

Uncertain Knowledge and Reflective Epistemology

Our knowledge forms a highly interconnected and dynamically changing body of propositions. One obviously important way that knowledge changes is via rational inference, based either upon new insight into the content of what we already know or upon new knowledge provided by the senses. The most obvious codification of the acceptability of inference driven knowledge growth is the so-called known entailment closure principle, the principle that if S knows that p and knows that p implies q then S knows that q, or, more formally:

$$(KC) [Kp \ \& \ K(p \rightarrow q)] \rightarrow Kq.$$

(KC) is intended to be a principle about knowledge, not about belief. In fact, the belief version of (KC):

$$(BC) [Bp \ \& \ B(p \rightarrow q)] \rightarrow Bq$$

is not very plausible and in fact seems intuitively far *less* plausible than (KC). How can this be, given that knowing implies believing? This can look paradoxical since although (KC) does not imply (BC), it does imply a version of (BC) restricted to those beliefs which are known, and there seems to be no reason why this should be so if (BC) is not correct in general. I think that the resolution of this minor paradox is to be found in the fact that while (BC) is an empirical claim (KC) is a normative one. Our response to the normative principle is founded on a usually implicit assumption that the empirical doxastic conditions underlying it have been met. But, in particular cases, these conditions may *not* be met.

Thus a worry that (KC) could fail for ‘mechanical reasons’¹. It is possible to doubt (BC) simply because it is possible to imagine irrational believers. That is, it would seem possible to imagine someone who just fails to believe q despite believing both p and $p \rightarrow q$. If p and $p \rightarrow q$ are assumed to also be known, then (KC) fails as well. And of course this sort of thing can and does happen. The plight of those who fail to believe what *obviously* follows from what they already believe is all too well known. One of my favourite illustrations comes from the Darwin Awards web page, where we find this sad tale:

An assistant plant manager for Blacklidge Emulsions died when he used an acetylene torch to cut a hole in a 10,000 gallon tank of asphalt emulsion. He was

¹ My thanks to Josh Mozersky for pressing the importance of this upon me.

attempting to visually survey the amount of emulsion that remained in the tank, but ‘no safety precautions were taken before the cutting operation began,’ stated an OSHA representative. ‘[His] attention was twice called to a warning sign on the side of the structure which stated the contents were combustible. In complete disregard of safety procedures,’ the erstwhile manager ‘lit an acetylene torch and began cutting, causing an explosion that blew him 93 feet away.’

We can explain such lapses in various ways and there is no doubt that we are all irrational believers from time to time. This fact does not impugn (KC) however. We can try to express (KC)’s intended meaning better by restating it with the satisfaction of the empirical doxastic conditions explicitly spelled out as:

$$(KC') [Kp \ \& \ K(p \rightarrow q)] \rightarrow [Bq \rightarrow Kq]$$

This formula attempts to state certain conditions sufficient for transforming belief into knowledge. But even (KC') fails to capture fully the intended point of (KC). It falls victim to cases where a subject knows that p, knows that p implies q, believes q but nevertheless fails to know q since the subject does not believe q *because of* believing both p and p implies q. Suppose that a Martian Q-ray is adventitiously producing the belief that q in a subject who, in the current situation, simply would not think to draw the inference to q from the belief that p and the belief that p implies q. Although it would be very hard to spell out exactly how the input beliefs are to be ‘epistemically appropriately linked’ to the output of the inference², my intuition is that if the inputs are not appropriately ‘used’ to generate the output, then such a subject does not count as knowing q. But the point of (KC) was to express the idea that genuine inferences from what is known generate further knowledge, not to make an empirical claim about what inferences actually will be made by any subject under any particular circumstances. Also subject to this understanding, we could strengthen (KC') as follows: $[Kp \ \& \ K(p \rightarrow q)] \rightarrow [Bq \leftrightarrow Kq]$, which follows simply because knowledge entails belief. This version states that under the conditions specified in the antecedent belief is *exactly enough* to constitute knowledge.

However, interesting though they may be, these caveats and complexities won't substantially affect the argument in what follows, which will attempt to exploit the fact that in

² After all, it would be easy to imagine that it is the presence of the belief that p and the belief that p implies q which triggers the Q-ray in a subject that would not otherwise come to form the belief that q. And yet if the Martians installed this device specifically in order to improve our subject's epistemic performance it might well serve as an epistemologically virtuous replacement ‘inference engine’ in a ‘logically challenged’ subject.

some circumstances, just believing the true will suffice for knowledge. So, bearing them in mind, we can proceed to examine (KC). (KC) would appear to be a very conservative principle. By the second conjunct of the antecedent, knowledge is claimed to expand only via inferential connections which are themselves known and used. If the inferential connection linking p with q is not even believed the subject will simply never notice that there is any logical connection between p and q . If the connection is believed but not known, it cannot license new *knowledge* for if not known the connection is either false or unjustified³, and these faults will infect the output of the inference. Note also that in this case, it is only expansion from knowledge that p , via a cognitive process which uses the $(p \rightarrow q)$ connection, to knowledge of q that is blocked. It is still quite possible that q is known, if there is another – either inferential or sensory – route to it.

So (KC) would at first glance seem to be entirely innocuous. Indeed, if we could not rely upon it a huge number of our knowledge claims would be put in jeopardy, and knowledge growth would be largely paralyzed. We would also lose one of the primary means of *justifying* our knowledge, that of appeal to the fact that some new piece of knowledge can be inferred from what is already known.

Nonetheless, a number of contemporary epistemologists (e.g. Goldman 1976, Nozick 1981, Dretske 1970, Audi 1998) have argued, each in somewhat different ways, that (KC) is false. Generally speaking, the denial of (KC) serves some ulterior motive, as we shall see below.

But I want to begin by briefly considering an argument against (KC) which is more straightforward, put forth by Audi (1998, pp. 169 ff.). Audi considers a variant of the following example. Suppose I've added up some numbers, checked the sum twice, and come thereby to know that the sum is 5,687. Call this proposition P . Now, P implies that my wife's claim that the sum is not 5,687 is wrong. So according to (KC) I also know this: my wife's claim that the sum is not 5,687 is wrong. That seems unexceptionable, even if, as Audi asks us to add to the tale, my wife is the better arithmetician. But Audi claims that P also implies a conditional of the form: if my wife says that the sum is not 5,687 then she is wrong. (KC) would then license knowledge of this conditional. But such a conditional is best regarded as a kind of subjunctive conditional. My wife has made no such claim and such a claim is not part of my current epistemic situation,

³ The term 'unjustified' is used here without theoretical presupposition. Something transforms mere true belief into knowledge, and at least a crucial part of that something is lacking in the envisaged case. I label the lack here by the term 'unjustified'.

whose parameters condition whether or not I do know P. Such a conditional must be evaluated in the way an explicitly counterfactual conditional is evaluated, namely – to speak in the jargon of possible worlds – we must consider the nearest possible world, or the possible world most ‘appropriately’ similar to the actual world save for this difference, in which my wife does claim that the sum is not 5,687. This is a world in which my epistemic situation is different from my actual epistemic situation. It is not obvious that in that world she would be wrong, because it is not obvious that in that world I know P or that P is even true (one might well think that in the nearest world where she makes this claim she is right and I am the one who made the mistake, since she is, by hypothesis, the better arithmetician). Thus it is not the case that P implies this conditional, and this sort of example poses no threat to (KC). Although it is probably a needless complication to mention, it is of course right to say that P implies the *material* conditional: if my wife says that the sum is not 5,687 then she is wrong. This is trivially true merely because of the falsity of the conditional’s antecedent, and granting my knowledge of it leads nowhere.

