KNOWLEDGE
and its
LIMITS

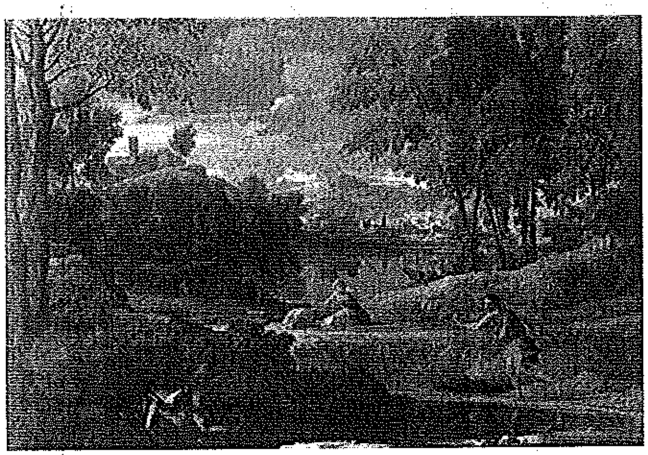

Timothy Williamson

OXFORD

5
Margins and Iterations

5.1 KNOWING THAT ONE KNOWS

One can know something without being in a position to know that one knows it. We reached that conclusion using the form of argument developed in the previous chapter, for by a gradual process one can gain or lose knowledge. Similarly, one can know that one knows something without being in a position to know that one knows that one knows it, for by a gradual process one can gain or lose knowledge that one knows. This chapter explores such limits to our ability to iterate knowledge. They stem from our need of margins for error in much of our knowledge. Those limits make problems for common knowledge, in which everyone knows that everyone knows that everyone knows that . . .. Chapter 6 will apply the results to suggest a diagnosis of the paradox of the Surprise Examination and related puzzles.

We first consider in some detail a variant argument against the luminosity of the condition that one knows something. One can know without being in a position to know that one knows.

Looking out of his window, Mr Magoo can see a tree some distance off. He wonders how tall it is. Evidently, he cannot tell to the nearest inch just by looking. His eyesight and ability to judge heights are nothing like that good. Since he has no other source of relevant information at the time, he does not know how tall the tree is to the nearest inch. For no natural number i does he know that the tree is i inches tall, that is, more than $i-0.5$ and not more than $i+0.5$ inches tall. Nevertheless, by looking he has gained some knowledge. He knows that the tree is not 60 or 6,000 inches tall. In fact, the tree is 666 inches tall, but he does not know that. For all he knows, it is 665 or 667 inches tall. For many natural numbers i, he does not know that the tree is not i inches tall. More precisely, for many natural numbers i, he does not know the proposition expressed by the result of replacing 'i' in 'The tree is not i inches tall' by a numeral designating i. We are not concerned with knowledge of propositions expressed by sentences in which i is designated by a definite description, such as 'the height of the tree in inches', for he may not know which number fits the description.

To know that the tree is *i* inches tall, Mr Magoo would have to judge that it is *i* inches tall; but even if he so judges and in fact it is *i* inches tall, he is merely guessing; for all he knows it is really $i-1$ or $i+1$ inches tall. He does not know that it is not. Equally, if the tree is $i-1$ or $i+1$ inches tall, he does not know that it is not *i* inches tall. Anyone who can tell by looking that the tree is not *i* inches tall, when in fact it is $i+1$ inches tall, has much better eyesight and a much greater ability to judge heights than Mr Magoo has. These reflections do not depend on the value of *i*. For *no* natural number *i* is the tree $i+1$ inches tall while he knows that it is not *i* inches tall. In this story, Mr Magoo reflects on the limitations of his eyesight and ability to judge heights. Mr Magoo knows the facts just stated. Consequently, for each relevant natural number *i*:

(1_i) Mr Magoo knows that if the tree is $i+1$ inches tall, then he does not know that the tree is not *i* inches tall.

We could make the case for (1_i) even stronger by reducing the interval of an inch to something much smaller, perhaps a millionth of an inch, but that should not be necessary. To make the conditional 'If the tree is $i+1$ inches tall, then he does not know that it is not *i* inches tall' as uncontentious as possible, we can read 'if' as the truth-functional conditional, the weakest of all conditionals. In effect, it merely denies the conjunction 'The tree is $i+1$ inches tall and he knows that it is not *i* inches tall'.

