To appear in F. Jackson and M. Smith, *The Oxford Handbook of Analytic Philosophy*, Oxford: Oxford University Press

Knowledge and Scepticism

Timothy Williamson

Scepticism is a disease in which healthy mental processes run pathologically unchecked. Our cognitive immunity system, designed to protect our conception of the world from harmful errors, turns destructively on that conception itself. Since we have false beliefs, we benefit from the ability to detect our mistakes; removing our errors tends to do us good. Our cognitive immunity system should be able to destroy bad old beliefs, not just prevent the influx of new ones. But that ability sometimes becomes indiscriminate, and destroys good beliefs too. That can happen in several ways. I start by considering two of them.

(I) Sometimes we can detect an error in ourselves by suspending belief in the proposition at issue, while we assess its truth or falsity on the basis of our remaining beliefs and any new ones that we may form in the process of investigation. For example, someone may test his belief that Toronto is the capital of Canada by consulting an encyclopaedia, as a result of which he forms the belief that it says that Ottawa is the capital of Canada, which he combines with his independent belief that

whatever the encyclopaedia says about such matters is true to conclude that Ottawa rather than Toronto is the capital of Canada. If he cannot find the encyclopaedia, he may find himself unable to decide on the basis of what is left whether Toronto is the capital, but still abandon his old belief on the grounds that it was unwarranted. The original belief survives the test only if it can be recovered, starting from what is left when it is suspended.

Since belief is not a purely voluntary matter, we cannot just abandon a belief at will, but we can still avoid reliance on it in a particular inquiry. The phrase 'suspend belief' will be applied to that case too. In this sense, suspending belief in *p* is consistent with believing *p*. The suspension test often takes a dialectical form. One tries to convince an opponent who does not already accept the belief; one needs a starting-point that the opponent will accept (in order not to "beg the question"). Even when we have no opponent, we can imagine one.

The test works properly only if we also suspend belief in various propositions closely related to the original one: for example, that Toronto is the large capital (which entails the original proposition that it is the capital), that one knows that Toronto is the capital (which also entails the original proposition), that either it is snowing or Toronto is the capital (which is entailed by the original proposition), and that the work of reference will say that Toronto is the capital (which is logically independent of the original proposition). One's opponent could not be expected to accept such propositions; one's belief in them depends on one's belief in the original proposition. If we were allowed to use them in answering the question, the test would be too easy to be useful.

As we become more reflective, we apply the test to more abstract and general beliefs, to beliefs ever more central to our whole system of beliefs. We test our old belief that we have souls, or that what looks bent is bent, or that no proposition entails its own negation. To apply the test

properly to such beliefs, we must suspend increasingly large numbers of related beliefs, thereby leaving a decreasing proportion of our original belief system unsuspended. Moreover, once we start testing our beliefs in the reliability of our usual methods of belief acquisition, we must also suspend our use of those methods, for otherwise our beliefs in their reliability would pass the test too easily. To test your belief that someone is trustworthy, it is not enough to ask him whether he is trustworthy. How much is left when we apply the test to our belief that we perceive an external world?

Sooner or later, we arrive at beliefs so central that their suspension would leave us too cognitively impoverished to accomplish anything useful. For example, if we test the belief that there are good reasons, we must suspend any belief that something is a good reason, so we cannot acknowledge anything as a good reason to assert or to deny that there are good reasons, or even to leave the question unanswered. Beyond some point, applying the test does more harm than good to our belief system. But our disposition to apply the test to any proposition in which we are seriously interested has no internal inhibiting mechanism. The test evolved under conditions of less than maximal reflectiveness. In most practical contexts, the propositions in which we are seriously interested are peripheral enough to our belief system to make the test harmless or beneficial. In philosophy, by contrast, we are interested in propositions so central to our belief system that the test has devastating results when properly applied: it leaves us with virtually nothing. We realize that, yet still feel that it would be blinkered and dogmatic to refuse to test our most cherished beliefs in the way in which we are accustomed to test less cherished beliefs, especially when a real or imagined opponent challenges us to do so.

One form of scepticism exploits our vulnerability to the suspension test. The sceptic drives

us to apply the test to ever more central beliefs, with predictable consequences. Insofar as we are willing and able to abide by the results of the test, we are forced step by step to abandon our whole system of beliefs. In effect, the sceptic's challenges force us down a regress of justifications.

(II) Sometimes, without suspending my belief in any proposition, I reach a conclusion about the objective correlation, or lack of it, between the conditions under which I believe some proposition and those under which it is true. A sufficiently poor correlation indicates that the proposition is not known to be true, for some kind of reliability is a necessary (if insufficient) condition for knowledge. Suppose, for example, that my review of my past performance indicates this: my believing one day that it will rain the next day provides no evidence at all for rain the next day. My weather predictions do no better than chance. Then my belief that it will rain tomorrow hardly constitutes knowledge that it will rain tomorrow.

We can be slightly more precise by speaking in probabilistic terms. Let us say that a proposition q is *evidence for* a proposition p if and only if the conditional probability of p on q is higher than the unconditional probability of p ($\text{Prob}(p \mid q) > \text{Prob}(p)$): that is, q raises the probability of p.[1] Thus if the conditional probability that it will rain given that I believe that it will rain is no higher than the unconditional probability that it will rain, then that I believe that it will rain is not evidence that it will rain, and the status of the belief as knowledge is endangered. More generally, where Bp is the proposition that one believes p, if $\text{Prob}(p \mid Bp)$ $\text{Prob}(p)$, then Bp is not evidence for p, and the epistemic status of belief in p is under threat.

A special case of the apparent undermining occurs when p is certain conditional on the assumption that it is *not* believed, for it follows that Bp is not evidence for p. In symbols, if $\text{Prob}(p \mid$

~*Bp*) = 1 then Prob(*p* | *Bp*) Prob(*p*).[2] Roughly: since all ~*p* possibilities are *Bp* possibilities, adding the information *Bp* cannot reduce the proportion of ~*p* possibilities amongst unexcluded possibilities. Although this case might seem too special to matter, it is crucial to many sceptical arguments.

The sceptic makes me imagine a scenario in which, although my experience and beliefs are relevantly similar to my actual experience and beliefs, I am dreaming, or being manipulated by an evil demon, or floating as a brain in a vat, so that very many of my beliefs are false.[3] Since I actually believe that I am not in the sceptical scenario, in the sceptical scenario I also believe that I am not in the sceptical scenario; in no possible situation am I in it without believing that I am not in it.[4] Thus I am certainly not in it, conditional on my not believing that I am not in it. Symbolically, if *s* is the proposition that I am in the sceptical scenario, Prob(~*s* | ~*B*~*s*) = 1. Consequently, Prob(~*s* | *B*~*s*) Prob(~*s*). In other words, that I believe that I am not in the sceptical scenario is not evidence that I am not in the sceptical scenario, for it does not raise the probability that I am not in it. This casts doubt on my supposed knowledge that I am not in the sceptical scenario.

