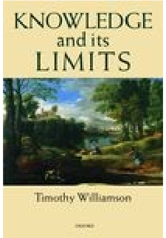

Knowledge and its Limits
Timothy Williamson
https://doi.org/10.1093/019925656X.001.0001
Published: 2002 **Online ISBN:** 9780191598678 **Print ISBN:** 9780199256563

Search in this book

CHAPTER

4 Anti-Luminosity

Timothy Williamson

https://doi.org/10.1093/019925656X.003.0005 Pages 93–113
Published: October 2002

Abstract

We are often conceived as cognitively at home with conditions that are *luminous* in roughly the sense that whenever they obtain we know or are in a position to know that they obtain; mental states such as feeling cold or pain are often thought to provide examples of luminous conditions. This chapter argues that there are no non-trivial luminous conditions, and therefore that we suffer from a kind of cognitive homelessness. The argument involves consideration of gradual processes in which small changes are below our level of discrimination. It is related to, but not the same as, sorites paradoxes, for example about how many grains make a heap. The result provides the basis for an objection to the attempt that Michael Dummett has made to characterize linguistic meaning in terms of assertability rather than truth.

Keywords: assertability, cognitive homelessness, discrimination, Dummett, luminous, meaningpain, paradoxes, sorites, truth

Subject: Epistemology, Metaphysics

Collection: Oxford Scholarship Online

4.1 Cognitive Homes

One source of resistance to the conception of knowing as a mental state is the idea that one is guaranteed epistemic access to one's current mental states. According to that idea, one must be in a position to know whether one is in a given mental state, at least when one is attending to the question. When one asks oneself whether one knows a given proposition, one is not always in a position to know the answer. Section 1.2 responded to the objection by arguing that many uncontentious examples of mental states are the same as knowing in this respect. Nevertheless, some are inclined to think that a central core of mental states must be different. If S belongs to that core, then whenever one attends to the question one is in a position to know whether one is in S. In that sense, knowing would not be a core mental state. This chapter argues that there is no central core of mental states in that special sense. That conclusion will be a corollary of a far more general result about the limits of knowledge.

There is a constant temptation in philosophy to postulate a realm of phenomena in which nothing is hidden from us. Descartes thought that one's own mind is such a realm. Wittgenstein enlarged the realm to everything that is of interest to philosophy.[1] That they explained this special feature in very different ways hardly needs to be said; what is remarkable is their agreement on our possession of a cognitive home in which everything lies open to our view. Much of our thinking—for example, in the physical sciences—must operate outside this home, in p. 94 ↳ alien circumstances. The claim is that not all our thinking could be like that.

To deny that something is hidden is not to assert that we are infallible about it. Mistakes are always possible. There is no limit to the conclusions into which we can be lured by fallacious reasoning and wishful thinking, charismatic gurus and cheap paperbacks. The point is that, in our cognitive home, such mistakes are always rectifiable. Similarly, we are not omniscient about our cognitive home. We may not know the answer to a question simply because the question has never occurred to us. Even if something is open to view, we may not have glanced in that direction. Again, the point is that such ignorance is always removable.

The aim of this chapter is to argue that we are cognitively homeless. Although much is in fact accessible to our knowledge, almost nothing is inherently accessible to it. However, it is first necessary to sharpen the issue, to make it more susceptible to argument.

4.2 Luminosity

As in previous chapters, it is convenient to frame the discussion in terms of conditions, which obtain or fail to obtain in various cases. A case depends on a subject (referred to by 'one'), a time (referred to by the present tense), and a possible world. Although conditions are expressed by sentential clauses, they are not propositions as the latter are usually conceived, just because they are open with respect to person, place, and perhaps other circumstances, too. We often use clauses like that, as in 'When it rains, it pours'. The domain of cases will be taken to include counterfactual as well as actual possibilities. Since the cases on which the arguments below rely are physically and psychologically feasible, issues about the bounds of possibility are not pressing.

Conditions are coarsely individuated by the cases in which they obtain: they are identical if they obtain in exactly the same cases. This raises a delicate issue when we say that someone knows that a condition C obtains, for C may be presented in different guises. Under which guise is C known to obtain? If the condition that one is drinking water is the condition that one is drinking H_2O, because they obtain in the same cases, it does not seem to follow that one knows that the condition that one is drinking water obtains if and only if one knows that the condition that one is drinking H_2O obtains, for one may not know that water is H_2O. Fortunately, in a context in which the only relevant presentation of the condition C is as the condition that

p. 95 one is F, knowing that C ↳ obtains can be identified with knowing that the condition that one is F obtains, which is in turn only trivially different from knowing that one is F. We can therefore often leave the reference to guises tacit.

We will also use the notion of being in a position to know. To be in a position to know *p*, it is neither necessary to know *p* nor sufficient to be physically and psychologically capable of knowing *p*. No obstacle must block one's path to knowing *p*. If one is in a position to know *p*, and one has done what one is in a position to do to decide whether *p* is true, then one does know *p*. The fact is open to one's view, unhidden, even if one does not yet see it. Thus being in a position to know, like knowing and unlike being physically and psychologically capable of knowing, is factive: if one is in a position to know *p*, then *p* is true. Although the notion of being in a position to know is obviously somewhat vague and context-dependent, it is clear enough for present purposes. The vagueness and context-dependence are in any case primarily the result of fudging in attempts to defend the views to be criticized below.

A condition C is defined to be *luminous* if and only if (L) holds:

(L) For every case α, if in α C obtains, then in α one is in a position to know that C obtains.

Since being in a position to know is factive, the converse of (L) holds for any condition C, so the conditional in (L) could just as well be a biconditional. The picture is that a luminous condition always shines brightly enough to make its presence visible. However, (L) does not say that C must obtain independently of our dispositions to judge that C obtains; for all (L) says, the condition might obtain in virtue of those dispositions.

A realm in which nothing is hidden is a realm in which all conditions are luminous. Our question is: what conditions, if any, are in fact luminous?