Suppose, for a reductio ad absurdum, that the condition that one knows a proposition is luminous: if one knows it, then one is in a position to know that one knows it. We may also assume that, in the case at hand, for each proposition *p* pertinent to the argument, Mr Magoo has considered whether he knows *p*. Consequently, if he is in a position to know that he knows *p*, he does know that he knows *p*. Thus:

(KK) For any pertinent proposition *p*, if Mr Magoo knows *p* then he knows that he knows *p*.

Statement (KK) is a special case of the general 'KK' principle that if one knows something then one knows that one knows it, but sufficiently restricted to avoid many of the objections to the latter (for some of which see Sorensen 1988: 242). For example, (KK) does not imply by iteration that if *p* is pertinent then Mr Magoo has every finite number of iterations of knowledge of *p*, for it has not been granted that if *p* is pertinent then so too is the proposition that he knows *p*. The pertinent propositions are just those that occur in the argument below, which form a strictly limited set. Statement (KK) is also immune to the objection that a simple creature without the concept *knows* might still know, but would not know that it knew, for Mr Magoo has the concept *knows*.

We may legitimately assume that in the example Mr Magoo has been reflecting on the height of the tree and his knowledge of it so carefully that he has drawn all the pertinent conclusions about its height that follow deductively from what he knows; he has thereby come to know those conclusions. Let us consider a time at which that process is complete. We can therefore assume:

(C) If p and all members of the set X are pertinent propositions, p is a logical consequence of X, and Mr Magoo knows each member of X, then he knows p.

Of course, (C) is not justified by some general closure principle about knowledge. We often fail to know consequences of what we know, because we do not know that they are consequences. Statement (C) is simply a description of Mr Magoo's state once he has attained reflective equilibrium over the propositions at issue, by completing his deductions. Since Mr Magoo's deductive capacities do not fully enable him to overcome the limitations of his eyesight and ability to judge heights, and he knows that they do not, (1_i) remains true for all i.

By (KK), we can infer (3_i) from (2_i):

(2_i) Mr Magoo knows that the tree is not i inches tall.

(3_i) Mr Magoo knows that he knows that the tree is not i inches tall.

Now, let q be the proposition that the tree is i+1 inches tall. By (1_i), Mr Magoo knows $q \supset \sim(2_i)$; by (3_i), he knows (2_i). Now, $\sim q$ is a logical consequence of $q \supset \sim(2_i)$ and (2_i). Consequently, by (C), (1_i) and (3_i) imply that Mr Magoo knows $\sim q$:

(2_{i+1}) Mr Magoo knows that the tree is not i+1 inches tall.

Consequently, from (KK), (C) and (2_i) we can infer (2_{i+1}). By repeating the argument for values of i from 0 to 665, starting from (2_0) we reach the conclusion (2_{666}):

(2_0) Mr Magoo knows that the tree is not 0 inches tall.

(2_{666}) Mr Magoo knows that the tree is not 666 inches tall.

Statement (2_{666}) is false, for the tree is 666 inches tall and knowledge is factive. Thus, given the premises (1_0), . . ., (1_{665}), (2_0), (C), and (KK), we can deduce the false conclusion (2_{666}). Therefore, at least one of (1_0), . . ., (1_{665}), (2_0), (C), and (KK) is to be rejected. Premise (1_i) has already been defended for all i, and (2_0) is obviously true. Consequently, either (C) or (KK) is to be rejected.

Could we reject the assumption (C) that Mr Magoo's knowledge of the pertinent propositions is deductively closed? Assumption (C) is true if deduction is a way of extending one's knowledge: that is, if knowing $p_1, \ldots, p_n$, competently deducing q, and thereby coming to believe q is in general a way of coming to know q. Call that principle *intuitive closure*. Since by hypothesis Mr Magoo satisfies the conditions for the intuitive closure principle to apply, rejecting (C) is tantamount to rejecting intuitive closure. Robert Nozick's counterfactual analysis of knowledge is famously inconsistent with intuitive closure, but that is usually taken as a reason for rejecting the analysis, not for rejecting closure. Chapter 7 will provide arguments against counterfactual conditions on knowledge even of quite a weak kind; a fortiori they are arguments against Nozick's analysis.