Of course, sceptics rarely formulate their arguments in terms of the probability calculus. Nevertheless, such calculations provide a convenient way of capturing the widespread intuition that one's belief that one is not in the sceptical scenario is unreliable.[5]

So far nothing has been said about the kind of probabilities at issue. They are not pure subjective probabilities (degrees of belief), for I may be certain that I believe *p*: but if Prob(*Bp*) = 1 then automatically Prob(*p* | *Bp*) = Prob(*p*), so *Bp* would not be evidence for *p* whatever the epistemic status of belief in *p*.[6] Rather, we need probabilities that give some weight to possibilities in which *p* is not believed, in order to isolate the effect of the information that it is believed. They are

possible in the sense that they could have obtained, even if we are sure or know that they do not obtain. We need not treat the probabilities as purely *a priori*, for we can investigate empirical correlations between our belief states and states of the world *a posteriori*, perhaps in the spirit of Quine's naturalized epistemology.[7] The results of such investigations might nevertheless shake our claims to knowledge from within, as with my beliefs about tomorrow's weather. The argument can use such empirical probabilities; it is independent of their exact values.

Given the kind of probability at issue, we may assume that the proposition that I believe that I am not in the sceptical scenario has probability less than one (Prob($B{\sim}s$) < 1), because I could have been convinced by the sceptic to abandon my belief that I am not in the sceptical scenario. We may further assume that, conditional on my believing that I am not in the sceptical scenario, it is not certain that I am not in it (Prob($\sim s \mid B{\sim}s$) < 1), because I could have been in it and would then have believed that I was not in it. Given these two plausible extra assumptions, the information that I believe that I am not in the sceptical scenario actually *lowers* the probability that I am not in it, and therefore raises the probability that I am in it (Prob($\sim s \mid B{\sim}s$) < Prob($\sim s$), so Prob(s) < Prob($s \mid B{\sim}s$)). That I believe that I am not in the sceptical scenario is evidence that I *am* in it.

Call belief in p *truth-indicative* if one's belief in p is evidence for its truth (Prob($p \mid Bp$) > Prob(p)), and *falsity-indicative* if one's belief in p is evidence for its falsity (Prob($p \mid Bp$) < Prob(p)). Belief that one is not in a sceptical scenario is not just not truth-indicative; it is falsity-indicative. Such a belief looks to be in bad shape. Falsity-indicativeness is a better criticism than lack of truth-indicativeness. After all, if p is a tautology, the probability calculus guarantees that Prob(p) = Prob($p \mid Bp$) = 1, so belief in p is neither truth-indicative nor falsity-indicative; that does not show that one cannot know p. The absence of falsity-indicativeness looks more plausible as a

necessary condition for knowledge than truth-indicativeness.

Could one respond to the problem by sacrificing the *recherché* philosophical belief that one is not in the sceptical scenario, whilst retaining the familiar beliefs with which one negotiates the everyday business of life? Very many of my everyday beliefs entail that I am not in the sceptical scenario. I am watching television only if I am not a brain in a vat. We tend to regard deduction as a paradigmatic way of extending our knowledge. If I know that I am watching television, I can come to know that I am watching something by competently deducing the conclusion that I am watching something from the premise that I am watching television. Similarly, if I know that I am watching television, why cannot I come to know that I am not in the sceptical scenario by competently deducing that I am not in it from the premise that I am watching television? I have no difficulty in performing the deduction itself. Contrapositively, if I cannot know that I am not in the sceptical scenario, how can I know that I am watching television? We thus appear to lose most of the knowledge that we ordinarily take ourselves to have.

The problem for an everyday proposition p is not that Bp lowers the probability of p. That I believe that I am watching television surely raises the probability that I am watching television. The problem is less direct. We have a compelling principle something like this:

CLOSURE(K) If one knows $p_1, ..., p_n$ and believes q in the light of competent deduction of q from $p_1, ..., p_n$ then one knows q.

The phrase 'in the light of competent deduction' here is intended to apply both when one's belief in q originates in the deduction and when it originated otherwise but is now sustained by that deduction

and perhaps by many others. Now if *p* entails *q*, belief in *p* may be truth-indicative even if belief in *q* is falsity-indicative. But if a necessary condition for knowing *p* is that belief in *p* not be falsity-indicative, then the effect of CLOSURE(K) is to strengthen that necessary condition: one knows *p* only if belief in no proposition *q* that *p* entails is falsity-indicative (given that there is no problem about believing *q* in the light of competent deduction from *p*). It is this stronger putatively necessary condition for knowing that everyday propositions flout, because they entail that one is not in a sceptical scenario.[8] A similar argument can be given with the proposition that one is not in a specified sceptical scenario replaced by the bare proposition that one does not falsely believe *p* (~(~*p* & *Bp*)), which is also a logical consequence of *p*.

If the problem only concerned knowledge, some philosophers might respond by abandoning claims to knowledge and retreating to claims of justified true belief instead.[9] Knowing *q* is not necessary for having a justified true belief in *q*, for Gettier (1963) showed that having a justified true belief in *q* is not sufficient for knowing *q*. Since false premises entail some true conclusions, one can believe a true proposition *q* on the basis of competent deduction of *q* from a false proposition *p*; if one's belief in *p* is justified, so too is one's belief in *q*; thus one has a justified true belief in *q*, but does not know *q*, because one's belief in *q* depends on a false premise. The idea of conceding scepticism about knowledge while resisting it about justified true belief assumes that one can have a justified belief in a proposition *q* without knowing anything else to justify *q*. For the sake of argument, let us grant that dubious assumption.[10] The argument in the Gettier case assumed a closure principle for justified belief rather than knowledge, like this:

CLOSURE(JB) If one has justified belief in p_1, ..., p_n and believes *q* in the light of competent

deduction of q from p_1, ..., p_n then one has justified belief in q.[11]