Some examples will help. Pain is often conceived as a luminous condition, in the sense that if one is in pain, then one is in a position to know that one is in pain (for a recent discussion see McDowell 1989). The definition of luminosity gives scope to finesse some of the more obvious objections to claims of this kind. Thus people who lack the concept of pain—perhaps because their concepts carve up the space of possible sensations in an alternative way—and so never know that they are in pain, may still count as being in a position to know that they are in pain. Perhaps more primitive creatures are sometimes in pain without possessing any concepts at all; if they count as not even being in a position to know that they are in pain, a counterexample to luminosity might still be avoided by a stipulation that the subject of a case must be a possessor of concepts.

p. 96 Two claims of luminosity are implicit in the following passage from Michael Dummett (1978: 131):

> It is an undeniable feature of the notion of meaning—obscure as that notion is—that meaning is *transparent* in the sense that, if someone attaches a meaning to each of two words, he must know whether these meanings are the same.[2]

Thus if two words have the same meaning for one, then one is in a position to know that they have the same meaning; if the words have different meanings for one, then one is in a position to know that they have different meanings. Dummett does not even make the qualification ‘in a position to’; what of a subject who has never compared the two words? The two claims of luminosity are genuinely distinct, for the premise that whenever a condition C obtains one is in a position to know that C obtains does not entail the conclusion that whenever C does not obtain one is in a position to know that C does not obtain. If whenever one is awake one is in a position to know that one is awake, it does not follow that whenever one is not awake one is in a position to know that one is not awake (such asymmetries are discussed in Chapter 8).

Strictly, of course, having the same meaning and having different meanings are contraries, not contradictories, since both require the words to be meaningful.

Other conditions for which luminosity is often claimed are those of the form: it appears to one that A. When there really is an oasis ahead, one may not be in a position to know that there really is an oasis ahead but, it is supposed, when there at least appears to one to be an oasis ahead, one must be in a position to know that there at least appears to one to be an oasis ahead.

4.3 An Argument Against Luminosity

Consider the condition that one feels cold. It appears to have about as good a chance as any non-trivial condition of being luminous. Nevertheless, there is reason to think that it is not really luminous at all. This section presents the argument, and section 4.6 generalizes it. Sections 4.4 and 4.5 discuss objections.

Consider a morning on which one feels freezing cold at dawn, very slowly warms up, and feels hot by noon. p. 97 One changes from feeling cold ↳ to not feeling cold, and from being in a position to know that one feels cold to not being in a position to know that one feels cold. If the condition that one feels cold is luminous, these changes are exactly simultaneous. Suppose that one's feelings of heat and cold change so slowly during this process that one is not aware of any change in them over one millisecond. Suppose also that throughout the process one thoroughly considers how cold or hot one feels. One's confidence that one feels cold gradually decreases. One's initial answers to the question 'Do you feel cold?' are firmly positive; then hesitations and qualifications creep in, until one gives neutral answers such as 'It's hard to say'; then one begins to dissent, with gradually decreasing hesitations and qualifications; one's final answers are firmly negative.

Let $t_0, t_1, \dots, t_n$ be a series of times at one millisecond intervals from dawn to noon. Let α_i be the case at t_i ($0 \leq i \leq n$). Consider a time t_i between t_0 and t_n, and suppose that at t_i one knows that one feels cold. Thus one is at least reasonably confident that one feels cold, for otherwise one would not know. Moreover, this confidence must be reliably based, for otherwise one would still not *know* that one feels cold. Now at t_{i+1} one is almost equally confident that one feels cold, by the description of the case. So if one does not feel cold at t_{i+1}, then one's confidence at t_i that one feels cold is not reliably based, for one's almost equal confidence on a similar basis a millisecond later that one felt cold is mistaken. In picturesque terms, that large proportion of one's confidence at t_i that one still has at t_{i+1} is misplaced. Even if one's confidence at t_i was just enough to count as belief, while one's confidence at t_{i+1} falls just short of belief, what constituted that belief at t_i was largely misplaced confidence; the belief fell short of knowledge. One's confidence at t_i was reliably based in the way required for knowledge only if one feels cold at t_{i+1}. In the terminology of cases, we have this conditional:

(1_i) If in α_i one knows that one feels cold, then in α_{i+1} one feels cold.

Note that (1_i) is merely a description of a stage in a specific process; it does not purport to be a general principle about feeling cold. Statement (1_i) is asserted for each i from 0 to n 1, which is not to say anything about cases other than $\alpha_0, \dots, \alpha_n$.

Suppose that the condition that one feels cold is luminous. Then in any case in which one feels cold, the condition that one feels cold obtains, so one is in a position to know that the condition that one feels cold obtains, so one is in a position to know that one feels cold; since by hypothesis one is actively considering p. 98 the matter, one therefore does know that one feels cold. We therefore have this conditional: ↳

(2_i) If in α_i one feels cold, then in α_i one knows that one feels cold.

Now suppose:

(3i) In α *i* one feels cold.

By modus ponens, (2_i) and (3_i) yield this:

(4i) In α *i* one knows that one feels cold.

By modus ponens, (1_i) and (4_i) yield this:

(3i+1) In α *i*+1 one feels cold.

The following is certainly true, for α0 is at dawn, when one feels freezing cold:

(30) In α0 one feels cold.

By repeating the argument from (3_i) to (3_{i+1}) *n* times, for ascending values of *i* from 0 to *n* 1, we reach this from (3_0):

(3n) In α *n* one feels cold.

But (3_n) is certainly false, for α *n* is at noon, when one feels hot. Thus the premises (1_0), . . . , (1_n 1), (2_0), . . . , (2_n 1), and (3_0) entail a false conclusion. Consequently, not all of (1_0), . . . , (1_n 1), (2_0), . . . , (2_n 1), and (3_0) are true. But it has been argued that (1_0), . . . , (1_n 1) and (3_0) are true. Thus not all of (2_0), . . . , (2_n 1) are true. By construction of the example, one knows that one feels cold whenever one is in a position to know that one feels cold, so (2_0), . . . , (2_n 1) are true if the condition that one feels cold is luminous. Consequently, that condition is not luminous. Feeling cold does not imply being in a position to know that one feels cold.