A different objection occasionally made to intuitive closure is that even if one's premises are individually probable enough to count as known, one's conclusion might not be. For a logical consequence of several propositions may be less probable than each of them. If there are a million tickets in the lottery and only one wins, each proposition of the form 'Ticket i does not win' has a probability of 0.999999, yet the conjunction of all those propositions has a probability of 0. But that objection misconceives the relation between probability and knowledge; however unlikely one's ticket was to win the lottery, one did not know that it would not win, even if it did not (see also section 11.2). No probability short of 1 turns true belief into knowledge. Chapter 10 provides a very different understanding of the connection between knowledge and probability; it does not threaten intuitive closure.

The appeal to probability is in any case unavailing, for the argument can be reworked so that (C) is applied only to single-premise inferences; if q is a logical consequence of p then q is at least as probable as p. For the considerations that supported (1_i) also support:

(4o) Mr Magoo knows that (for all natural numbers m (if the tree is m+1 inches tall then he does not know that it is not m inches tall) and (the tree is not 0 inches tall)).

Parentheses have been inserted to clarify scope. Now suppose, for some given i:

(4_i) Mr Magoo knows that (for all natural numbers m (if the tree is m+1 inches tall then he does not know that it is not m inches tall) and (the tree is not i inches tall)).

By (KK) we have:

(5_i) Mr Magoo knows that he knows that (for all natural numbers *m* (if the tree is *m*+1 inches tall then he does not know that it is not *m* inches tall) and (the tree is not *i* inches tall)).

But Mr Magoo knows with certainty that if he knows a conjunction then the first conjunct is true and he knows the second. Thus:

(6_i) Mr Magoo knows that (for all natural numbers *m* (if the tree is *m*+1 inches tall then he does not know that it is not *m* inches tall) and (he knows that the tree is not *i* inches tall)).

But (C) for single-premise deductions applied to (6_i) gives:

(4_{i+1}) Mr Magoo knows that (for all natural numbers *m* (if the tree is *m*+1 inches tall then he does not know that it is not *m* inches tall) and (the tree is not *i*+1 inches tall)).

The inference from (4_i) to (4_{i+1}) is the required sorites step. If we iterate it for each *i* from 0 to 665, starting with (4_0), we reach:

(4_{666}) Mr Magoo knows that (for all natural numbers *m* (if the tree is *m*+1 inches tall then he does not know that it is not *m* inches tall) and (the tree is not 666 inches tall)).

Statement (4_{666}) is false, for the tree is 666 inches tall. Thus the problem does not depend on applying (C) to deductions with more than one premise.

We should in any case be very reluctant to reject intuitive closure, for it *is* intuitive. If we reject it, in what circumstances can we gain knowledge by deduction? Moreover, the closely related anti-luminosity argument in section 4.3 did not assume closure in any form, which suggests that it is not the crucial premise.

A different objection to the argument is that vagueness is somehow to blame. Section 4.5 discussed the same objection. Since the reasons for dismissing it are the same as before, they will not be repeated in detail here. The crucial point is that the premises of the argument are not justified by vagueness in 'know' but by limits on Mr Magoo's eyesight and his knowledge of them. In checking that (1_i) remains true when 'know' is sharpened, we must be careful because 'know' occurs twice in (1_i), which ascribes to Mr Magoo knowledge that he could express in the words 'If the tree is *i*+1 inches tall, then I do not know that the tree is not *i* inches tall'. But if we sharpen 'know' by stipulating a high standard for its application, we make that conditional harder to falsify and therefore easier to know, because the only occurrence of 'know' in the sentence is negative. Since (1_i) was clearly true prior to the sharpening, it therefore remains true afterwards; we may legitimately assume that Mr

Magoo has considered the sharpened sense of 'know'. That will not improve his eyesight. The argument does not rely on the vagueness of 'know'.

Given (C) and (KK) as auxiliary premises, there is a valid argument with otherwise true premises and a false conclusion. Premise (C) is accepted. Therefore, (KK) is to be rejected. Mr Magoo knows something pertinent without knowing that he knows it. Since (KK) follows from the assumption that the condition that one knows a proposition is luminous and background assumptions about Mr Magoo, the luminosity assumption is false. As in section 4.5, we can check that rejecting luminosity really does meet the difficulty by constructing a formal model of (C), $(1_0), \ldots, (1_i), \ldots, (2_0)$ and the negation of (2_{666}) (Appendix 2 has more details).