Since deduction preserves truth, CLOSURE(JB) entails that if one has justified true belief in p_1, ..., p_n and believes q in the light of competent deduction of q from p_1, ..., p_n then one has justified true belief in q. But the sceptic might argue that if one has a justified belief that belief in p is falsity-indicative, one does not also have a justified belief in p. In particular, if I have a justified belief that belief that I am not in the sceptical scenario is falsity-indicative, then I do not have a justified belief that I am not in the sceptical scenario. But since my belief that belief that I am not in the sceptical scenario is falsity-indicative is justified (or no less justified than my belief that I am not in the sceptical scenario), my belief that I am not in the sceptical scenario is not justified. Given CLOSURE(JB), this conclusion is the thin end of the wedge for a general scepticism about justified belief, just as, given CLOSURE(K), the conclusion that I do not know that I am not in the sceptical scenario is the thin end of the wedge for a general scepticism about knowledge. For I can believe that I am not in the sceptical scenario in the light of competent deduction from an everyday belief; by CLOSURE(JB), if the everyday belief were justified, so would be my belief that I am not in the sceptical scenario. By hypothesis, the latter belief is not justified. Therefore, my everyday belief is also not justified. Thus, even if the justification of belief does not depend on knowledge, sacrificing knowledge of everyday truths is not enough to appease the sceptic's style of argument, for it devours justified belief in those truths as well.

(I) and (II) represent quite distinct routes into scepticism. Unlike (I), (II) does not depend on the suspension of belief. Even if one is certain that one is not in the sceptical scenario, and reasons

accordingly, one must still concede that the belief is falsity-indicative. But the differences between (I) and (II) do not preclude mixed strategies for testing one's beliefs by a combination of the two methods. For example, one might use testing by strategy (I) to undermine the belief that one is not in the sceptical scenario of strategy (II). In practice, mixing often takes the form of opportunistic switching between the two strategies: the sceptic employs whichever is not under current examination. Still, less disreputable mixed strategies also exist. But many sceptics prefer one strategy to the other. Sceptics motivated by internalist concerns about justification may find strategy (I) more compelling than strategy (II). Sceptics motivated by externalist concerns about reliability may find strategy (II) more compelling than strategy (I).

Is scepticism of one kind or another irresistible? Of course, we can laugh off the sceptical arguments, refusing to take them seriously and continuing to believe what we always believed. But that sounds like a mere refusal to do one's job as a philosopher, or one's duty as a reflective person. A more interesting suggestion is that although sceptics speak truly, they do so by creating special contexts of utterance in which their words express truths: by making sceptical scenarios discursively salient they drive up the contextually relevant standards for the correct application of the term 'know' to a point at which they apply to virtually nothing, even though they have wide application as used in everyday contexts. Thus in a conversation about travel someone says truly 'Mary knows that the train leaves at noon', while simultaneously in an epistemology class someone else is saying of the same person, equally truly, 'Mary does not know that the train leaves at noon', for sceptical scenarios are relevant only in the latter context. According to *contextualism*, sentences involving 'know' express different propositions with different truth-values as used in different contexts, because the reference of 'know' varies with the context of utterance, even if the

reference of every other constituent is held fixed.[12] A similar contextualist line might be taken about other epistemic terms, such as 'justified'.

Once sceptics have manipulated the context, in the epistemology seminar, contextualists are apt to console themselves with the thought that although most denials of "knowledge" in that context of scepticism are correct, in everyday contexts many assertions of "knowledge" are also correct. For example, although 'I do not know that there is a whiteboard in the room' expresses a truth as uttered in the seminar, 'He knows that there is a whiteboard in the room' expresses a truth as uttered simultaneously outside the seminar about the same person. But that thought underestimates the gravity of the situation in which the sceptic has put contextualists, on their own analysis. For since 'know' is a factive verb, the truth of 'He knows that there is a whiteboard in the room', as uttered outside the seminar, requires the truth of 'There is a whiteboard in the room' as uttered outside the seminar and therefore as uttered inside the seminar too (since no context-dependence in the latter sentence is relevant here). Thus the consoling thought commits them to claims such as 'There is a whiteboard in the room', while their contextualism commits them to claims such as 'I do not know that there is a whiteboard in the room'. Consequently, in the seminar, they are committed to this:

(MK) There is a whiteboard in the room and I do not know that there is a whiteboard in the room.

But (MK) is Moore-paradoxical: although (MK) could be true, it is somehow self-defeating to assert (MK). Thus contextualists are not entitled to the consoling thought. In the seminar, they should not say that 'He knows that there is a whiteboard in the room' expresses a truth outside the

seminar. They might say 'I do not know that the sentence "He knows that there is a whiteboard in the room" does not express a truth outside the seminar'. But since they must also admit 'I do not know that the sentence "He knows that there is a whiteboard in the room" expresses a truth outside the seminar', they are not identifying some positive feature of the everyday situation; they are simply confessing ignorance of its status. A similar argument applies to contextualism about justification. At least in the seminar, contextualists are at the sceptic's mercy and cannot look outside for help.[13]

Should we surrender as much as contextualists do to the sceptic's arguments? Let us first investigate how radical those arguments are. They are often assumed not to touch beliefs about one's own mental states. For example, if I feel dizzy, I know that I feel dizzy, even if I do not know that there is a whiteboard in the room. It is the status of one's beliefs about the world external to one's present state of mind that is supposed to be in question. But it is not obvious that the sceptic's styles of argument are so limited.

Using strategy (I), I start by testing my belief that I feel dizzy. I suspend it, and others relevantly related to it, but then immediately recover them by my usual methods for forming beliefs about my current mental states. Then, however, I decide to test my belief that my methods for forming beliefs about my current mental states are reliable. Once I have suspended that belief, and others relevantly related to it, and the use of the methods themselves, can I really recover them from what remains? *Perhaps* I can, by some clever philosophical argument. But then I decide to test my belief that my methods of abstract argument are reliable. After all, it is not as though most clever philosophical arguments turn out to be sound. Once I have suspended my belief that my methods of abstract argument are reliable, and others relevantly related to it, and the use of the methods themselves, how can I recover them from what remains? Not by a clever philosophical argument,

for the use of such an argument is not part of what remains, on pain of begging the question. Thus strategy (I) has no internal limitation that would protect my belief that I am feeling dizzy. If moderate sceptics apply strategy (I) to beliefs about the external world but not to beliefs of the kinds just discussed, they seem to be using a method of argument when they like its conclusions and avoiding it when they do not. Such opportunism is less intellectually impressive than the willingness of extreme sceptics to deny that I know, or even have a justified true belief, that I am feeling dizzy.