4.4 Reliability

Since (1_0), . . . , (1_n 1) are the key premises in the argument of the last section against luminosity, it is prudent to pause and reconsider the argument for (1_i).

The argument applies reliability considerations to degrees of confidence. These degrees should not be equated with subjective probabilities as measured by one's betting behaviour. For assigning a very high subjective probability to a false proposition does not by itself constitute any degree of unreliability at all, in

the sense relevant to knowledge. Suppose that draws of a ball from a bag have been made. The draws are ↳ numbered from 0 to 100. You have not been told the results; your information is just that on each draw *i*, the bag contained *i* red balls and 100*i* black balls. You reasonably assign a subjective probability of *i*/100 to the proposition that draw *i* was red (produced a red ball), and bet accordingly. You know that draw 100 was red, since the bag then contained only red balls, even if the proposition that draw 99 was red—to which you assign a subjective probability of 99/100—is false. That does not justify a charge of unreliability against you. Intuitively, for any *i* less than 100, your bets do not commit you to believing outright that draw *i* was red. Your outright belief may be just that the probability on your evidence that draw *i* was red is *i*/100, which is true. On draw 100, unlike the others, you can form the belief on non-probabilistic grounds that it was red. What incurs the charge of unreliability is believing a false proposition outright, not assigning it a high subjective probability.

What is the difference between believing *p* outright and assigning *p* a high subjective probability? Intuitively, one believes *p* outright when one is willing to use *p* as a premise in practical reasoning. Thus one may assign *p* a high subjective probability without believing *p* outright, if the corresponding premise in one's practical reasoning is just that *p* is highly probable on one's evidence, not *p* itself. Outright belief still comes in degrees, for one may be willing to use *p* as a premise in practical reasoning only when the stakes are sufficiently low. Nevertheless, one's degree of outright belief in *p* is not in general to be equated with one's subjective probability for *p*; one's subjective probability can vary while one's degree of outright belief remains zero. Since using *p* as a premise in practical reasoning is relying on *p*, we can think of one's degree of outright belief in *p* as the degree to which one relies on *p*. Outright belief in a false proposition makes for unreliability because it is reliance on a falsehood. The degrees of confidence mentioned in the argument for (1_i) should therefore be understood as degrees of outright belief.

The argument for (1_i) assumes that the underlying basis on which one believes that one feels cold changes at most slightly between t_i and t_{i+1}, for otherwise an error in the belief at t_{i+1} might not threaten the reliability of the belief at t_i. For example, if one believes inferentially at t_{i+1} and not at all inferentially at t_i, false belief at t_{i+1} might well be consistent with knowledge at t_i. Apparent gradualness in the process does not guarantee gradualness at the underlying level (Wright 1996: 937). Nevertheless, we can choose an example in which there is gradualness at the underlying level too, and that will suffice for a counterexample to (L). The basis on which one judges that one feels cold need not change suddenly as one gradually becomes colder.

p. 100 The invocation of reliability does not presuppose that whether one ↳ feels cold is independent of one's dispositions to judge that one does. Luminosity is often supposed to rest on a constitutive connection between the obtaining of the condition and one's judging it to obtain, but the effect of such a connection would be to make reliability less contingent, not to make unreliability consistent with knowledge.

The concept of reliability is notoriously vague. If one believes *p* truly in a case α, in which other cases must one avoid false belief in order to count as reliable enough to know *p* in α? There is no obvious way to specify in independent terms which other cases are relevant. This is sometimes known as the *generality problem* for reliabilism. Some have argued that the generality problem is insoluble and that reliabilist theories in epistemology should therefore be abandoned (Conee and Feldman 1998). Let us concede for the sake of argument that the generality problem is indeed insoluble. It does not follow that appeals to reliability in epistemology should be abandoned. For the insolubility of the generality problem means that the concept of reliability cannot be defined in independent terms; it does not mean that the concept is incoherent. Most words express indefinable concepts; 'reliable' is not special in that respect. Irrespective of any relation to the concept *knows*, we clearly do have a workable concept *is reliable*; for example, historians sensibly ask which of their sources are reliable. The concept is certainly vague, but most words express vague concepts; 'reliable' is not special in that respect either. The concept *is reliable* need not be precise to be related to the concept *knows*; it need only be vague in ways that correspond to the vagueness in *knows*. No reason has emerged to doubt the intuitive claim that reliability is necessary for knowledge.

If one believes *p* truly in a case α, one must avoid false belief in other cases sufficiently similar to α in order to count as reliable enough to know *p* in α. The vagueness in 'sufficiently similar' matches the vagueness in 'reliable', and in 'know'. Since the account of knowledge developed in Chapter 1 implies that the reliability condition will not be a conjunct in a non-circular analysis of the concept *knows*, we need not even assume that we can specify the relevant degree and kind of similarity without using the concept *knows*. To suppose that reliability is necessary for knowledge is not to suppose that the concept *knows* can be analysed in terms of the concept *is reliable*, for it may be impossible to frame other necessary conditions without use of the concept *knows* whose conjunction with reliability is a necessary and sufficient condition for knowledge (see section 1.3).

We cannot expect always to apply a vague concept by appeal to rigorous rules. We need good judgement of
p. 101 particular cases. Indeed, even when we can appeal to rigorous rules, they only postpone the moment ↳ at which we must apply concepts in particular cases on the basis of good judgement. We cannot put it off indefinitely, on pain of never getting started. The argument for (1_i) appeals to such judgement. The intuitive

idea is that if one believes outright to some degree that a condition C obtains, when in fact it does, and at a very slightly later time one believes outright on a very similar basis to a very slightly lower degree that C obtains, when in fact it does not, then one's earlier belief is not reliable enough to constitute knowledge. The earlier case is sufficiently similar to the later case. One's earlier reliance on C has too much in common with one's later reliance on it. The use of the concept *is reliable* here is a way of drawing attention to an aspect of the case relevant to the application of the concept *knows*, just as one might use the concept *is reliable* in arguing that a machine ill serves its purpose. The aim is not to establish a universal generalization but to construct a counterexample to one, the luminosity principle (L). As with counterexamples to proposed analyses of concepts, we are not required to derive our judgement as to whether the concept applies in a particular case from general principles.