Mr Magoo cannot identify the particular proposition for which (KK) fails. In general, one cannot knowingly identify a particular counterexample to the KK principle in the first person present tense. If I know that I both know *p* and do not know that I know *p*, I must know the first conjunct of that conjunction (since knowing a conjunction entails knowing its conjuncts), that is, I must know that I know *p*, so the second conjunct is false, so I do not know the conjunction after all (since knowledge is factive); Chapter 12 discusses this kind of argument in more depth. The point may help to explain the seductiveness of the KK principle.

The crucial features of the example are common to virtually all perceptual knowledge. Thus the argument generalizes to show that our knowledge is pervaded by failures of the KK principle. To the informed observer, hearing gives some knowledge about loudness in decibels, and touch about heat in degrees centigrade. When I smell the milk I have some knowledge of the number of minutes since it was opened; when I taste the tea I have some knowledge of how many grains of sugar were put in. The point generalizes to knowledge from sources beyond present perception, such as memory and testimony. This is partly because they pass on inexact knowledge originally derived from past perception, partly because they add further ignorance themselves. How long was my last walk in steps? How long was someone else's walk, described to me as 'quite long'? In each case the possible answers lie on a scale, which can be divided so finely that if a given answer is in fact correct, then one does not know that its neighbouring answers are not correct, and one can know that one's powers of discrimination have that limit. The argument then proceeds as in the case of the distant tree.[1]

[1] The argument of section 5.1 is similar in form to the argument used by Nathan Salmon (1982: 238–40, 1986, 1989) against the S4 principle $\Box p \supset \Box\Box p$ for metaphysical

5.2 FURTHER ITERATIONS

We can generalize the argument of section 5.1 to further iterations of knowledge. We define them inductively. One knows^0 *p* if and only if *p* is true. For any natural number *k*, one knows^{k+1} *p* if and only if one knows^k that one knows *p*. To know^1 *p* is to know *p*, to know^2 *p* is to know that one knows *p*, and so on.

For any *k*, we can argue in parallel with section 5.1 that one can know^k something without being in a position to know that one knows^k it. For if we make suitably modified assumptions about the height and distance of the tree, Mr Magoo's eyesight, his knowledge of its limitations, and his powers of reflection, we can construct a situation in which these modified assumptions are true for a given *k* and all *i*:

($1_i{}^k$) Mr Magoo knows^k that if the tree is *i*+1 inches tall, then he does not know that the tree is not *i* inches tall.

($2_0{}^k$) Mr Magoo knows^k that the tree is not 0 inches tall.

(C^k) If *p* and all members of the set X are pertinent propositions, *p* is a logical consequence of X, and Mr Magoo knows^k each member of X, then he knows^k *p*.

Now make these two assumptions, for a given number *i*:

($2_i{}^k$) Mr Magoo knows^k that the tree is not *i* inches tall.

(KK^k) For any proposition *p*, if Mr Magoo knows^k *p* then he knows^{k+1} *p*.

Since knowing^{k+1} is equivalent to knowing^k that one knows, ($2_i{}^k$) and (KK^k) entail:

($3_i{}^k$) Mr Magoo knows^k that he knows that the tree is not *i* inches tall.

Assumptions ($1_i{}^k$), ($3_i{}^k$), and (C^k) entail:

necessity. The form is valid; the question in each case is whether the premises are plausible. Williamson 1990a: 129 suggests that the plausibility of one of Salmon's premises comes from a source that generates sorites paradoxes; Salmon 1993 disagrees. It is hard to adjudicate disputes about whether intuitions have a common source. The structure of Salmon's premise in itself does not commit him to a sorites paradox. On the other hand, if the validity of the S4 principle is built into our conception of unrestricted metaphysical possibility and necessity, those who find the major premises of sorites paradoxes plausible would also be likely to find Salmon's premise plausible. The analogue of Salmon's argument may be sound for the restricted notions of possibility and necessity discussed in section 5.3, which might also help to account for the plausible appearance of the premises in Salmon's original version.

($2_{i+1}{}^k$) Mr Magoo knows^k that the tree is not i+1 inches tall.

Suppose that the tree is in fact n inches high. By repeated application of the argument from ($2_i{}^k$) to ($2_{i+1}{}^k$), starting with $2(_0{}^k)$, we reach:

($2_n{}^k$) Mr Magoo knows^k that the tree is not n inches tall.