With strategy (II), my belief that I feel dizzy looks at first sight better off, for how could one construct a sceptical scenario in which I appear to myself to feel dizzy without actually feeling dizzy? However, let SS be a sceptical scenario in which I am a brain in a vat but appear to myself to be walking on a sandy beach; in SS, I neither feel dizzy nor appear to myself to feel dizzy nor believe myself to feel dizzy; in SS, I believe that I am not in SS. How things appear to me in SS is utterly different from how they actually appear to me. In actuality, as well as in SS, I believe that I am not in SS. Both in actuality and in SS, my belief that I am not in SS is falsity-indicative, for exactly the same reason that applies to any other belief that one is not in a sceptical scenario. My belief that I am not in SS lowers the probability that I am not in SS. The crucial point for strategy (II) is that I believe in the sceptical scenario that I am not in it, whatever my actual mental state.[14] Thus if not being falsity-indicative is necessary for a belief to constitute knowledge, I do not know that I am not in SS. Similarly, if not being falsity-indicative is necessary for a belief to be justified, my belief that I am not in SS is not justified. But that I feel dizzy entails that I am not in SS. I believe that I am not in SS in the light of that competent deduction. Consequently, by CLOSURE(K) and CLOSURE(JB), I neither know nor have a justified belief that I feel dizzy. Thus my belief that I feel dizzy is vulnerable to strategy (II).

Moderate sceptics will protest that they have been misinterpreted. In effect, they concede that a falsity-indicative belief *can* constitute knowledge, or be justified. So even for a sceptical scenario that matches one's actual situation in its appearance to one, they cannot claim that its agreed feature, that in it one believes that one is not in it, by itself prevents one from knowing or having a justified true belief that one is not in it. If such knowledge or justified true belief is blocked, it must be for a more specific reason.

A natural thought is that I can *discriminate* between the actual situation and SS, but cannot discriminate between the actual situation and a sceptical scenario SS* in which I am a brain in a vat but the world appears to me just as it actually does. For example, both in the actual situation and in SS* I feel dizzy; in SS I do not feel at all dizzy. Of course, if my mental state in the actual situation differed from my mental state in SS only in some respect that was concealed from me, that would not enable me to discriminate between the actual situation and SS. The point is that I am aware of whether I feel dizzy. Both in the actual situation and in SS*, I know that I feel dizzy. In SS, I know that I do not feel dizzy. Even if I did not know in SS that I did not feel dizzy, knowing in the actual situation that I feel dizzy enables me to know that I am not in SS, for I know by definition of 'SS' that in SS I do not feel dizzy, and from the premises that I feel dizzy and that in SS I do not feel dizzy I can competently deduce the conclusion that I am not in SS. Given that I make the deduction in my actual situation, by CLOSURE(K) I know that I am not in SS.

By parallel reasoning, if I do know in the actual situation that I have hands, that enables me to know that I am not in SS*, for I know by definition of 'SS*' that in SS* I lack hands, and from the premises that I have hands and that in SS* I lack hands I competently deduce the conclusion that I am not in SS*. By CLOSURE(K), I know that I am not in SS*.[15] But the moderate sceptic

who allows that I know that I feel dizzy does not allow that I know that I have hands. Why is my belief that I feel dizzy supposed to be epistemically better off than my belief that I have hands? The salient difference between them is that we have been given a possible situation in which I falsely believe that I have hands but no possible situation in which I falsely believe that I feel dizzy. If the sceptic is assuming that the mere possibility of error precludes knowledge, why should we accept that premise?

We should not take for granted that even beliefs about one's current mental state are infallible. Someone feels dizzy, is aware of doing so and complains about it to others, who react with concern. Gradually the dizzy feeling wears off, but he is enjoying all the attention too much to notice; he needs to believe that it is deserved. If he now asks himself whether he feels dizzy, he still answers sincerely that he does; he believes falsely that he feels dizzy. The story contains no impossibility. Of course, if one believes truly that one feels dizzy, then no sceptical scenario is possible in which one is in exactly the same mental state but believes falsely that one feels dizzy, for feeling dizzy is an aspect of one's mental state. But if one's own current mental state is not perfectly transparent to one, then sceptical scenarios are relevant in which one is in a slightly different mental state and falsely believes that one feels dizzy. Confirmation of this lack of transparency comes from the non-transitivity of indiscriminability. When my mental state changes only gradually, I go through a long series of mental states M_0, M_1, M_2, ..., M_n, such that for no i can I discriminate between M_i and M_{i+1} (as presented to me at those times): I notice no change from one state to the next. Nevertheless, I can easily discriminate between M_0 and M_n; they are manifestly different. It follows that M_i and M_{i+1} are qualitatively distinct for at least some i: for since qualitative identity is a transitive relation, if M_i and M_{i+1} were qualitatively identical for every i then M_0 and M_n would be

qualitatively identical too, which they are not. Therefore, at least two mental states are qualitatively distinct but indiscriminable (as presented to the subject at the relevant times). Thus there are unnoticeable differences between mental states.[16]

Why are sceptics often reluctant to apply their arguments to knowledge of one's own current mental states? Start from the notion of rationality, as a criterion of praise and blame. A central requirement of rationality is that one should proportion one's belief to the evidence. It is praiseworthy to do so, blameworthy to do otherwise. Rational and irrational agents are responsible to the standard of rationality. But if agents are not in a position to know whether they are conforming to a given standard, it seems unfair to hold them responsible to that standard.[17] Thus it is natural to suppose that people are in a position to know whether they are conforming to the standard of rationality. In sceptical scenarios, subjects are still culpable for their irrationality (if any), even though they are blameless for their ignorance. In particular, they should be in a position to know whether they are proportioning their belief to their evidence, and therefore in a position to know what their evidence is, even in a sceptical scenario. Consequently, their evidence cannot be what would count as evidence in a normal scientific or legal context — the result of an experiment, the testimony of a witness — for one can be blamelessly deceived about such matters in a sceptical scenario. Only one's current mental states seem to have a chance to survive this filtering process. Thus they come to occupy a privileged role as all that is left to count as evidence. They are what other beliefs are measured against. Once the evidence base is attenuated that far, it is hardly surprising that external world scepticism should look plausible. But not even one's own current mental states really meet the exacting condition. For reasons of the kind noted above, one is not always in a position to know whether one is in a given mental state, for example, whether one feels

dizzy.