Within the limits just explained, we can nevertheless see how a reliability condition on knowledge is consonant with the role of knowledge in the causal explanation of action, as described in sections 2.4 and 3.4. Knowledge is superior to mere true belief because, being more robust in the face of new evidence, it better facilitates action at a temporal distance. Other things being equal, given rational sensitivity to new evidence, present knowledge makes future true belief more likely than mere present true belief does. This is especially clear when the future belief is in a different proposition, that is, when the future belief can differ in truth-value from the present belief.

Some hunters see a deer disappear behind a rock. They believe truly that it is behind the rock. To complete their kill, they must maintain a true belief about the location of the deer for several minutes. But since it is logically possible for the deer to be behind the rock at one moment and not at another, their present-tensed belief may be true at one moment and false at another. By standard criteria of individuation, a proposition cannot change its truth-value; the sentence ‘The deer is behind the rock’ expresses different propositions at different times. In present terminology, it is logically possible for the unchanging condition that the deer is behind the rock to obtain at one moment and not at another. If the hunters know that the deer is behind the rock, they have the kind of sensitivity to its location that makes them more likely to have future true beliefs about its location than they are if they merely believe truly that it is behind the rock. If we are to explain why
p. 102 they ↳ later succeeded in killing the deer, given the foregoing situation, then it is more relevant that they know that the deer is behind the rock than that they believe truly that it is behind the rock.

The role of knowledge in the explanation of action exploits a kind of reliability. If at time *t* on basis *b* one knows *p*, and at a time *t** close enough to *t* on a basis *b** close enough to *b* one believes a proposition *p** close enough to *p*, then *p** should be true. The argument of section 4.3 allows us to pick *t**, *b**, and *p** arbitrarily close to *t*, *b*, and *p* respectively. We can make the time interval between t_i and t_{i+1} as short as we like. Since the relevant beliefs are in the obtaining of the same condition at those times, they will be correspondingly close. Since, as noted above, the beliefs can also be assumed to change in basis only gradually, their bases too will be correspondingly close. A well-chosen example will verify (1_0), . . . , $(1_n 1)$ and thereby provide the required counterexample to (L). A reliability condition on knowledge facilitates the role that knowledge does in fact play in the causal explanation of action. The appeal to such a condition does not depend only on brute intuition; it fits the independently motivated conception of knowing as a mental state.

4.5 Sorites Arguments

An obvious doubt arises about the argument of section 4.3. The reasoning is very reminiscent of that in sorites paradoxes. If with 0 hairs on one's head one is bald, and, for every natural number *i*, with *i* hairs on one's head one is bald only if with *i* + 1 hairs on one's head one is bald, then for any natural number *n*, however large, it follows that with *n* hairs on one's head one is bald. The reasoning may therefore be suspected of concealing a mistake just like the concealed mistake in sorites reasoning, whatever that is. Does the argument illicitly exploit the vagueness of ‘feels cold’ or ‘know’?

The doubt can be made more specific. If the conclusion of the argument is false, then either not all the premises are true or the reasoning is invalid. Given (1_0), . . . , $(1_n$ 1) and the straightforwardly true (3_0) as auxiliary premises, the argument derives (2_0), . . . , $(2_n$ 1) from the supposed luminosity of the condition at issue and uncontested background assumptions, and then uses modus ponens to reach the straightforwardly false (3_n). By reductio ad absurdum, luminosity is rejected. On any reasonable view of vagueness, this reasoning shows that the luminosity claim is less than perfectly true, given that (1_0), . . . , $(1_n$ 1) are perfectly true.

p. 103 On some accounts, the rule of modus ponens fails to preserve less than perfect truth, because it sometimes leads from almost perfectly true premises to a conclusion that is not even almost perfectly true. But modus ponens should still preserve perfect truth. Within degree-theoretic semantics, a pseudo-conditional can be defined for which a conditional statement is perfectly true if and only if its consequent is at worst slightly less true than its antecedent (Peacocke 1981: 127). For present purposes, however, we can legitimately stipulate that the conditional to be used in the argument is of the more conventional kind for which the conditional statement is perfectly true if and only if the consequent is at least as true as the antecedent.

On other accounts, the rule of reductio ad absurdum is problematic because an assumption can have perfectly false consequences without itself being perfectly false, and therefore without having a perfectly true negation. Nevertheless, an assumption with perfectly false consequences is still less than perfectly true. Moreover, it is arguable that vagueness requires no revision of classical logic at all.[3]

For the purposes of this chapter, it would suffice to argue that the luminosity claim is less than perfectly true, for then it will have perfectly false consequences, which should discourage its application to philosophy. Thus the way for the defender of (perfect) luminosity to use the connection with sorites paradoxes is by arguing that not all of (1_0), . . . , $(1_n$ 1) are perfectly true, and using the vagueness of some relevant term to explain away their plausibility. Of course, the argument for (1_i) would remain to be addressed. Fortunately, however, the strategy can be tested more directly. For if (1_0), . . . , $(1_n$ 1) are in effect the premises of a sorites paradox, then sharpening the relevantly vague expressions should make at least one of them clearly false, just as sharpening the term ‘bald’ by stipulating a cut-off point gives the conditional ‘With *i* hairs on one's head one is bald only if with *i* + 1 hairs on one's head one is bald’ a clearly false instance. Does the same happen here?