However, most of the denials of (KC) stem from one basic background motivation: fear of skepticism. And one central method of denying (KC) is based upon analyses of knowledge that essentially deploy counterfactuals. According to such analyses, very roughly speaking, I know p if it is the case that if p were not true then I would not believe it. This is often called ‘tracking’: I track p if such counterfactuals hold. It is also often labelled ‘sensitivity’: I know that p if I am *sensitive* to the truth of p, where, again, sensitivity is defined in terms of the above kind of counterfactuals. Then a nice technical feature of counterfactuals makes it easy for closure to fail. The reasoning is parallel to that which revealed the problem with Audi’s example. The counterfactual $\sim x \gg y$ (read as ‘if x were not the case then y would be the case’) is true if in the *nearest* (most relevantly similar) possible world where x is false, y is true. Suppose that $\sim x \gg y$ is true, that x implies z, and that x is actually true. It clearly does not follow that in the *nearest* world where z is true, y is true, because that world could be a world where both x and y are false. For example, the counterfactual ‘if I had sunk my life savings in Nortel stocks in 2000 and hadn’t sold, I’d now be ruined’ seems like a true counterfactual. But using all my money to buy (and hold) Nortel shares in 2000 logically implies investing some money in the stock market in 2000. Nonetheless, it is evidently not true that if I had invested some money in the stock market in 2000 I’d now be ruined.

This idea applies in epistemology as follows. Suppose I know that I have hands. According to the tracking analysis of knowledge, this means that in the nearest world where I don't have hands I would not believe that I have hands. Perhaps in that world I've had a terrible accident at the sawmill, and of course after the accident I would no longer believe I have hands. But, back in the actual world, I also know that having hands entails that I am not a brain-in-a-vat. Anti-closure epistemologists hold that it does not follow that I know that I am not a brain-in-a-vat, because in the nearest world where I am a brain-in-a-vat I still believe that I am not a brain-in-a-vat (that is, I do *not* track being a brain-in-a-vat)⁴. Notice that in that world I don't have hands, so any entailment from having hands to not being in a vat is completely irrelevant. This result is thought to be a positive one since it is assumed to be the case that the skeptic is right that in some important sense we after all *don't* know that we are not brains-in-vats. The denial of (KC) allows for this without destroying such quotidian knowledge as that we have hands.

Despite its technical possibility, denying (KC) is obviously a desperate measure, but desperate situations sometimes require desperate measures and perhaps the threat of skepticism is a desperate enough situation. But one might be forgiven for thinking that denying (KC) leads to more than merely counterintuitive consequences. Consider Dretske's famous zebra case (Dretske 1970)⁵. Imagine that I'm at the zoo, looking into the zebra paddock at some zebras. I'm no expert on zebras, but the animals before me are black and white striped, horse-like creatures, in a zoo-cage labelled 'zebras'. Dretske suggests that commonsense and intuition agree that I know these animals are zebras. Now, let us assume further that as a matter of fact, I could not tell a zebra from a mule that had been cleverly painted to mimic a zebra. Call such mules 'fake zebras'. Dretske contends that I don't know that the zebras in front of me aren't fake zebras. Lots of knowledge related intuitions frequently exploited in skeptical argumentation can be brought to bear in support of this. For example, such fake zebras are *indistinguishable* from real zebras so far

4 As stated, this counterfactual is rather difficult to evaluate. It's fairly plausible to think that the *nearest* world in which I am a brain-in-a-vat is one in which I've been inserted there for medical or, conceivably, entertainment reasons, and would have been informed of the procedure. We must therefore understand 'being a brain-in-a-vat' as the specific skeptical hypothesis that I am a brain-in-a-vat under conditions of *deception* or some other initiating condition which precludes me believing that I am a brain-in-a-vat. Of course, I've given a question begging definition of these conditions, but presumably there is an acceptable, if less economical, specification available.

5 Dretske was not advancing a tracking account of knowledge but rather a closely related type of epistemological theory, now labelled *contextualism*. Such theories will be discussed below. The particular example discussed here, however, fits perfectly into the tracking theory framework.

as I am concerned, so I cannot ‘rule out’ the fake zebra alternative. In terms of the counterfactual based theory of knowledge, it seems reasonably clear that if these were indeed fake zebras, I’d still believe they were zebras. So I’m not *tracking* zebreness that far into the depths of logical space. My beliefs just are not *sensitive* to this kind of, admittedly outré, possibility.

But now let’s also suppose that I’ve read my tracking theorists. I take it that no viable philosophical theory of knowledge can be such that the theory works only if those to whom it applies are not aware of that theory. Thus I can, and if the theory is right even *should*, correctly reason that while I do know that these animals are zebras, I don’t know that these animals aren’t fake zebras. I also know that fake zebras are non-zebras. I should conclude that I don’t know that these animals aren’t non-zebras. But not being a non-zebra is logically equivalent to being a zebra. So I have to admit that I don’t know that these animals are zebras. Now, this is a flat out contradiction; it is impossible for me both to know that these are zebras and not to know these are zebras. But although this line of reasoning is indeed compelling, and properly so in my view, it is no *formal* threat to one who denies (KC), since this reasoning itself depends upon a logical variant of (KC), namely:

$$(KCN) [\sim Kq \ \& \ K(p \rightarrow q)] \rightarrow \sim Kp^6.$$

Nonetheless, this example of reasoning from one’s own awareness of the consequences of a tracking style theory of knowledge should be very disturbing. For the reasoning involved is personally compelling; it is obviously a correct inference from my presumed situation and the tracking account of knowledge. Thus it would be rational for me to accept it. But it is never rational to accept an outright contradiction. So I must believe that the claim that I don’t know that these aren’t fake zebras is a false claim, or else give up the claim that I know that they are zebras. That is, on either disjunct, I must believe that the tracking theory of knowledge is false. I think that if it is demonstrably rational to believe that theory T is false then that, rather decisively, counts against T.