Should we conclude that we have no evidence at all? A more plausible alternative is to be less purist about praise and blame: we can be responsible to a standard even if we are sometimes not in a position to know whether we are conforming to it. Thus we need not always be in a position to know what our evidence is; that is consistent with the claim that *p* is part of one's evidence only if one knows *p*, for one is not always in a position to know what one's knowledge is. What would count as evidence in a scientific or legal context may therefore count as genuine evidence after all. We refuse to be stampeded by sceptical scenarios into attenuating our evidence base according to the sceptic's demands.[18]

The proposed view is inconsistent with various forms of internalism, which guarantee us special epistemic access to our own states with respect to specified matters of justification. Whatever those states are, we are not always in a position to know whether we are in them. But the argument does not refute weaker internalist theses, according to which we are *typically* in a position to know whether we are in the relevant states, and the knowledge may be of a special nonobservational kind. Nor does the argument show that justified belief is not necessary for knowledge.[19]

The terminology of 'internalism' and 'externalism' suggests a Cartesian picture, on which the mind forms an inner world, to which it has direct access, while access to the external world is at best indirect. Given that picture, one can debate whether matters of justification belong wholly to the inner world. But the prior question is whether the picture is correct. On an alternative picture, paradigmatic mental states include cognitive and conative relations to independently existing objects in one's environment: loving honey, seeking honey, seeing honey, tasting honey, remembering

honey. Since mental states have non-mental constituents (such as honey), they do not constitute an inner world in any useful sense. Believing that there is honey in the hive is a mental state that involves a relation to one's environment, for one's belief to be about honey; knowing that there is honey in the hive is a mental state that involves a further relation to the environment, since one can be in it only if there is honey in the hive. Like belief states, knowledge states can play an essential role in causal explanations of action.[20] Although one is not always in a position to know whether one knows, for virtually no mental state is one always in a position to know whether one is in it. On this view, one should not concede to the sceptic that one could be in one's current total mental state even if most of one's beliefs were false, for if one knows *p* and knowing *p* is a mental state, one could not be in one's current total mental state without knowing *p*.

When sceptics try to *argue* that one does not know (or does not have a justified belief) that one is not in a sceptical scenario, they tacitly assume deeply problematic epistemological principles. Scepticism's great strength is that sometimes we ourselves cannot see how honestly to dissent. It can feel like mere empty bravado to insist that one does know, or have a justified true belief, that one is not in the sceptical scenario. However firmly one wishes to resist the sceptics, one may be betrayed from within by one's own intuitions. Of course, at other points in the dialectic sceptics will insist that mere intuitions have no rational force. But no inconsistency in sceptics' treatment of intuitions shows anti-sceptics how to live with the sceptical tendency of their own intuitions. Sceptics are troublemakers who can disrupt our position without having a coherent position of their own, by presenting us with considerations to which we cannot find a response that we find satisfying.[21] If they are sick, they infect us with their sickness. Although some people have more natural immunity than others, probably few epistemologists feel no conflict at all within themselves between sceptical

and anti-sceptical tendencies.

What should our attitude be to our sceptical intuitions, if we have them? It may help to consider other cases in which intuitions conflict. What should someone do in a practical dilemma? Presenting the case with one emphasis elicits the strong intuition that she should f ; presenting the same case with another emphasis elicits the strong intuition that it is not the case that she should f . A contextualist might say that the sentence 'She should f ' expresses a true proposition as uttered in the context of the first presentation but a different and false proposition as uttered in the context of the second presentation. Yet that attempted resolution of the conflict feels quite unconvincing. In the context of the first presentation, it is intuitive that the sentence 'She should f ' expresses a true proposition *even as uttered in the context of the second presentation*; in the context of the second presentation, it is intuitive that the same sentence expresses a false proposition *even as uttered in the context of the first presentation*. The agent must choose between f ing and not f ing. The practical question is whether the proposition *s*f that 'I should f ' expresses in her context is true. A speaker in another context who uses 'She should f ' to express a proposition that does not match *s*f in truth-value thereby loses touch with the practical question. For that question, the agent's context has primacy. To focus on the practical question, one lets 'She should f ' express a proposition materially equivalent to *s*f . Consequently, in different contexts with a focus on the practical question 'She should f ' expresses propositions materially equivalent to each other, and the relevant form of contextualism fails.

One might suppose that 'know', unlike 'should', has no practical aspect. Arguably, however, the question 'Do I know *p*?' is closely related to the practical questions 'Should I assert *p*?' and 'Should I believe *p*?' (Williamson 2000a). Perhaps this practical aspect of 'know' gives

primacy to the subject's context: if third-person or past-tense 'knowledge' ascriptions must express propositions materially equivalent to that expressed by the first-person present-tense 'I know *p*' as uttered by the subject, then they must express propositions materially equivalent to each other, and again the relevant form of contextualism fails.

Contextualism supplies a perfectly general strategy for resolving any apparent disagreement whatsoever. Since some disagreements are genuine, we should not always follow that strategy. The conflict of intuitions does not always disappear on further reflection. At least some intuitions are mistaken.[22] Moreover, they are explicably, not blankly, mistaken. Whether the agent should f depends on the balance between many complex, unquantifiable, subtly interacting considerations, some genuinely weighing one way, others genuinely weighing the other way. The concept of what should be done provides no algorithm for weighing all these factors against each other or integrating them into a final verdict. It is no surprise that one can make each verdict intuitive by highlighting considerations that favour it; what else would one expect? The cases need not even be borderline. For even if the considerations on one side heavily outweigh those on the other, a skilled presenter can still present the outweighed considerations so strikingly that they appear to outweigh the others.

Surely the same phenomenon can occur for epistemic concepts too. Whether one knows (or one's belief is epistemically justified) also depends on the balance between many complex, unquantifiable, subtly interacting considerations. In many realistic cases, some weigh one way, others the other way. The concept of knowledge or of epistemic justification provides no algorithm for weighing all these factors against each other or integrating them into a final verdict. No wonder that the skilled sceptic can present the considerations that favour a negative verdict so vividly that they intuitively appear to outweigh the considerations on the other side. It does not follow that the

sceptic is right, even in the context of the epistemology seminar; the case may not even be borderline. Nor does it follow that the sceptic is wrong, even in the everyday context. As before, the intuitions that predominate in one context spill over to judgements about the truth-values of sentences as uttered in the other context. In the everyday context, it is intuitive that someone in the epistemology class who says 'Mary does not know that the train leaves at noon' is overestimating her epistemic difficulties. In the epistemology class, it is intuitive that someone in the everyday context who says 'Mary knows that the train leaves at noon' is underestimating her epistemic difficulties. Although such data are not decisive against contextualism, they tend to support the non-contextualist explanation.