The relevantly vague expressions in (1_i) are ‘feels cold’ and ‘knows’. We can sharpen ‘feels cold’ by using a physiological condition to resolve borderline cases. Let us assume that the subject of the process has no access to the technology needed to determine whether the physiological condition obtains, and so is not in a position to know whether it does. These stipulations in no way weaken the argument for (1_i). The
p. 104 considerations about reliability remain as cogent as before, for they were ↳ based on our limited powers of discrimination amongst our own sensations, not on the vagueness of ‘feels cold’. It might be objected that the sharpening violates the intended meaning of ‘feels cold’. However, that would not undermine the contrast between (1_i) and the major premise of a sorites paradox. For *any* complete sharpening of ‘bald’ yields a clearly false instance of the principle ‘With *i* hairs on one's head one is bald only if with *i* + 1 hairs on one's head one is bald’, even if it violates the intended meaning of ‘bald’ by, for example, falsifying the

converse downwards principle ‘With $i + 1$ hairs on one's head one is bald only if with i hairs on one's head one is bald’. By definition, the sharpened term applies wherever the unsharpened term clearly applied and fails to apply wherever the unsharpened term clearly failed to apply; thus, on any sharpening, ‘With 0 hairs on one's head one is bald’ is true and ‘With i hairs on one's head one is bald’ is false for a suitably large number n, so for some number i the conditional ‘With i hairs on one's head one is bald only if with $i + 1$ hairs on one's head one is bald’ is false. Thus even the truth of (1_0), . . . , $(1_n\ 1)$ on a sharpening of the vague terms that violates their intended meaning is enough to differentiate them from the premises of a sorites paradox.

The vague expression ‘knows’ remains. Sharpen it by tightening up its conditions of application: in the new sense it is not to apply in borderline cases for knowing in the old sense. It does not matter whether it applies in borderline cases of borderline cases for the old sense. If anything, this strengthens the argument for (1_i), by building more into its antecedent. It does not help one to know whether one feels cold. Indeed, one need not even be aware of the stipulation about ‘know’, for it is made by the theorist, not by the subject.

The stipulations will not make ‘feels cold’ and ‘knows’ perfectly precise; no feasible sharpening could do that. Fortunately, perfect precision is not necessary. We need only sharpen those expressions enough to resolve the finitely many borderline cases that actually arise in the argument. Such sharpening has the opposite effect to that predicted by the assimilation of the argument against luminosity to sorites reasoning; (1_i) becomes more not less plausible. The argument is not just another sorites paradox.

Nevertheless, the argument against luminosity might be thought to commit a subtler fallacy of vagueness. A defender of (2_i) might take the vagueness of its constituent terms to be essential to its truth, and explain the plausibility of (1_i) by assigning it a status short of perfect truth, while conceding that all of (1_0), . . . , (1_{n-1}) are true on some sharpenings, such as those considered above. The critic might take any sharpening that p. 105 falsifies (2_i) to violate the intended meanings of the vague terms, on ↳ the grounds that those meanings make (2_i) analytic. On such a view, some unsharpened (1_i) would be almost but not quite perfectly true, because its consequent would be almost but not quite as true as its antecedent. The reliability conditions adduced in favour of (1_i) would be treated as almost but not quite perfectly correct. No justification has been provided for not treating them as perfectly correct, but let that pass. For the concession is in any case inadequate. The defender of (2_i) must reject the following variation on (1_i):

(1Pi) If it is perfectly true that in α *i* one knows that one feels cold, then it is perfectly true that in α *i*+1 one feels cold.

For if (2_i) is perfectly true, then the perfect truth of its antecedent implies the perfect truth of its consequent:

(2Pi) If it is perfectly true that in α *i* one feels cold, then it is perfectly true that in α *i* one knows that one feels cold.

Statements $(1P_i)$ and $(2P_i)$ give an argument from the perfect truth of (3_i) to the perfect truth of (3_{i+1}), and therefore from the uncontested perfect truth of (3_0) to the perfect truth of (3_n); but the falsity of (3_n) is uncon-tested.

The critic will presumably treat $(1P_i)$ like (1_i), claiming that for some number *i*, it can be perfectly true that in α *i* one knows that one feels cold, but slightly less than perfectly true that in α *i*+1 one feels cold. Can there be such an *i*? If it is less than perfectly true that in α *i*+1 one feels cold, then there is a strict standard by which it is false in α *i*+1 that one feels cold; so, by that standard, in α *i*+1 one is fairly confident of what is false, that one feels cold. If so, it is less than perfectly true that in α *i* one knows that one feels cold, if the reliability considerations are to be assigned any positive weight at all. To put the argument more directly, if it is perfectly true that in α *i* one knows that one feels cold, then it is perfectly true that one achieves the level of reliability necessary for knowing, and therefore perfectly true that in α *i*+1 one feels cold. Thus the objection to $(1P_i)$ fails, and $(1P_0)$, . . . , $(1P_n 1)$ suffice for an argument that not all of (2_0), . . . , (2_{n-1}) are perfectly true. Invoking degrees of truth will not protect claims of perfect luminosity.

The point is reinforced by the observation that, once the luminosity assumption is dropped, (3_n) does not follow in classical logic from (1_0), . . . , $(1_n 1)$ and (3_0). To see this, pick j and k such that $0 \leq j < k < n$; for each i, evaluate 'One feels cold' as true in $\alpha\ i$ if and only if $i \leq k$, and otherwise as false; evaluate 'One knows that one feels cold' as true in $\alpha\ i$ if and only if $i \leq j$, and otherwise as false. On this evaluation, (1_i) is always true, for if
p. 106 the antecedent is true, then $i \leq j < k$, so $i + 1 \leq k$, so the ↳ consequent is true. Statement (3_0) is true because $0 < k$. Statement (3_n) is false because $k < n$. We can extend this evaluation in the manner of the standard semantics for modal logic by treating cases like possible worlds and 'One knows that . . . ' like 'It is necessary that . . . '. The foregoing evaluation results if one defines a case $\alpha\ h$ to be accessible from a case $\alpha\ i$ if and only if $|\ h\ i| \leq k\ j$, evaluates 'One knows that A' as true at a case $\alpha\ i$ if and only if 'A' is true at all cases accessible from $\alpha\ i$, and evaluates 'one feels cold' as before. Since a classical evaluation makes (1_0), . . . , $(1_n 1)$ and (3_0) true and (3_n) false, the latter does not follow from the former in classical logic. Contrast the sorites paradox: for any n, 'With n hairs on one's head one is bald' does follow in classical logic from 'With 0 hairs on one's head one is bald' and conditionals of the form 'With i hairs on one's head one is bald only if with $i + 1$ hairs on one's head one is bald'. Once luminosity is denied, conditionals of the form (1_i) generate no paradox.