It is also worth noting the somewhat curious fact that at least an intermediate step in the

6 This is quite obvious for the basic versions of (KC) and (KCN). The case where we think in terms of the slightly more realistic (KC’) is somewhat more complicated. A little elementary logic then shows us that we can derive the somewhat peculiar formula: $[\sim Kq \ \& \ K(p \rightarrow q)] \rightarrow (Bq \rightarrow \sim Kp)$. In terms of our skeptical example, the consequent translates to if S believes that he is not a brain in a vat then he does not know he has hands (just because, of course, he fails to *know* that he is not a brain in a vat as specified in the antecedent).

above reasoning is immune to the criticism levelled by the tracking theorist against (KC). This is the step from something's not being a non-zebra to its being a zebra. This depends upon a logical *equivalence*: x is not a non-zebra iff x is a zebra. If X and Y are logically equivalent then they hold in exactly the same set of possible worlds and the technical feature of counterfactuals adduced above cannot come into play. That is, the principle that $Kp \ \& \ K(p \leftrightarrow q) \rightarrow Kq$ *does* hold even in a tracking theory of knowledge. This is more peculiar than it looks at first. For any transformation of my knowledge that depended upon this principle would seem actually to work via the inferential connection from p to q , the reverse implication sanctioned by the equivalence would seem to be irrelevant to my deducing q from p and $q \leftrightarrow p$. That is, we would normally think that reasoning from an equivalence requires deploying one of the dual implications (that is the way we teach 'natural deduction' after all). If (KC) is denied however, this cannot be the right story of how we expand knowledge via known equivalences, and that seems very strange indeed.

But perhaps I, as a good tracking theory acolyte, should be even more self-conscious, and explicitly reason throughout in terms of the tracking theory of knowledge. Then, if I accept that I know that these are zebras then I can infer that I do track zebranness. Of course, I don't think that I track zebranness throughout all of logical space because I am well aware of the damaging effect on my knowledge of the painted mule hypothesis, or the brain-in-a-vat hypothesis, or various other more or less powerful skeptical hypotheses. On pain of contradiction, I cannot conclude that I both do and do not track zebranness. Instead, I must infer something like the dual claim that I do track zebranness within logical region A but I do not track zebranness beyond region A. However, if my knowledge claim is to be sustained I must then maintain that tracking within region A is sufficient for knowledge. But evidently I also track non-non-zebranness within that region. So I also know that these aren't fake zebras. But I don't know this, by hypothesis. So we get the same contradiction. The contradiction can be avoided if we take knowledge claims to be relative to some *relevant* region of logical space (as is, presumably, the ultimate point of the tracking theory). Relative to region A I do indeed know that these animals are not fake zebras, but relative to a larger region, B, I do not know this. Thus knowledge claims are not 'absolute' and we should rewrite the apparent paradox in something like the following way:

- (1) I know_A that these are zebras and I do not know_B that these are zebras.

This shows that the tracking theory of knowledge essentially connects to – perhaps in a way

reduces to – another theory of knowledge to which I now turn. The so-called *contextualist* theory of knowledge, which agrees with tracking approaches in denying (KC), makes the semantical claim that the word ‘knowledge’ has a meaning which varies with the context of its utterance (see Lewis 1996, Dretske 1970). According to a contextualist account of knowledge, the constraints on the appropriateness of a claim to knowledge vary with the ‘epistemic demands’ created by the context in which that claim is made. In ‘ordinary’ contexts my claim to know that the animals in the cage before us are zebras is completely acceptable, but, for example, your introduction of the possibility that these are fake zebras (painted mules) shifts the context, raising the epistemic stakes so to speak and making my knowledge claim harder – in fact impossible – to sustain. In the new context, my inability to distinguish real from fake zebras undermines my knowledge claim.

Thus the kind of argument given above is evaded since the ‘contradiction’ inferred contains an equivocation: the meaning of ‘know’ is quite different in the claim ‘I know this is a zebra’ from its meaning in the claim ‘I don’t know that this is not a fake zebra’ because of an implicit difference in ‘epistemic context’. We can explicitly bring out these contexts and write the seemingly paradoxical claim in this way, parallel to the formulation above, save for a broadening of the notion of the relevant contexts at issue (which we initially restricted to regions of logical space):

(2) I know[c1] that these are zebras but I don’t know[c2] these are zebras.

This is no more a contradiction than would be a claim that I am 50 years old and I am 49 years old, where the time of the claims is different or a different speaker makes each claim. Perhaps a better example, used by Lewis, would be the apparently contradictory claim: the road is flat and it is not flat. ‘Flat’ has a context relativity that is supposed to be almost exactly similar to that proposed for ‘knows’⁷. We could remove the paradox by rewriting the claim as ‘the road is flat (compared to most roads) and it is not flat (compared to most tables)’.

I think that self-conscious reflection on my own knowledge in the light of a contextualist approach still leads to serious trouble (see DeRose 1995). What is the appropriate context for just thinking about, as opposed to *talking* about, one’s knowledge? And how is that context set in the absence of any conversational interchange? Here I am, looking at the zebras. Do I know they are

⁷ See also Unger (1975) who denies that there is anything more than a pragmatic contextual relativity to ‘know’ and who defends skepticism on the basis, roughly, that the meaning of ‘knowledge’ requires absolute certainty in knowers. I would agree with Unger’s first claim but can see no reason to concede the second.

zebras? Anyone might start to wonder about this, albeit philosophers are more prone to such fits than others. What sets the context for ‘know’ in such thoughts? Suppose I actually start to think about fake zebras, and have to admit to myself that I couldn’t tell the difference between them and real zebras. Does that raise the stringency of my epistemic context, forcing me to rescind my knowledge claim? Does that mean that if I’m so unimaginative that such possibilities never enter my mind that my knowledge claim is in a better position? Ignorance is bliss, it seems. More serious consideration of knowledge would seem to involve thinking about more of the conceivable undermining possibilities, and if this raises the standards of knowledge then any serious consideration of knowledge would destroy it in just the way skeptics traditionally maintain (see Schaffer, forthcoming). It is in any case clear that the ‘context’ of a knowledge claim extends beyond what the thinker or the interlocutors in a conversation can or do think of. For intuitions about the fate of knowledge claims when there is unknown undermining evidence are clear: if this zoo has unfortunately lost almost all its zebras but has covered up this embarrassment with painted mules then intuition seems clear that I don’t know that what’s before me is a zebra, even if, by some lucky chance, it really is one.

Perhaps the contextualist is better off than the tracking theorist. For the latter, serious consideration of knowledge undercuts the theory; for the former, it undercuts knowledge itself. But neither position is ultimately viable, and the contextualist position seems to end up positively embracing skepticism insofar as serious epistemological reflection inevitably raises the contextual stakes facing knowledge claims.

In fact, I doubt that ‘know’ is really a term that changes its meaning from epistemic context to epistemic context, though its ‘use conditions’ obviously do change with context. That is however an almost universal feature of the pragmatics of language⁸. I fear instead that what the contextualist is at bottom endorsing is that what I would call hypothetical undermining evidence is sufficient to undercut claims to knowledge. As already noted, it certainly does seem right to deny that I know that these are zebras if indeed it is true that there are lots of painted mules about even if, by chance (even good chance), the particular animal I’m looking at is in fact a real zebra.