If we are sensitive to many complex considerations when we assess matters of knowledge or justification, does it follow that our concept of knowledge or our concept of justification is analysable in terms of our concepts of those complex considerations? No. Suppose that we grasped the concept of knowledge or of justification simply by means of examples of situations in which someone knows or is justified and of the opposite. When we assessed matters of knowledge or justification, we should therefore compare the case at hand with our paradigms of knowledge or justification and of ignorance or lack of justification. The comparative similarity would depend on many complex considerations, to which we should need to be sensitive, but our concept of knowledge or of justification would not be analysable in terms of our concepts of those complex considerations. The very act of comparison might cause us to notice similarities and differences that we had not previously conceptualised. If recognising a face does not require conceptual analysis, why should recognising knowledge or justification do so?

No doubt a purely paradigmatic account of epistemic concepts is too simple. For example,

it seems essential to grasping the concept of knowledge that one appreciate the entailment from *S knows that P* to *P*. Nevertheless, paradigms involving perception, memory, testimony and deduction may play a large role in our grasp of the concept of knowledge. That is not to say that the concepts of perception, memory, testimony and deduction are in any sense prior to the concept of knowledge: without some grasp of epistemic considerations, one would be unable to appreciate the difference between the successful, knowledge-providing cases of perception, memory, testimony and deduction and the failures. If all that sounds messy, it may be the normal position for empirically applicable abstract concepts such as *time, causation, body, mind, meaning, society, intention, choice,* and *government*.

When Gettier showed that knowledge is more than justified true belief, many epistemologists reacted by trying to construct an alternative analysis of the concept of knowledge. Since knowledge implies true belief, but not *vice versa*, they sought whatever it is that, added to true belief, makes knowledge. Some proposed further constraints on justification: for example, that the premises as well as the conclusion of a justifying argument should be true. Others proposed more external constraints on the relation between the belief state and the truth of the belief: for example, a causal, reliable, lawlike or counterfactual connection. Each particular analysis is generally agreed by a majority of epistemologists to be refuted by counterexamples; it does not even give necessary and sufficient conditions for knowledge. Some continue to seek. This was the usual fate of programmes for the analysis of concepts in twentieth century philosophy, which is somewhat surprising on the assumption that we grasp such analyses in grasping the concepts. But why accept that assumption? The regress of analysis must stop somewhere, presumably; why should it not stop with concepts like the concept of knowledge? There need be no concept that, added to the concept

of true belief, makes the concept of knowledge. Quite generally, if the concept C applies only where the concept C* applies, and not *vice versa*, it does not follow that the concept C is analysable in terms of the concept C* and other concepts. *Same person* applies only where *weighs the same* applies, and not *vice versa*, but the concept of personal identity is not to be analysed in terms of the concept of identity of weight and other concepts. Although nobody has proved that the concept of knowledge is not a complex combination of simpler concepts, we have no good reason to think that it is such a combination.[23]

Suppose that we do apply the concept of knowledge on the basis of something more like a fallible recognitional capacity than a conceptual analysis. Our susceptibility to both sceptical and anti-sceptical intuitions demonstrates that our willingness to apply the term 'know' varies strongly with context. Contextualists try to reconcile the intuitions by postulating that what relation we use the term to refer to varies accordingly; we have seen reason to doubt their account. A more pessimistic interpretation is that the term 'know' expresses no coherent concept: there is no relation that the associated "recognitional capacity" is a capacity to recognise. We "recognise" that we know the truth of various everyday propositions, we "recognise" that we do not know that we are not in a sceptical scenario in which those propositions are false, and we "recognise" that knowledge is closed under competent deduction. The constraints are all built into the concept of knowledge, but no relation could satisfy them all. The pessimist could argue on similar grounds that even the epistemic term 'justification' expresses no coherent concept.[24]

The pessimist seems to rely on an over-simple view of how recognitional capacities determine reference. For example, our most primitive forms of reference to some spatial properties and relations of shape, size and distance are presumably determined by our corresponding

perceptual recognitional capacities, even though, notoriously, the structure of our visual system renders us susceptible to various spatial illusions and paradoxes. No assignment of reference would make all the perceptual appearances veridical, but we are not tempted by the conclusion that our most primitive spatial terms do not express coherent concepts. Some spatial properties and relations play a sufficiently large role in the explanation of the perceptual appearances for us to count as referring to them, even though they falsify some of those appearances; the structure of our perceptual systems can take up the explanatory slack. Analogously, the relation of knowing may play a sufficiently large role in the explanation of our epistemic intuitions for us to count as referring to it, even though it falsifies some of those intuitions; perhaps the structure of the mechanisms underlying our epistemic judgements can take up the explanatory slack. The intuition that one cannot know that one is not in a sceptical scenario may thus be a conceptual illusion. Of course, we want a good reason to blame that as the conceptual illusion, rather than blaming the intuition that many ordinary people know the truth of many ordinary propositions, or the intuition that knowledge is closed under competent deduction. But the reason might be that classifying the sceptical intuitions as illusions involves postulating far less extensive illusions than would be involved in classifying the anti- or non-sceptical intuitions as illusions.[25]

We might go further by conjecturing a constraint on an assignment of reference: that in normal circumstances it should tend to make our perceptual beliefs (those in which we take perceptual appearances at face value) count as knowledge (not just as true). Such a constraint does not amount to a cast-iron guarantee that our perceptual beliefs are not massively false; a coin with the tendency to land heads and tails equally often may happen to land heads every time it is tossed. We lack perceptual knowledge only if we are very unlucky; the sceptic was still wrong to argue that

we could not have had perceptual knowledge. Analogously, we may conjecture, a constraint on an assignment of reference is that in normal circumstances it should tend to make our intuitive beliefs (those in which we take conceptual appearances at face value) count as knowledge. In particular, it should tend to make our intuitive epistemological beliefs count as epistemological knowledge. They cannot all count, but the exceptions should as far as possible be minimised.[26] The sceptic was wrong to argue that we could not have known that we had perceptual knowledge. Although general constraints on interpretation do not explain how we know in any particular case, they do suggest that the odds are stacked against the sceptic.[27]

Notes

1 The conditional probability Prob(p | q) may be defined in terms of unconditional probabilities as Prob(p & q)/Prob(q) when Prob(q) is non-zero. Note that, in the sense defined in the text, q may be evidence for p even if q itself is unknown or even false: the relation between p and q is a purely conditional one.

2 Proof: If Prob(p | ~Bp) = 1 then Prob(p | Bp) = Prob(Bp)Prob(p | Bp) + (1 Prob(Bp))Prob(p | Bp) Prob(Bp)Prob(p | Bp) + Prob(~Bp)Prob(p | ~Bp) = Prob(p). The result holds even on the weaker assumption that Prob(p | Bp) Prob(p | ~Bp).