Consistently with all this, we can postulate a more general phenomenon of which both vagueness and failures of luminosity independent of vagueness are special cases (Williamson 1994b and below). On such a view, the epistemological principles underlying (1_i) are important for vagueness too, but it does not follow that all their manifestations involve vagueness. Indeed, the epistemological principles by themselves imply no specific theory of vagueness.

4.6 Generalizations

Section 4.3 argued that a specimen condition—that one feels cold—is not luminous. How far does the argument generalize?

The argument assumed nothing specific about the condition of feeling cold. It extends to the examples of supposedly luminous conditions mentioned in section 4.2. Since pain sometimes gradually subsides, for example, an argument against the luminosity of the condition that one is in pain can be modelled on the argument against the luminosity of the condition that one feels cold, without any structural revisions. It is not perfectly true that whenever one is in pain, one is in a position to know that one is in pain. That one is in pain does not imply that one is in a position to know that one is in pain. Similarly, two synonyms can gradually diverge in meaning, as a mere difference in tone grows into a difference in application. The structure of the argument against luminosity is just as before. That two words have the same meaning for
p. 107 one does not imply that one is in a position to know that they have the same ↳ meaning for one. Equally, that they have different meanings for one does not imply that one is in a position to know that they have different meanings for one. The argument also applies to the condition that things appear to one in some way, for example, that it looks to one as though there is a purple patch ahead. Cases in which things appear to one in some way can gradually give way to cases in which they do not appear to one in that way. That they appear to one in that way does not imply that one is in a position to know that they appear to one in that way.

The condition that things appear to one in some way is often supposed to be a paradigm of what is called *response-dependence*. Unfortunately, that phrase is used in many senses, few of them clear. If the response-dependence of a condition means only that whether it obtains has *some* constitutive dependence on whether one is disposed to judge that it obtains, then response-dependence does not entail luminosity, although non-luminosity does constrain what forms of dependence a condition can exhibit (see Williamson 1994b: 180–4 for the case of colour). But if ‘response-dependent’ is so defined that a response-dependent condition must be luminous, then the conditions that are standardly taken as paradigms of response-dependence are none of them response-dependent.

Further applications of the argument involve conditions on one's knowledge. Since one can gain or lose knowledge gradually, we can use the argument to show that, for most propositions *p*, neither the condition that one knows *p* nor the condition that one does not know *p* is luminous. One can know *p* without being in a position to know that one knows *p*, and one can fail to know *p* without being in a position to know that one fails to know *p*. Chapters 5 and 8 respectively discuss these applications in more detail.

On what general features of a condition does the argument against luminosity depend? As it stands, it requires the condition to obtain in some cases and not in others. Thus it is ineffective against a condition that obtains in all cases or in none. Given a sufficiently restrictive understanding of what a case is, that might include the Cartesian condition that one exists, or even that one thinks. It does not include the condition that one is thinking about one's existence, for one does that in some cases and not in others on any reasonable understanding of what a case is.

A condition that obtains in no case, the impossible condition, is automatically luminous; (L) holds vacuously. Is a condition that obtains in every case, the necessary condition, luminous too? It is luminous as
p. 108 presented in a simple tautological guise, if cases are restricted to those in ↳ which the subject has the concepts to formulate the tautology. It is not luminous as presented in the guise of an a posteriori necessity, or an unproved mathematical truth, or if the cases include some in which one lacks appropriate concepts.

The argument also requires the possibility of a change from cases in which the condition obtains to cases in which it does not. Thus it would not be effective against an eternal condition, which always obtains if it ever obtains: for example, the condition that one felt cold at midnight on New Year's Eve 1999. However, many eternal conditions, including that one, permit a change from cases in which one is in a position to know that they obtain to cases in which one is not in a position to know that they obtain. Such a condition cannot be luminous, for since it obtains in the earlier cases in which one is in a position to know that it obtains (because being in a position to know is factive), it also obtains in the later cases in which one is not in a position to know that it obtains (because the condition is eternal). Thus an eternal condition is luminous only if one cannot change from being in a position to know that it obtains to not being in such a position. There are candidates for such conditions. For example, if a subject S is always in a position to know that she is S—which is not to say that she must know her own name—then anyone who is ever in a position to know that the condition that one is S obtains is always in a position to know that it obtains, because the only such person is S herself. Perhaps the argument could be extended to show that not even this condition is luminous, by consideration of a science-fiction process in which someone else is gradually replaced by S. However, no such extension will be attempted here. Such examples do not seriously threaten the idea that only trivial conditions are luminous.

The argument also assumes that one is considering the relevant condition under the relevant guise throughout the process. Consequently, it does not apply to some conditions on one's considerations. For example, let C be the condition that one is entertaining the proposition that it is raining, and let G be the guise under which C has just been presented here. To consider C under G is to consider as such the condition that one is entertaining the proposition that it is raining; in so doing, one thereby entertains the proposition that it is raining, so C obtains. Thus one cannot gradually pass from cases in which C obtains to cases in which C does not obtain while considering C under G throughout the process. Although one can gradually pass from cases in which C obtains to cases in which C does not obtain, one does not consider C under G in the late stages of the process. For all the argument shows, C is luminous: if one is entertaining the
p. 109 proposition that it is raining, then one is ↳ in a position to know that one is entertaining the proposition that it is raining. When one is entertaining a slightly different proposition *p*, one does not have a high degree

of false belief that one is entertaining the proposition that it is raining; one has a high degree of true belief that one is entertaining *p*, since the belief derives its content from *p* itself. Thus the argument does not apply to examples in which one considers the condition only when it obtains. Such examples constitute a very minor limitation on the generality of the argument. In any case, we may conjecture that, for any condition C, if one can move gradually to cases in which C obtains from cases in which C does not obtain, while considering C throughout, then C is not luminous. The conjecture is discussed further in section 5.2.