⁸ Thus it is a pragmatic feature of ‘flat’ that one can correctly say that the table is flat, the rule being something like: call a surface flat if it is flat *enough* not to mislead your interlocutors in the present context (flat enough to set a table is one thing, to play billiards another). It seems to me that such a pragmatic rule would be incomprehensible if the *semantics* of ‘flat’ was itself contextually variable. That is, an absolute sense of ‘flat’ is needed to explain the pragmatics of the term.

That would be a typical case of undermining evidence which I do not possess. The fact that I am unaware of this evidence does not insulate my knowledge claim from its corrosive effect. And, closer to the typical skeptical hypothesis, if it really was the case that, say, half the time I was in some kind of suitably realistic virtual reality chamber then my claim to know that these animals are zebras would also be undercut. At least, it is obvious that if I came to believe that half the time I was in such a VR chamber then I would seriously doubt that I was seeing real animals. And since in general undermining evidence which I do not possess can undercut my knowledge claims, a *genuine* VR chamber possibility would seem to also be sufficient to destroy my claim to knowledge whether I am aware of it or not.

But now consider the claim that the mere – entirely hypothetical – possibility that these could be painted mules, or that I could be in a VR chamber, undercuts my knowledge claim. Why should we accept this? The answer, generally speaking, is that it is claimed (and, so far as it goes, the claim is quite correct) that I could not tell the difference between being in the VR chamber and being in the actual world, and if I can't tell the difference then I have no right to claim knowledge which depends upon the presumed truth of just one of these experientially indistinguishable scenarios⁹. A question for contextualists then is this: if raising such possibilities simply in conversation about knowledge raises the contextual stakes enough to undercut a knowledge claim, why doesn't the *fact* that these possibilities obtain as possibilities count as ordinary undermining evidence? But I think the important question is why indistinguishability is relevant, either as undermining evidence or as a contextual stake raiser, if the envisioned possibility is entirely hypothetical?

Suppose we are in court, and Jake has been (correctly as a matter of fact) accused of murder. There are any number of eye witnesses to the crime who watched Jake shoot Nancy, Jake's DNA has been found on both the murder weapon and the victim, as have his fingerprints, there is a crystal clear videotape of Jake committing the crime, and so on. It would be a laughably ludicrous defence to point out that it is hypothetically possible that Jake has a twin (same DNA,

⁹ There is an obvious and interesting similarity between the skeptical tactics here and the claim that underdetermination of theory by evidence undercuts scientific realism (if this is taken to be the view that we have reasonable evidence that our current theories are at least approximately true). However, there is a serious disanalogy: the pessimistic induction by which we note that so far all scientific theories save the current ones have turned out to be quite radically mistaken. This cannot be said of the ordinary claims of interest to philosophical skepticism.

same fingerprints, same scars on his face, etc.). With all due respect to philosophy, I can think of few situations where the contextual stakes on knowledge claims would or should be higher than at a murder trial. Certainly the typical undergraduate philosophy class does not seem to be one¹⁰. On the other hand, it would be an excellent defence (or at least the beginning of an excellent defence) to prove in court that Jake actually does have such a twin. In fact, it would not be a bad defence if there were witnesses of a certain degree of credibility who just *claimed* that Jake had such a twin.

The point is that purely hypothetical possibilities carry no weight in the court. Why should they carry weight in epistemology? The philosophical skeptic might argue that the appeal to a legal setting is irrelevant since a finding of guilt in court is not dependent upon *knowledge* but rather upon the suspect's guilt being 'beyond reasonable doubt'. So it would remain correct, in light of the hypothetical skeptical possibility of a perfect twin, to say that no one really knows if Jake is guilty. This seems to imply that skeptical doubts in this case, and by implication skepticism in general, are *unreasonable* and admitted to be such by the skeptic. That's enough for me and, unless philosophers are in possession of a higher standard of reason than the law courts, should be enough for anyone to reject skepticism. This is not to say that we should identify what is beyond reasonable doubt with a piece of knowledge. For, of course, it sometimes happens that what is beyond reasonable doubt turns out to be false, and hence not known. This is not a problem for the analysis of knowledge, even if it is a serious one for some who now rest in jail, or on death row. More significantly, undermining evidence which no one possesses can always undercut knowledge claims even if the belief in question is true and beyond reasonable doubt. But the undercutting is done by genuine undermining evidence, not merely hypothetical possibilities.

I think one part of the reason merely hypothetical undermining evidence has been taken to be epistemologically relevant is a common and traditionally sanctioned confusion between knowledge and certainty. Suppose that we insist on this principle:

(3) If S knows P then S must be absolutely certain of P.

¹⁰Perhaps there is a strange kind of reverse effect at work here. In a philosophy class (or a philosophy discussion) the actual stakes on knowledge claims are so *low* that there is no need to take seriously any particular claim and for that reason it is easy to overestimate the effect of hypothetical undermining scenarios. What I mean is that the knowledge claims under discussion are, in the worst sense, academic (save, supposedly, for the distinctively philosophical claims about knowledge itself). In the context of the class discussion, nothing hinges on whether I view myself as awake or dreaming. Should the firebell ring, reality will reassert itself.

Here we must first note the obvious ambiguity between an epistemic sense of ‘absolute certainty’ and a psychological sense. The former should be understood in terms of what must be true or cannot be false: P is epistemically absolute certain just in case P is necessarily true in the present epistemic context. Thus it is absolutely certain that I exist right now; even though this is clearly not a necessary truth *tout court*, it cannot be false in my current epistemic context. On the other hand, it is not epistemically absolutely certain that I am a human being, for I could discover that I was an alien orphan (nonetheless, I claim that I *know* that I am a human being). I also have psychological absolute certainty that I am conscious. The link between the two notions is that I am psychologically certain of what I cannot conceive of being false (in my current epistemic situation) and I take what I cannot conceive to be false as epistemically absolute certain. But even if it is true, unlikely as it seems to me, that I cannot be mistaken about psychological certainty, perhaps because of an infallible introspective access to my own beliefs, I *can* be mistaken about what is epistemically certain, for I may well be wrong about what can and cannot be false in my current epistemic context¹¹.

Claims of absolute certainty *can* be undercut by merely hypothetical undermining evidence, since such claims entail that one believe that there is absolutely no way that one could be wrong. But even a purely hypothetical consistent alternative that explains one’s beliefs shows that there is at least one way in which one could be wrong, and hence shows that one should not be absolutely certain of one’s beliefs. Unsurprisingly, almost all of our beliefs fall prey to one or another purely hypothetical undercutting scenario. Hence we should withhold absolute certainty from almost all of our beliefs.

Skeptical arguments typically begin by first presenting an hypothetical undercutting scenario which successfully undermines our absolute certainty in some proposition, P, but then they go on to claim that *therefore* P is not known by either an explicit or an implicit appeal to some principle akin to (3).