3 The description 'the sceptical scenario' is not being defined, circularly, in terms of the condition 'I believe that I am not in the sceptical scenario'. That in the sceptical scenario I believe that I am not in it is a consequence, not a part, of the definition.

4 In a much-discussed application of semantic externalism, Putnam (1981: 1-21) argued that causal constraints on reference ensure that the brain in the vat cannot entertain the proposition that it is not a brain in a vat. However, even if the argument works for some sceptical scenarios, it seems to fail for others, such as that of a recently envatted brain (Smith 1984). For a recent exchange on semantic externalist replies to scepticism and

further references see DeRose 2000 and Williamson 2000b.

5 The poor correlation between $\sim s$ and $B\sim s$ is more usually articulated by means of the counterfactual conditional $s \mathrel{\square\!\!\rightarrow} B\sim s$ is true (if I were in the sceptical scenario, I would believe that I was not in it). Nozick's influential counterfactual analysis of knowledge makes the truth of $\sim p \mathrel{\square\!\!\rightarrow} \sim Bp$ (or a complex variant thereof) necessary for knowing p, so the falsity of $s \mathrel{\square\!\!\rightarrow} \sim B\sim s$ precludes knowing $\sim s$ (1981: 172-288). Nozick's analysis is subject to counterexamples (see the essays in Luper-Foy 1987, Williamson 2000a: 147-63 and references therein). The use of conditional probabilities shows that a similar problem for knowing $\sim s$ can be articulated without counterfactuals. One advantage of the condition $\mathrm{Prob}(p \mid Bp) \gg \mathrm{Prob}(p)$ over $\sim p \mathrel{\square\!\!\rightarrow} \sim Bp$ is that the former holds while the latter fails in cases involving a correct judgement of quantity with a large margin for error by a largely reliable subject who nevertheless has a systematic tendency to very slight misjudgements; intuitively, that tendency does not justify a sceptical verdict in such cases (Williamson op. cit.).

6 The point is close to the problem of old evidence in Bayesian epistemology: if $\mathrm{Prob}(e) = 1$ then $\mathrm{Prob}(h \mid e) = \mathrm{Prob}(h)$, where e is the evidence and h an hypothesis. See Earman 1992: 119-35, Howson and Urbach 1993: 403-8, Maher 1996 and Williamson 2000a: 220-1.

7 See Quine 1969 and Kornblith 1994.

8 Dretske 1970 and Nozick 1981 contain scepticism by rejecting the closure principle. On Nozick's analysis, one can even know *p & q* without knowing *p*. Non-closure is currently unpopular. Contextualists (see below) typically argue that closure holds in any given context, but may appear to fail if the context changes in the course of the argument (DeRose 1995).

9 For examples of this marginalisation of knowledge see Craig 1990: 272, Earman 1993: 37, Kaplan1985 and Wright 1991: 88.

10 Unger (1975: 197-249) and Williamson (2000a: 184-208) argue that justification depends on knowledge.

11 Some hold that a belief is justified if and only if it has a probability greater than some threshold constant *c* less than 1 (for example, 90%), on some appropriate probability distribution. They will object to CLOSURE(JB) that sometimes p_1 and p_2 each have probabilities greater than *c* while their conjunction p_1 & p_2 has probability less than *c*. On their view, someone who both believes the conjuncts and believes the conjunction on the basis of competent deduction from the conjuncts has justified belief in the conjuncts but not in the conjunction (see Kyburg 1970). Consequently, they would restrict CLOSURE(JB) to single-premise deductions (if *q* is a logical consequence of *p*, the probability of *q* is at least as high as the probability of *p*). This restriction does not affect the applications of CLOSURE(JB) in the text, which are to single-premise deductions, in both the Gettier case

and the sceptical argument. Of course, the probabilistic standard for justified belief has highly questionable consequences: one may have justified belief in each member of an obviously inconsistent set of propositions: for example, 'Ticket number *n* will not win the lottery' for each natural number $n < 1{,}000{,}000$ and 'For some natural number $n <$ 1,000,000, ticket number *n* will win the lottery'. Another putative objection to CLOSURE(JB) is that one might be justified in believing various complex propositions (perhaps on the basis of testimony) but then competently deduce a contradiction from them and somehow get oneself to believe it on that basis (perhaps aided by a paraconsistent logician) without having a justified belief in the contradiction. But this objection misreads CLOSURE(JB) as a diachronic principle. Its intended reading is synchronic: if at time *t* one has justified belief in $p_1, ..., p_n$ and at *t* one believes *q* on the basis of competent deduction of *q* from $p_1, ..., p_n$ then at *t* one has justified belief in *q*. Once one has deduced the contradiction, one no longer has justified belief in all the premises (modulo the issue raised in the first objection). For difficulties in formulating logical constraints on rationality see Harman 1986.

12 Unger 1986, Cohen 1988, DeRose 1995, 1999 and Lewis 1996 propose versions of contextualism; Schiffer 1996 offers criticism. For two recent exchanges see Feldman 2001, Cohen 2001, Williams 2001 and Williamson 2001. See Pryor 2001 for a brief survey of different forms of contextualism. The view that the application of 'know' depends on the context in which the subject of the ascription (rather than the ascriber) is thinking or speaking (Williams 1996) is not contextualism in the present sense. Although the subject

and the ascriber are identical in first-person ascriptions, the two roles are still distinct. Contextualism has predecessors in accounts on which knowing *p* is a matter of ruling out (in some sense to be explained) relevant alternatives in which *p* is false. (Goldman 1976, Dretske 1981).

13 See also Williamson 2001.

14 Schiffer 1996 uses a related example.

15 See Klein 1995 for related discussion.

16 The argument is developed more rigorously in Williamson 2000a: 93-134.

17 Compare Bonjour 1985.

18 See Williamson 2000a: 164-208 for this view, and Fumerton 2000 for a critique.

19 For versions of internalism see Bonjour 1985 and Fumerton 1995. A recent critique is Goldman 1999.

20 See Williamson 2000: 21-92 on knowing as a mental state. McDowell 1995 and Gibbons 2001 defend similar views. The idea is widespread that the contents of propositional

attitudes depend constitutively on the subject's environment (Burge 1979; compare Putnam 1973). The proposed view of knowing extends the environmental dependence to the attitudes to those contents.

21 Compare Feyerabend 1978: 143.

22 I ignore various forms of relativism.

23 Shope 1983 surveys earlier post-Gettier attempts to analyse knowing.

24 Schiffer 1996 proposes a form of pessimism.

25 On the epistemology of intuition see Bealer 1999 and DePaul and Ramsey 1998.

26 Compare Davidson 1986. The idea that belief aspires to the condition of knowledge (not just of truth) is sketched in Williamson 2000a: 41-8.