Luminous conditions are curiosities. Far from forming a cognitive home, they are remote from our ordinary interests. The conditions with which we engage in our everyday life are, from the start, non-luminous.

4.7 Scientific Tests

To be physically and psychologically capable of knowing *p* is not sufficient, even given *p*, for being in a position to know *p*; one may be in the wrong place. Thus it does not follow from the non-luminosity of a condition that there are cases in which, although it obtains, one is not physically and psychologically capable of knowing that it obtains. Nevertheless, it is natural to ask, if one is not in a position to know in a case α that one then feels cold, how is one to know in some other case β that in α one feels cold? Must or can there be such a case β? Analogous questions arise about other non-luminous conditions. The argument of section 4.3 leaves them open. It is consistent with, but does not entail, the possibility of a physiological technique by which one could subsequently discover that one had been feeling cold in α.

The hypothetical technique faces difficulties. Suppose that feelings of cold and hot are found generally to be correlated with a measurable physiological variable V. We must discover which values of V are associated with the condition that one feels cold. They include the values associated with the condition that one is in a position to know that one feels cold. But if they included only those values, the condition that one feels cold would be luminous, which it is not. We are not in a position to know which further values of V are associated with that condition. Our problem is that we cannot calibrate the physiological measurement of feeling cold. Even if measurements of V were perfectly precise—which ↳ they will not be—they would not answer the original question. Attempts to measure other ordinary conditions face similar problems. Their non-luminosity prevents us from perfectly calibrating instruments to detect whether they obtain.

It might still be held to be metaphysically possible to find out whether one feels cold by the testimony of a literal or metaphorical *deus ex machina*. But it certainly cannot be assumed without argument that if an ordinary condition obtains in a case α, then in some possible case it is known that the condition obtains in α. Section 12.5 discusses the issue further.

4.8 Assertibility Conditions

The failure of luminosity impinges on Michael Dummett's arguments for an anti-realist theory of meaning, which explains meanings in terms of the conditions under which speakers are warranted in using sentences assertively, by contrast with a realist theory of meaning, which explains meanings in terms of the conditions under which sentences express truths. Of course, Dummett's anti-realist does not make the extreme claim that every condition is luminous. All parties can accept that stone age men lived when the moon caused the tides, although they were not in a position to know that the moon caused the tides. The connection between luminosity and anti-realism is a subtler one.

Dummett objects to the realist's truth-conditional theory of meaning that it violates a necessary connection between meaning and use. To understand a sentence is to know what it means. If, as Dummett's realist holds, meanings are truth-conditions, then speakers of a language know the truth-conditions of its sentences.[4] Knowing the truth-condition of a sentence *s* cannot consist merely in being disposed to say something of the form '*s* is true if and only if P'; one must also understand the biconditional, and an infinite regress looms. In the basic case, one's knowledge of the truth-condition must be implicit. If one could always ↳ recognize whether it obtained, then knowledge of the truth-condition of *s* might consist in a willingness to assert *s* just when the truth-condition obtained. Dishonesty, shyness, and other complications are assumed to have been somehow filtered out. However, the realist insists that the truth-conditions of some sentences obtain even though no speaker of the language can recognize that they obtain. Dummett argues that the realist has no substantial explanation of what knowing that sentences have those truth-conditions consists in. The proposed remedy is that the meaning of a sentence should be given by its assertibility-condition rather than by its truth-condition. Thus to understand a sentence is to know its assertibility-condition, and this knowledge can consist in a willingness to assert the sentence just when the assertibility-condition obtains.[5]

p. 111

The remedy fails if the objection to truth-conditional theories of meaning applies equally to assertibility-conditional theories of meaning. Thus Dummett's argument requires that when an assertibility-condition obtains, competent speakers of the language can recognize that it obtains. He acknowledges that requirement: 'The conditions under which a sentence is recognized as true or false . . . have, by the nature of the case, to be conditions which we can recognize as obtaining when they obtain' (1981: 586; compare 1991: 317 and 1993: 45–6). That is, when a recognition-condition obtains, we can recognize that the recognition-condition obtains. Dummett evidently intends the recognition-condition for the truth of a sentence to be its assertibility-condition, which yields the thesis that when an assertibility-condition obtains, we can recognize that it obtains. But recognizing is coming to know, and Dummett's 'can' may be glossed as 'is in a position to'. Thus Dummett requires assertibility-conditions to be luminous.

The argument against luminosity in section 4.3 generalizes to assertibility-conditions. For example, it can gradually cease to be assertible that it is raining. By the argument, that it is assertible that it is raining does not imply that one is in a position to know that it is assertible that it is raining. Even in the mathematical case, in which Dummett uses the proof-based intuitionistic semantics as a paradigm of an assertibility-conditional theory of meaning, proofs can be understood or forgotten gradually.[6] By the argument, that one has a proof of a mathematical assertion does not imply that one is in a position to know that one has a ↳ proof of it. Thus assertibility-conditions have the very feature that is supposed to lay truth-conditions open to Dummett's attack.[7]

p. 112

An assertibility-conditional theory of meaning is likely to distinguish between canonical and non-canonical warrants for assertion, for example, between having a proof and having been told by a reliable informant that there is one. The recursive semantics will be formulated in terms of canonical warrants; a non-canonical warrant will be explained as an entitlement to believe that there is a canonical warrant. The argument applies whether or not 'warrant' is qualified by 'canonical'.