¹¹ The sense of ‘epistemic context’ intended here is not that of a linguistic context such as contextualists appeal to. Rather, an epistemic context is a set of beliefs currently held. We assess what is possible within an epistemic context by seeing what is compatible with that set of beliefs. To test whether my existence must be true in the present epistemic context, check whether there is a possible world where I have all the beliefs I currently have but am non-existent. Obviously, there is no such world. But there is possible world in which I believe what I currently believe but am not a human being (the world where I am an alien orphan, for example). Only if I am imaginative enough to come up with the alien orphan hypothesis, or some equivalent, will my psychological certainty in that proposition be shaken. No matter the strength of my imagination, however, it is just a fact that my being a human being is *not* epistemically absolutely certain.

But (3) is far from obviously correct. It is an analytic truth about knowledge that

$$(4) \Box(Kp \rightarrow p)$$

and we can interpret this in terms of certainty (or some subject, S's, certainty, symbolised with \blacksquare) as well:

$$(5) \blacksquare(Kp \rightarrow p).$$

But from these innocuous truths about knowledge it trivially does *not* follow that either

$$(6) Kp \rightarrow \Box p$$

or

$$(7) Kp \rightarrow \blacksquare p.$$

So is it true that whatever is known is certain? I think it is obvious that such a principle is much too strong. Virtually nothing should be believed with absolute certainty, but it does not follow that virtually nothing is known. English even has a way to mark out the distinction between knowledge and certainty with the expression 'known with certainty' or 'certain knowledge' and I think we should respect the implied distinction between knowledge and certain knowledge.

The skeptic (and, I fear, the contextualist) trades upon another pragmatic feature of our use of the word 'knowledge', namely its function as an 'epistemic intensifier'. We often do use claims of knowledge to express our certainty (or professed certainty) in a proposition and thus hopefully to encourage others to accept what we believe. My beliefs are merely a matter of my psychology and do not necessarily correspond to anything in the world. But knowledge has the wonderful property that it is guaranteed to be true (by analytic necessity) independent of any facts about my psychological makeup. Thus, if I can get you to accept that I *know* that P, you cannot but accept that P. Since I can use the word 'know' to indicate my certainty of belief, merely hypothetical undermining scenarios can interfere with epistemic intensification. They return the debate about P to an issue of my psychology, to why I do believe what I do and to what I would believe in the scenario under consideration. And the upshot of bringing forth the hypothetical undermining scenario is to show that I would believe P even if P was false inside that scenario. My attempt at epistemic intensification is thus in some way stymied: I was trying to show that P was certain and here you've found a way to show that P is not absolutely certain. But this does

not show that I don't actually know P. It just shows I shouldn't be absolutely certain of P (we could say that it shows I don't have 'certain knowledge' that P).

A quick look on google entertainingly reveals the epistemic intensification role of knowledge claims. Some random samples: On a UFO advocacy site we find a report of a sighting that includes this remark: 'I don't know why we saw it, but I know something is out there, I don't believe it, I know it.' On a somewhat bizarre religious site: 'We know there is reincarnation. We don't believe it: we know it'. A more common turn of phrase involves contrasting 'think' (the common colloquial alternative to 'believe') with 'know'. Here we find such gems as this remark from a lunatic who maintains that it was the novelist Stephen King who murdered John Lennon: 'I don't think it, I know it'. Or again, from aging hollywood star Rita Moreno: 'I think it's very important to stay in shape. I think it adds to one's longevity. In fact, I don't think it, I know it.' And so on. In every case it would obviously be uncharitable as well as pedantic to pounce upon these speakers of ordinary English with 'so you admit it ... you *don't* believe ...'. The point of these remarks is of course not to deny belief in some proposition, but to intensify one's public epistemic commitment to it¹².

In more normal, non-philosophical, situations of epistemic assessment the raising of undermining possibilities in order to undercut epistemic intensification often really does threaten knowledge claims, because in normal situations the undermining scenario is not merely hypothetical (which is why someone would think it worthwhile to raise it in the first place). In normal situations, the aim is to show that there is some actual undermining evidence; it is not sufficient to show that there is some hypothetical, as opposed to genuinely possible, undermining scenario (for there is *always* such a scenario, at least in cases of empirical knowledge).

Consider presenting the dreaming argument to some beginning philosophy students. It's not all that hard to get such students to think that maybe they don't, after all, know that they are not dreaming. But I think we're tricking them (and, sadly, ourselves too). What we are showing them is that it is not absolutely certain that they are not dreaming, and in this we do, and should,

¹² This intentional mischaracterization of the the relation between belief and knowledge can also be used to report on the strength of the epistemic commitments of others, as in a review of a recent book on witchcraft (Telegraph Newspaper, 16/03/2003) where, in an attempt to mitigate the wickedness of witch persecutors, it is said: 'we need to remember that, for early modern Europeans, the existence of evil magic was a foregone conclusion, as self-evident as the earth's orbit around the sun is for us. They did not 'believe' that witches existed: they knew it, and they acted accordingly'. The scare-quotes of course reveal the author's awareness of the tactic in use.

succeed. This loss of certainty is however not occasioned by the skeptical argument as such, but merely by pointing out, or reminding, the students that sometimes we dream but think we are awake. Thus it is open to an astute student to reply something like this:

... while I now see that it is not absolutely certain that I am awake, you have not given me any evidence, or even any reason to think that there is any evidence, that I am in fact, or even might be, dreaming now. Thus I see no reason whatsoever to withdraw my claim or to entertain any 'extra' doubts that I am awake. I would be an idiot to give up my belief that I am awake merely because there is a purely hypothetical undermining scenario in which I would be having these experiences even though I was just dreaming.

This student would have a point. And if the student *would* be an idiot to give up his belief that he is awake, then it is rational for him to persist in the belief. Since he is in fact awake, it would appear that he still retains his knowledge even if he is perhaps somewhat less certain that he is awake than he was before class began. Furthermore, the philosophical discussion just engaged in does not improve his epistemic position (at least with respect to knowing whether he is dreaming; perhaps it makes him a better thinker as well as providing other epistemic benefits of doing philosophy). He knew he was awake when he came into the class, and he still knows it when he leaves the class. Although there is a hypothetical possible scenario in which his beliefs are false (or true by accident) which is indistinguishable from his actual epistemic situation, this is irrelevant to his knowledge claim. For him, in his actual epistemic situation, believing that he is awake amounts to knowing that he is awake.

One part of the student's reply might be worrisome. If I am not absolutely certain that P then I ought to believe that not-P is possible. So why do I have the student reply that no evidence that he *might* be dreaming has been given. Isn't the admitted fact that the proposed dream scenario is indistinguishable from waking experience a kind of evidence that one might be dreaming? No. Evidence is, minimally, information that does or would raise the probability of some relevant proposition. Thus, for example, the information that I rolled an even number is, in the absence of any other information, some evidence in favour of the proposition that I rolled a six. But the fact that some proposed dream scenario would be indistinguishable from reality has no effect on the probability that one is dreaming.

