$ 'for all practical purposes.' Of course, I may still know p . Perhaps it is possible to revise the account to accommodate alleged 'a posteriori' necessary truths. As mentioned above, this paper's focus is on contingent beliefs, and I leave these issues for another day.

10. Keith DeRose (1996) gives an analogous but substantially different account of how such utterances work.

11. Suppose Henry studies this essay and sees clearly that the proposition his earlier utterance expressed is true. He will still 'speak with the vulgar,' because doing so is more informative and practical than the alternative.

12. This denies Alvin Goldman's position in Goldman, 1976. The Henry example is borrowed from this seminal article (only Goldman's pseudo-barns are papier maché facades).

13. Where p is the object of an everyday 'I know that p ' claim, I don't know that p is true, of course; but if p is true, what I actually assert by uttering the sentence 'I know that p ' will in many cases be the case.

14. Fred Dretske also attacks PCk, but acknowledges that he gives a 'weak form of argument' against it (Dretske 1970, p. 1019).

15. Suppose tracking is necessary but insufficient for knowledge. I don't know I'm not a BIV, plainly. However, my belief that I have hands may also fail to be knowledge, even though it tracks the truth. For there may be additional necessary conditions for knowledge that it doesn't satisfy (e.g. that I have no belief which isn't knowledge that I know it entails). We haven't generated a counterexample to PCk. Suppose tracking is sufficient but unnecessary for knowledge. I know that I have hands, plainly. But my belief that I'm not a BIV may be knowledge, too. For there may be other sufficient conditions for knowledge that it does satisfy (e.g. that I have knowledge that I know entails it). Again, we have no counterexample to PCk.

16. I whimsically disjoin this scenario to any new belief when I can sensibly do so, mutatis mutandis.

17. If D had been false, it would have been because the Venusians didn't bribe Sally to deceive me. As stipulated above, no one else is waiting in the wings to bribe her, and Sally, unbribed, isn't going to come up with a Porsche. So if the Venusians don't bribe her (and lend her a Porsche), she's not going to show me her 'new Porsche,' or even tell me she owns one--not that I'd take her word for it! So I wouldn't have believed D if it had been false. Well, what if D had been true? Note that the second disjunct is true at the closest possible world at which D is true; hence I believe D in that slightly different world (suppose

the Venusians lend Sally a different color Porsche).

18. Note that the plausible thesis that tracking is necessary for knowledge, conjoined with the apparently necessary PCk, yields the skeptical conclusion that I don't know I have hands.

19. This definition paraphrases Bruce Brower's in Brower, forthcoming. Contextualists accept PCk, and so hold that I know I'm not a BIV in the context in which the claim that I know I have hands is true. My own view isn't contextualist, because it maintains the standard for knowledge is invariant, and that everyday 'I know that p' claims aren't claims to know p and so don't engage PCk.

20. If my evidence doesn't rule out the closest possible not-p world, then I can't tell whether or not this is that world; as I can't tell whether or not p is true, for all I know it's false. I believe there is a principled exception to this 'sensitivity requirement' on knowledge, however. If p entails that I'm positioned epistemically so that I can immediately discern certain facts, one of which is that p, and not-p entails that I'm at a disadvantage on account of which I might believe p and not be able to detect my mistake, then I can know p even though I might believe p if it were false (see Stone, 1984). I know it when I'm wide awake (because I'm wide awake) even though I might believe I'm wide awake if I were asleep and dreaming. To my knowledge, however, no non-contextualist who admits exceptions denies the sensitivity requirement itself: exceptions either 'prove the rule,' as above, or can be eliminated by modifying the requirement (see Nozick 1981, p. 179).

21. The 'sensitivity' terminology was introduced by DeRose (see DeRose, 1995).

22. DeRose 1995, p. 36.

23. Lewis 1996, p. 559.

24. Cohen 1988, p. 114.

25. DeRose 1995, p. 14.

26. DeRose 1995, p. 16.

27. G. E. Moore points to his left hand and says 'Here is one hand,' and to his right hand, saying 'and here is another.'
(Moore 1959, p. 144) Moore is certain he knows that what the combination of words and gestures expresses is true. 'How absurd it would be to suggest that I did not know it...and that perhaps it was not the case! You might as well suggest that I do not know that I am now standing up and talking...'(p. 145) He observes that we accept proofs of this sort all the time; for instance, A can prove to B that there are three misprints on a page by pointing to them. What Moore misses is that such demonstrations generate only knowledge 'for all practical purposes,' which is why it would be absurd to doubt that they work. Moore hasn't proved he knows he has hands, only that it would be absurd to deny his everyday claim 'I know I have hands' is true.

28. As we saw in section I, I know 'for all practical purposes' that there are three chairs at the table, but not that I'm not a BIV--that's something I just assume tenaciously. PCK' (the analogue of PCK adapted for knowledge 'for all practical purposes) isn't necessarily true.

29. A non-standard reading of what we assert when we say 'Grass is green' is wrongheaded because it is entirely unmotivated.

30. Suppose I believe p truly and, assuming this isn't a not- p world that conflicts with my WT, then, given E , it is indubitably certain that p . Then it is indubitably certain that I'm not in a not- p world that is consistent with my WT.

31. In paradigmatic instances of Gettier counterexamples, I have a justified false belief from which I infer something accidentally true, thereby generating a justified true belief that isn't knowledge. As the justification required for knowledge ('justification*') rules out every possible mistake, I can't have a justified* false belief, nor can I have a justified* belief that is accidentally true. What happens to knowledge 'for all practical purposes' in Gettier scenarios? Suppose S 's colleague Alice shows S her 'new BMW,' shows him the bill of sale, and so on. In fact it is an elaborate practical joke; Alice doesn't own a car. S believes falsely that Alice owns a BMW, and concludes that someone in his department owns one, which is true because of another colleague's secret BMW. Interestingly, what S asserts by saying 'I know that someone in my department owns a BMW' is true: S believes truly that someone in his department owns a BMW (let this belief= p), and, bracketing not- p worlds inconsistent with S 's WT (which includes the belief that his colleagues don't play weird epistemic pranks), E makes it indubitable that p . As the truth of S 's claim doesn't require that he not believe p at the closest not- p worlds, but only at the not- p worlds consistent with his WT, it is compatible with p being true by accident. Our judgement that ' S does not know that p ' is also true, for it enlists a WT consistent with such pranks and E doesn't rule them out. Again, no Gettier-related problems arise. (See note 8 for a solution to the lottery paradox.)

32. Aristotle, Metaphysics, Bk I, sec 2, 982b 13 in Apostle, 1979.

33. The Mush World isn't a complete chaos, or nothing would count as 'our inquiring for centuries.' It might be argued that this would require some successful 'folk psychological' explanations, e.g. that desires often cause behaviour. And if we are to survive and make a living, there would need to be some rough and ready 'laws' (e.g. that heat plus combustible material makes fire.) Perhaps so. But suppose that our scientific aspirations are otherwise frustrated, so that we remain almost entirely in the dark about why things happen as they do. Indeed, the universe is sufficiently chaotic, transient, and obscure that efforts at astronomy and even botanical classification come up empty. We never find law-like explanations of the sort that we now consider 'scientific,' so that, after hundreds of year of inquiry, science is less advanced than it was before Thales.

34. Thanks to Robert Berman, Keith Butler, Keith DeRose, Ed Johnson, Kathrin Koslicki, William G. Lycan, Alan Sussman, Leopold Stubenberg, and to anonymous referees from Philosophy and Phenomenological Research for comments. Special thanks to Judith Crane.