27 This chapter shamelessly favours my own prejudices and interests. Attempting nothing like a comprehensive survey of contemporary epistemology, I have left much significant work unmentioned. For a sense of the variety of contemporary epistemology see Greco and Sosa 1999. A compact survey is Pryor 1999. For most of the current approaches to the problem of scepticism see DeRose and Warfield 1999. Thanks to Alexander Bird and Nico Silins

for comments on an earlier draft. Exchanges with Keith DeRose and John Hawthorne have also been very helpful.

Bibliography

Bealer, G. 1999. 'The a priori', in Greco and Sosa 1999.

Bonjour, L. 1985. *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.

Burge, T. 1979. 'Individualism and the mental'. *Midwest Studies in Philosophy*, 4: 73-121.

Cohen, S. 1988. 'How to be a fallibilist'. *Philosophical Perspectives*, 2: 91-123.

Cohen, S. 2001. 'Contextualism defended: comments on Richard Feldman's "Skeptical problems, contextualist solutions"'. *Philosophical Studies*, 103: 61-85.

Craig, E. 1990. 'Three new leaves to turn over'. *Proceedings of the British Academy*, 76: 265-81.

Davidson, D. 1986. 'A coherence theory of truth and knowledge', in E. LePore (ed.), *Truth and Interpretation*. Oxford: Blackwell.

DePaul, M., and Ramsey, W. (eds.) 1998. *Rethinking Intuition: The Psychology of Intuition and its Role in Philosophical Inquiry*. Lanham, MD: Rowman & Littlefield.

DeRose, K. 1995. 'Solving the skeptical problem'. *Philosophical Review*, 104: 1-52.

DeRose, K. 1999. 'Contextualism: an explanation and defense', in Greco and Sosa 1999.

DeRose, K. 2000. 'How can we know that we're not brains in vats?'. *Southern Journal of Philosophy*, 38, supplement: 121-48.

DeRose, K., and Warfield, T. 1999. *Skepticism: A Contemporary Reader*. Oxford: Oxford University Press.

Dretske, F. 1970. 'Epistemic operators'. *Journal of Philosophy*, 67: 1007-23.

Dretske, F. 1981. 'The pragmatic dimensions of knowledge'. *Philosophical Studies*, 40: 363-78.

Earman, J. 1992. *Bayes or Bust?* Cambridge, MA: MIT Press.

Earman, J. 1993. 'Underdetermination, realism and reason'. *Midwest Studies in Philosophy*, 18: 19-38.

Feldman, R. 2001. 'Skeptical problems, contextualist solutions'. *Philosophical Studies*, 103: 87-98.

Feyerabend, P. 1978. *Science in a Free Society.* London: NLB.

Fumerton, R. 1995. *Metaepistemology and Skepticism*. Lanham, MD: Rowman & Littlefield.

Fumerton, R. 2000. 'Williamson on skepticism and evidence'. *Philosophy and Phenomenological Research*, 60: 629-35.

Gettier, E. 1963. 'Is justified true belief knowledge?' *Analysis*, 23: 121-3.

Gibbons, J. 2001. 'Knowledge in action'. *Philosophy and Phenomenological Research*, 62: 579-600.

Goldman, A. 1976. 'Discrimination and perceptual knowledge'. *Journal of Philosophy*, 73: 771-91.

Goldman, A. 1999. 'Internalism exposed'. *Journal of Philosophy*, 96: 271-93.

Greco, J., and Sosa, E. (eds.) 1999. *The Blackwell Guide to Epistemology*. Oxford: Blackwell.

Harman, G. 1986. *Change in View: Principles of Reasoning*. Cambridge, MA: MIT Press.

Howson, C., and Urbach, P. 1993. *Scientific Reasoning: The Bayesian Approach*, 2nd edn. Chicago: Open Court.

Kaplan, M. 1985. 'It's not what you know that counts'. *Journal of Philosophy* 82: 350-63.

Klein, P. 1995. 'Skepticism and closure: why the evil genius argument fails'. *Philosophical Topics*, 23: 213-36.

Kornblith, H. (ed.) 1994. *Naturalizing Epistemology*, 2nd edn. Cambridge MA: MIT Press.

Kyburg, H. 1970. 'Conjunctivitis', in M. Swain (ed.), *Induction, Acceptance, and Rational Belief*. Dordrecht: Reidel.

Lewis, D. 1996. 'Elusive knowledge'. *Australasian Journal of Philosophy*, 74: 549-67.

Luper-Foy, S. 1987. *The Possibility of Knowledge: Nozick and his Critics*. Totowa, NJ: Rowman & Littlefield.

Maher, P. 1996. 'Subjective and objective confirmation'. *Philosophy of Science*, 63: 149-74.

Nozick, R. 1981. *Philosophical Explanations*. Oxford: Oxford University Press.

Pryor, J. 2001. 'Highlights of recent epistemology'. *British Journal for the Philosophy of Science*, 52: 95-124.

Putnam, H. 1973. 'Meaning and reference'. *Journal of Philosophy*, 70: 699-711.

Putnam, H. 1981. *Reason, Truth and History*. Cambridge: Cambridge University Press.

Quine, W. 1969. 'Epistemology naturalized', in *Ontological Relativity and Other Essays*. New York: Columbia University Press.

Schiffer, S. 1996. 'Contextualist solutions to scepticism'. *Proceedings of the Aristotelian Society*, 96: 317-33.

Shope, R. 1983. *The Analysis of Knowing: A Decade of Research*. Princeton: Princeton University Press.

Smith, P. 1984. 'Could we be brains in a vat?'. *Canadian Journal of Philosophy*, 14: 115-

23.

Unger, P. 1975. *Ignorance: A Case for Scepticism*. Oxford: Oxford University Press.

Unger, P. 1986. 'The cone model of knowledge'. *Philosophical Topics*, 14: 125-78.

Williams, M. 1996. *Unnatural Doubts.* Princeton: Princeton University Press.

Williamson, T. 2000a. *Knowledge and its Limits.* Oxford: Oxford University Press.

Williamson, T. 2000b. 'Skepticism, semantic externalism, and Keith's mom'. *Southern Journal of Philosophy*, 38, supplement: 149-58.

Williamson, T. 2001. 'Comments on Michael Williams' "Contextualism, externalism and epistemic standards"'. *Philosophical Studies*, 103: 25-33.

Wright, C. 1991. 'Scepticism and dreaming: imploding the demon'. *Mind*, 100: 87-116.