I have access to certain sensory evidence (roughly speaking, my current experiences)

about the nature of the world around me; label that evidence S. Then the skeptic, advancing the dreaming scenario, notes that

$$\text{Prob}(S | R) = \text{Prob}(S | D),$$

where R is the proposition that one is awake and D is the proposition that one is dreaming (in the right, skeptic specified, exactly waking-like, sort of way). This fact, however, does not raise the probability of D. This is trivial to demonstrate. Let's label the proposition that $\text{Prob}(S | R) = \text{Prob}(S | D)$, Q. Now we can calculate $\text{Prob}(D | Q)$.

$$\text{Prob}(D | Q) = \text{Prob}(D \ \& \ Q) / \text{Prob}(Q).$$

But the dreaming hypothesis advanced by the skeptic is purposefully set up so that $\text{Prob}(Q) = 1$ (that's what it is for waking and dreaming to be set up so as to be completely indistinguishable). So $\text{Prob}(D | Q)$ reduces to

$$\text{Prob}(D \ \& \ Q) = \text{Prob}(D) \times \text{Prob}(Q|D).$$

Of course, whether we are dreaming or not, Q holds, so $\text{Prob}(Q | D) = 1$ as well, and hence

$$\text{Prob}(D | Q) = \text{Prob}(D).$$

Granting Q does not alter the probability of D at all. Recognising the mere possibility of the skeptical hypothesis is not evidence that I am dreaming.

Note that a very similar derivation shows that my current sensory evidence does not enhance the probability of R either. That is, I am not made more sure that I am *not* dreaming (in the skeptic specified sort of way), or that I am in the reality based situation, just by facing S. This would not be the case if $\text{Prob}(S | R)$ and $\text{Prob}(S | D)$ were not equal. Then one of R or D would be favoured by S. I take it that the former case is in fact the situation we are usually in; that is, I take it that our current sensory information fits in with the R hypothesis better than some more general, rather than the specially skeptically tailored, version of D. However, since our probability that R is already virtually 1, any increase occasioned by S will be infinitesimal. Things don't always work out in favour of R however, as the remarkable, but not all that uncommon, phenomenon of lucid dreaming reveals (see LaBerge 1985). In such a dream, one comes to realize that one is dreaming, typically by way of recognizing how seriously bizarre one's experiences are becoming. One faces sensory evidence that would be essentially *impossible*

if one was awake¹³. This shows there is at least sometimes a real issue of whether one is dreaming or not. It is curious and significant that it is the *skeptical* version of this question that ends up being unable to shift our confidence that we are dreaming or not!

The fundamental reason for this is that a proper epistemology should endorse a kind of ‘conservatism of belief’, which can be pungently expressed via the mechanic’s cliché ‘if it ain’t broke, don’t fix it’¹⁴. My current belief that I am not dreaming serves me well. It directly fits in with many other of my beliefs. It does not lead me astray and is amazingly well able to integrate with new evidence that I pick up during the day talking with others, reading the newspaper or browsing the web. This belief has passed all the epistemic tests necessary to count as knowledge (though it may begin to fail these tests at any time). Thus my belief it is enough to generate knowledge. Merely hypothetical possibilities which make no difference to what I should expect to experience are just irrelevant.

How then do we tell whether a scenario is ‘merely hypothetical’ and hence irrelevant to knowledge claims or is ‘genuinely possible’ and thus must be considered in deciding whether or not a knowledge claim is valid? Cases where the scenario is completely untestable and predicts nothing different from predictions based upon what we take reality to be are the obvious cases. In general, I doubt there is any algorithm that will decisively distinguish between genuine and hypothetical possibilities. But roughly speaking the answer is that genuine possibilities are assessed as such relative to the way we think the world is, that is, relative to our current epistemic situation. The idea that we are all brains in vats is not a real possibility relevant to knowledge claims because it involves overthrowing too much of what we think we already know for no epistemic benefit. Similarly for the idea, used above in the courtroom example, that a human might have an indistinguishable doppelganger. These ideas aren’t flat out impossible of course (they are, we might say, at least logically possible and many skeptical hypotheses, such as that of the brain in a vat, seem to be in fact rather more than merely logically possible). But they don’t merit serious consideration, given what we believe (and think we in fact know) about the way the

13 If it was the case that the experiences we were having were totally impossible if we were in the reality based scenario, that is, if $\text{Prob}(S | R) = 0$, then it would follow immediately that $\text{Prob}(R | S)$ would be 0 as well. In any case, $\text{Prob}(S | R)$ has to be extremely low to significantly shift our confidence in R since it starts out with such a high value.

14 Such conservatism can be defended in various ways and serves several purposes in cognitive life. See Harman (1986) or Cherniak (1986).

world works. But of course our ideas about how the world works are subject to change as we accumulate more evidence, and more theories to explain that evidence. Just as, in the small, weird events can quite properly make me seriously consider whether I'm dreaming or not, so too, in the large, weird events can overthrow more or less major parts of our view of the world, and open the door to new genuinely possible undermining hypotheses. Thus if we somehow discovered that it really was the case that fully 2% of the human population did have indistinguishable doppelgangers somewhere on the planet then new defences would become legitimate in a court of law and previously unassailable knowledge claims about who did what would sometimes be undermined. Note that such epistemic shifts are retroactive: yesterday everyone agreed that I knew that Jake did the deed, today it turns out that I never did know. But that hardly *shows* that I don't know anything today, unless and until some real underminers come to light. And it seems I retain my knowledge so long as there are no such underminers.

Another unavoidable complication is that the assessment of possibilities relative to my current epistemic situation is not merely a matter of what I happen to think about the possibility in question and my ability to reason from my current epistemic situation, for I may be irresponsibly remiss in what I manage or even *can* manage to think of. There are standards of rational assessment to which I must live up to; lack of attention, ignorance and stupidity do not help to preserve knowledge.

Finally, one might complain that just rationally fitting in with a general system of belief and evidence cannot be enough for knowledge. Quite so. What is needed in addition is simply the truth of the belief. For example, to return to our case of dreaming, since I am in fact not dreaming, and it is entirely reasonable for me to believe it and I do believe it then I do know that fact.

In order to spell out the relevant conception of belief conservatism a little more fully we need a notion of 'credible threats' to knowledge. Conservative belief systems should respond only to credible threats. To a first approximation, what is threatening about a credible threat, F, is that F is such that if S were to come to believe F, S would withdraw belief in P. Such facts might be 'mere possibilities', but only ones that a reasonable person ought to take into account given their epistemic situation (that is, they must be what we called genuine possibilities). If the credible threat is more than a possibility, but is an actuality then it is what has usually been called

undermining evidence. Thus a credible threat is either a fact which undermines belief or a possibility which a reasonable person ought to take into account and which undermines belief if taken into account. And note there might be facts which ‘open up’ undermining possibilities: for example, a news report which reveals there have been ‘reasonably credible’ allegations that the zoo has replaced most of its zebras with painted mules.