The anti-realist might reply that one's understanding of a sentence can consist in the fact that one is willing to assert it when and only when its assertibility-condition obtains, even if one does not know that it obtains. This reply concedes that assertibility-conditions fail Dummett's luminosity constraint; but then something is wrong with his argument for assertibility-conditional theories of meaning, which treats that constraint as binding.

A different reply is that if Dummett intends recognizability to be assertibility, then what he requires is only that when *p* is assertible, it is assertible that *p* is assertible. If misleading evidence sometimes warrants false assertions, then it might be assertible that *p* is assertible even when one is not in a position to know that *p* is assertible, so Dummett would not require assertibility-conditions to be luminous. This reply fails because the argument of section 4.3 can be generalized to an argument that no non-trivial condition obtains only when it is assertible that it obtains.[8]

p. 113 Dummett presents his argument as a challenge to the realist to explain what knowledge of realist truth-conditions consists in. He does not claim to prove that the realist cannot meet the challenge, although he denies that it has been met so far. He allows that it might be met in some areas and not in others. But he assumes that the anti-realist can easily meet the corresponding challenge, to explain what knowledge of assertibility-conditions consists in. If the foregoing argument is correct, that assumption is false; the anti-realist faces the same sort of difficulty as the realist does. The contrast between truth-conditions and assertibility-conditions is off the point.

Both truth-conditional and assertibility-conditional theories of meaning find it hard to meet Dummett's challenge because both truth-conditions and assertibility-conditions are non-luminous. They share this feature with every other kind of non-trivial condition that might be offered as the meaning of a sentence. Since trivial conditions are not serious candidates for the meanings of most sentences, a serious X-conditional theory of meaning will find it hard to meet Dummett's challenge, for any X. If any systematic theory of meaning can be cast as X-conditional for some X, then any systematic theory of meaning will find it hard to meet Dummett's challenge. If 'hard' turns out to be 'impossible', then failure to meet the challenge eliminates truth-conditional theories of meaning only if it eliminates all systematic theories of meaning. The challenge embodies extreme demands on a theory of meaning. We should not assume the possibility of a reductive explanation of what knowledge of meaning 'consists in' of the kind that Dummett demands.

On an anti-realist picture, thought initially engages with conditions whose *esse* is their *percipi*; if it later finds its laborious way to conditions of greater depth, it must do so from the starting point of that cognitive home. Assertibility-conditions are pictured as forming a cognitive home in language. They do not. Thought engages with conditions whose *esse* is distinct from their *percipi* as soon as it engages with any conditions at all; even perception does. Trivialities aside, there is nothing else to engage with. We have no cognitive home.

Notes

1 See Wittgenstein 1958: §126. Of course, Wittgenstein's 'we', unlike Descartes's, is collective. What is not hidden from each Cartesian subject is only its own thinking, not that of other Cartesian subjects. Wittgenstein is speaking of what is not hidden from any of us. The arguments below have not been, but could be, adjusted to the rejection in Wittgenstein 1969 of claims to know *p* when there is no question of doubting *p*; conversational inappropriateness is compatible with truth. When I feel cold with no question of doubt and I know that everyone else in the room feels cold, I usually know that everyone in the room feels cold. If so, I know that I feel cold.

2 See also Dummett 1981: 632 and 1993: 4. For a recent discussion see Boghossian 1994.

3 See Williamson 1994b on logic for vague languages. The arguments of the present chapter do not depend on the epistemic account of vagueness developed there, which requires no revision of classical logic.

4 Provision must naturally be made for non-declarative sentences, perhaps by a distinction between sense and force. Even for declaratives, knowledge of truth-conditions may not suffice for understanding. More generally, knowledge of meaning may not suffice for understanding. A reliable informant tells me that a sentence of Hungarian written on the board means that the cat sat on the mat. I know by testimony that the sentence means that the cat sat on the mat, but arguably I do not understand it because I do not know what each constituent word contributes to that meaning. Such complications will be ignored in what follows.

5 Dummett says that 'no difficulty can any longer arise over what such knowledge consists in' (1977: 375). Note that assertibility-conditions are not truth-conditions even on Dummett's anti-realist conception of truth (see n. 7 below).

6 On recognizing intuitionistic proofs as proofs see Weinstein 1983 and Pagin 1994.

7 In criticizing truth-conditional theories of meaning, Dummett often focusses on the *undecidability* of truth, that is, on speakers' lack of an effective procedure for coming to know whether a given truth-condition obtains. When the argument takes this form, it threatens assertibility-conditional theories of meaning too unless assertibility is decidable, that is, unless speakers always have an effective procedure for coming to know whether a given assertibility-condition obtains (Crispin Wright, 1992b: 56, endorses the decidability of assertibility). The decidability of assertibility does not entail the

decidability of truth, even on most anti-realist conceptions of truth, for the latter do not identify truth with assertibility in non-ideal cases. For example, intuitionists identify assertibility with possession of a proof, and truth with the existence of one (that is, with the possibility of possessing it); they deny the law of excluded middle just because they take truth to be undecidable. Even if assertibility were luminous, its decidability would follow only on the additional assumption that unassertibility is luminous too. Thus what the anti-realist requires of assertibility depends on whether what realist truth is accused of crucially lacking is luminosity or decidability, even though, given bivalence and the presence of classical negation in the object language—both assumptions acceptable to the realist—truth is decidable if and only if it is luminous (on an appropriate construal of 'in a position to know'). But since any decidable condition is luminous (on that construal), the points in the text apply to both forms of anti-realist argument.

8 A conditional of the form 'If in α^{i} it is assertible that C obtains, then in α^{i+1} C obtains' replaces (1_i) (Williamson 1995a). Chapter 11 defends a view on which the move in question makes no difference, because only knowing *p* warrants asserting *p*. Crispin Wright 1992b: 18 seems to assume the S4 principle for assertibility in arguing that truth and assertibility coincide in 'positive normative force'.