The knower need not have any awareness of these credible threats for his knowledge to be undermined. In fact, it would presumably be impossible for a reasonable person to know of a credible threat and still believe P (or at least to persist in thinking that they *know* P). And while not everybody is rational, people who ought to consider something a credible threat but who neglect it can still have their knowledge undermined.

A credible threat can undermine knowledge claims, but incredible threats cannot. I contend that skeptical arguments depend upon appeals to incredible threats and hence need not be taken seriously in assessing knowledge claims. But nonetheless these incredible threats are not *impossible*: I might be a brain in a vat, but it is rational for me to think otherwise, and there are no credible threats undercutting this belief. So I know it (since, as you will agree, it is *true* that I am not a brain in a vat).

Does this line of thought simply fly in the face of the historical success of skeptical arguments to at least worry people about the extent of their knowledge? Nothing I’ve said implies that people shouldn’t be concerned about the integrity of their belief systems: the unexamined belief system is not worth endorsing. What I say is that a serious examination of my belief system does not lead and should not lead me to, for example, give up my belief that I am not dreaming.

In defence of the persuasiveness of typical skeptical hypotheses, Keith DeRose remarks, about the possibility of being a brain in a vat (BIV):

for however improbable or even bizarre it may seem to suppose that I am a BIV, it also seems that I don’t know that I’m not one. How *could* I know such a thing? (1995, 2).

How could you? Simply by believing rationally that you are not a BIV, it being true that you are not a BIV, and there being no credible threats to that belief. It’s not hard to meet these conditions about being a BIV, unlike an attempt to know, say, that the Maple Leafs will not win the 2004

Stanley Cup. We should not confuse *knowing P* with something like *being able to prove beyond a doubt that P*. DeRose's remarks make sense if one rewrites them to read:

for however improbable or even bizarre it may seem to suppose that I am a BIV, it also seems that I could not prove beyond all doubt that I'm not one. How *could* I prove beyond a doubt that I am not a BIV?

I couldn't. But why think that this undermines my knowledge unless one is assuming already that knowing P implies being able to prove that P beyond any doubt, or being absolutely certain that P?

David Lewis says something quite similar to DeRose's comment:

If you are a contented fallibilist, I implore you to be honest, be naive, hear it afresh. 'He knows, yet he has not eliminated all possibilities of error.' Even if you've numbed your ears, doesn't this overt, explicit fallibilism *still* sound wrong? (1996, 420).

Lewis sounds reasonable because of an equivocation on 'possibilities'. Does he mean just any possibility, or does he mean what we called genuine possibilities? Wilful ignoring of the latter threatens knowledge, but ignoring the former is a proper, and indeed standard, hallmark of rationality. Nobody, anywhere has ever eliminated *all* possibilities of error (which include one's memory being unreliable for more than 500 milliseconds).

Thus we might briskly analyse knowledge something along these lines:

S knows that P iff

1. P is true
2. S believes P
3. It is reasonable for S to believe P in S's epistemic situation
4. There are no credible threats undermining S's belief

It is easy to see how the case of dreaming fits into this analysis. It is true that I'm not dreaming right now, I believe that I'm not dreaming, this belief is reasonable, and there is no credible threat undermining this belief. Therefore I know that I'm not dreaming. The same applies to the brain-in-a-vat hypothesis. One can also see that classic Gettier type cases of non-knowledge fail clause 4. The pieces of false information which figure in these cases are credible threats: if I came to

believe that some of the steps in my reasoning were based on such false information, I would come to doubt my conclusion (or *should* come to have such doubts).

One traditional kind of problem might seem to present a special difficulty for this view of knowledge, of which the lottery paradox is the prime example. Assuming that I have a losing lottery ticket, it seems that the above analysis might well license the claim that I *know* my ticket will lose. Suppose, for a moment, that it did. Obviously it would not follow that I know that every ticket will lose, since that is simply a false belief. The paradox is thus not about knowledge but rather about reasonable belief. I would venture, in agreement with many others, that it is *not* reasonable to conjoin arbitrary numbers of uncertain beliefs. In support of the unreasonableness of this I adduce the laws of probability, which apply straightforwardly to such situations and plainly show that arbitrary conjoining of the uncertain lowers the probability of the conjunction.

But in fact it isn't obvious that the analysis above shows that I (along with all the other ticket holders who believe their ticket will lose save perhaps one) do know my ticket will lose. Although the issues are intricate here, I think we should focus on an aspect of the phenomenology of lottery playing. Suppose that, on the night of the lottery draw, I am watching the numbers being selected on television. There is a palpable change in my epistemic position when I see that I have indeed lost even though I fully expected it. A possibility has been closed off. Why should this be? After all, there remains some chance that I still won but that there has been some kind of error in the report (perhaps they are by mistake playing a tape of last week's draw). Whatever that chance might be, there is a lottery with the same chance of winning, and I would not throw away a ticket on that lottery simply on the grounds that because of the lottery's size I *know* that I'll lose. I would venture to say that the difference in attitude depends upon the fact that the epistemic situation I take myself to be in assigns a definite and at least approximately calculable probability for winning the lottery within an explicit 'chance setup'. Under such circumstances, certain events with very small probabilities are real possibilities which undermine any claim to know that they will not occur. Similarly, I would deny that I know that I will arrive safely at the destination of an aircraft flight because we have some way to assigning more or less definite probabilities to the possible disasters that can occur. It seems that by purchasing the lottery ticket or by getting on a plane one changes one's epistemic situation in a way that opens up a new genuine possibility.

Thus we should distinguish non-genuine possibilities from possibilities with arbitrarily

low probability – some of the latter are genuine possibilities. It is tempting to regard those events which we are not absolutely certain will not (or cannot) occur but which are not genuine possibilities as events with an undefinable or imponderable probability. They are conceivable but there is no way for me to seriously deploy them in my current reasoning. They do not connect ‘properly’ to my epistemic situation or they are, in a way, incommensurable with my current epistemic situation as I take it. Applying this view to the fake zebra problem, it predicts that there should be a significant epistemic difference between idly thinking that it is hypothetically possible that the animal in front of is a painted mule and being given (credible) information that someone has replaced a zebra somewhere on the Earth with a painted mule. Assuming everything else is equal, this information (whether or not I am aware of it) might serve to make me doubt my claim to know this is a zebra. And, whether or not I am aware of this information, it has the power to undermine my knowledge. The case is difficult since I likely have other information that counts against *this* animal being the painted mule – perhaps I rightly think that the zebras are inspected by knowledgeable keepers every morning for example. (And, from the external point of view, if it is the case that the Toronto zoo has measures in place which preclude their zebras being replaced with painted mules, then the information that there is somewhere a painted mule standing in for a zebra will not impugn my knowledge.)

The ‘dynamics’ of belief must take into account that fact that one’s epistemic situation is sensitive to the presumed situations of others in a variety of more or less subtle ways. Suppose that, as I watch the zebras, somebody offers me a very large bet that the zebra standing by the gate is actually a painted mule. Just being offered this bet – which, one might naively think, I ought to regard as easy money – would itself alter my epistemic situation¹⁵. Anyone who seriously would make such a bet does so, one would suspect, on the basis of some ‘insider’ knowledge of undermining evidence.

When the day comes that it is reasonable to believe that people’s brains can be kept alive in vats and fed nerve signals which generate experiences indistinguishable from reality, then the brain-in-a-vat possibility will be a real possibility capable of undermining knowledge claims. We

¹⁵ Which reminds me of an old joke. A father is advising his son about the dangers of the big city. He says: ‘Son, if a man comes up to you and bets you 100 dollars that the Jack of Hearts will jump out of his deck of cards and spit cider in your eye ... don’t take the bet’. This suggests quite a bit about the way our epistemic situation changes upon receipt of new information.

will then be in an epistemic situation radically different from the one we are in now¹⁶. (Of course, if it actually is the case that I am in the world envisioned, but the victim of some awful deception, then I don't know now that I am not a brain-in-a-vat even if I'm not – that's just the result of ordinary undermining evidence which I do not possess.)

The view of knowledge defended here has the consequence that for many propositions, my believing them is exactly enough to turn them into knowledge, but this is not atypical of analyses of knowledge. After all, if there are even roughly specifiable conditions on knowledge and one of them is (as it must be) that the knower believe the proposition in question, then whenever the other conditions are met, just adding belief will create knowledge. But this account is also 'fallibilist' in the sense that my knowing that P does not require that P be absolutely certain, either psychologically or epistemically. Thus, although it sounds faintly paradoxical to put it like this, we can say of many propositions that I know them ... if they are true. I know that I am not a brain in a vat, if I'm not. This raises two final points: the status of meta-knowledge and the philosophical point of epistemology.

Being a reflective epistemologist, and, to return to the zebras, being willing to grant that it is not absolutely certain that these animals are zebras, I can truthfully say: if these are zebras then I know they are zebras. Since I do believe they are zebras, I reasonably infer that I know they are. Do I know that I know these are zebras? To answer that, we should just reapply the analysis. Suppose that I do know they are zebras. I believe that I know they are. It is reasonable for me to believe this (since I think it is reasonable to believe they are zebras, reasonable to believe that there is no undermining evidence, and I think my analysis is reasonable too). So if there is no credible threat undermining my belief that I know these are zebras then I know that I know these are zebras. Of course, if it was unreasonable for me to believe that there are no credible threats to my claim that I know, and I came to appreciate that, then I should withdraw my claim to know that I know. The interesting question is whether there are grounds for thinking there might be a credible threat against the claim 'I know these are zebras' which is not a credible threat against the claim 'these are zebras'. I doubt there are, and if not, then meta-knowledge claims are as easy to sustain as first-order knowledge claims and come for free once the first-order claim is secure. One might reply that if I know these are zebras then I can't be wrong, so to know that I know I'd

¹⁶An amazing fictional treatment of the epistemological worries arising in such a world can be found in Philip K. Dick's novel *Ubik*.

have to know that I can't be wrong, but I don't know that. But why not? I know that I can't be wrong *and* know these are zebras, just because knowledge implies truth. I am happy to grant that it is not absolutely certain that these are zebras, but that's another issue. The account offered here asserts that knowledge is compatible with the non-certainty of the known proposition. And in this case it is true that I'm not absolutely certain that I *know* these are zebras, even as I am not absolutely certain that these are zebras. That need not impugn either first or higher order knowledge claims.

Thus it is evident on this account that I cannot infallibly tell whether or not I know something, but this too should not be surprising. Unless we make the unreasonable demand that knowledge entail absolute epistemic certainty, most of our knowledge will be less than absolutely certain. This can corrupt my knowledge claims only if the hypothetical scenarios that undercut certainty are what I called genuine possibilities (or are actualities). If such scenarios are ungenthine but actual, I cannot be faulted for failing to take them into account, but my knowledge claims will be false. Still, there comes a time when I've done enough and it's up to the world to do its part – to line up with my reasonable beliefs. If it does line up, then I know. If it doesn't, then I don't know. Too bad, but there was nothing I could or should have done about it.

This does raise the question of what is so important about knowing what knowing is. Philosophers may take some intrinsic interest, and a characteristic delight, simply in the complexity of the concept of knowledge with all its subtle sensitivity to a host of cognitive and non-cognitive factors. That is a kind of parochial, purely internal, reason to study epistemology. But it's harder to see that the analysis of knowledge as such matters in a wider sense (any more than the study of the perhaps equally complex concept of, say, 'vehicle'). I would venture to guess that the original motivation behind this part of epistemology, going back at least to Plato, was quite misguided. The idea stemmed from noticing two simple features of knowledge: one, it is a mental state (or it seems to be a mental state) and, two, it guarantees the truth of what is known. If one supposes that mental states can be known by introspection, then we can tell what we know by self-examination. But once we know what we know, we can be absolutely certain of the *truth* of the known proposition without leaving the armchair. But this is a pipe dream.

If this account is on the right track, then the study of knowledge as such seems of relatively little moment. What matters are issues such as: what is it reasonable to believe in, and

about, a given epistemic situation or, in general, what is rationality. A very great deal of modern analytic epistemology can be read as devoted to this huge problem, and much has been accomplished. But even in the absence of a complete theory of rationality, it does not seem very risky to conclude that wholesale philosophical skepticism is an irrational doctrine which poses no threat to our knowledge and its growth via principles such as (KC)¹⁷.

William Seager
University of Toronto at Scarborough

References

- Cherniak, C. (1986). *Minimal Rationality*, Cambridge, MA: MIT Press (Bradford Books).
- DeRose, K. (1995). 'Solving the Skeptical Problem', in *Philosophical Review* v. 104, pp. 1-52.
- Dretske, F. (1970). 'Epistemic Operators', in *The Journal of Philosophy* v. 67, pp. 1007-23.
- Goldman, A. (1976). 'Discrimination and Perceptual Knowledge', in *The Journal of Philosophy* v. 73, pp. 771-91.
- Harman, G. (1986). *Change in View: Principles of Reasoning*, Cambridge: MA: MIT Press (Bradford Books).
- LaBerge, S. (1985). *Lucid Dreaming*, New York: St. Martin's Press.
- Lewis, D. (1996). 'Elusive Knowledge', in *Australasian Journal of Philosophy* v. 74, pp. 549-67; (reprinted in D. Lewis *Papers in Metaphysics and Epistemology*, Cambridge: Cambridge University Press, 1999, pp. 418-45; page reference above are to the reprint).
- Nozick, R. (1981). *Philosophical Explanations*, Cambridge, MA: Harvard University Press.
- Schaffer, J. (forthcoming). 'Skepticism, Contextualism and Semantic Self-Knowledge', in *Philosophy and Phenomenological Research*.
- Unger, P. (1975). *Ignorance: A Case for Scepticism*, Oxford: Clarendon Press.

¹⁷I would like to thank Tim Schroeder for sparking my interest in (KC) and its defence.