Since knowledge^k is as factive as knowledge, ($2_n{}^k$) is false. It was deduced from the assumptions ($1_0{}^k$), . . ., ($1_{n-1}{}^k$), ($2_0{}^k$), (C^k), and (KK^k). By construction of the example, ($1_0{}^k$), . . ., ($1_{n-1}{}^k$), ($2_0{}^k$), and (C^k) are true; therefore (KK^k) is false. The replies to objections to the argument follow the pattern of section 5.1. Thus one can know^k something without being in a position to know^{k+1} it. In other words, one can know^k something without being in a position to know that one knows^k it.

By contrast, some other objections to the general KK thesis do not threaten the corresponding generalization of (KK^k) for k>1. For example, a simple creature might know that it was snowing without knowing that it knows that it was snowing because the latter, unlike the former, requires it to have a concept of knowledge, which it lacks. But if k≥2 and one knows^k p, then one knows something concerning knowledge and so has the concepts needed for knowing^{k+1} p.

Can we combine all finite iterations of knowledge? One knows$^\omega$ p if and only if for every natural number k one knows^k p. Can we mimic the foregoing argument with ω in place of k? The premises of the reductio ad absurdum are these:

($1_i{}^\omega$) Mr Magoo knows$^\omega$ that if the tree is i+1 inches tall, then he does not know that the tree is not i inches tall.

($2_0{}^\omega$) Mr Magoo knows$^\omega$ that the tree is not 0 inches tall.

(C^ω) If p and all members of the set X are pertinent propositions, p is a logical consequence of X, and Mr Magoo knows$^\omega$ each member of X, then he knows$^\omega$ p.

(KK^ω) For any proposition p, if Mr Magoo knows$^\omega$ p then he knows$^\omega$ that he knows p.

For some n, the false conclusion is this:

($2_n{}^\omega$) Mr Magoo knows$^\omega$ that the tree is not n inches tall.

We might conclude on the basis of ($1_0{}^\omega$), . . ., ($1_{n-1}{}^\omega$), ($2_0{}^\omega$), and (C^ω) that Mr Magoo is a counterexample to (KK^ω). But that is the wrong moral to draw from this example, for (KK^ω) is a logical truth. If Mr Magoo knows$^\omega$ p, then for each natural number k he knows^{k+1} p, which is to know^k that he knows p, so he knows$^\omega$ that he knows p. Thus ($1_0{}^\omega$), . . .,

$(1_{n-1}{}^{\omega})$, $(2_0{}^{\omega})$, and (C^{ω}) entail the false conclusion $(2_n{}^{\omega})$ by themselves; one of them is false. Given a natural number k, we can construct an example in which $(1_0{}^k)$, . . ., $(1_{n-1}{}^k)$, and $(2_0{}^k)$ are true, by finite adjustments of the original case, which are clearly possible. An infinite adjustment turns out to be impossible. That does not undermine the morals drawn from the earlier versions of the argument. The crude point is that iterating knowledge is hard, and each iteration adds a layer of difficulty. Knowledge$^{\omega}$ involves infinitely many layers of difficulty. Under some conditions, that amounts to impossibility. The next section develops these remarks more systematically.

Knowledge$^{\omega}$ presents an interesting challenge to the generalized argument against luminosity in Chapter 3. Since it seems possible in principle to gain or lose knowledge$^{\omega}$, one might expect the argument to show that one can know$^{\omega}$ without being in a position to know that one knows$^{\omega}$. But that conclusion is problematic. For if one knows$^{\omega}$ p, then one knows each member of the set containing the proposition that one knows^k p for each natural number k; thus one knows the premises of a deductively valid argument to the conclusion that one knows$^{\omega}$ p; one is therefore in some sense in a position to know that one knows$^{\omega}$ p. The condition that one knows$^{\omega}$ p seems to be luminous.

The argument might be challenged on the grounds that we are not in a position to make inferences with infinitely many premises. Indeed, even when an inference has only finitely many premises, it is not obvious that we are always in a position to know that which follows deductively from what we know. Only in a rather attenuated sense are we in a position to know all the consequences of the axioms of Peano Arithmetic. However, this response is not wholly satisfying, for the original argument against luminosity made no appeal to limits on powers of inference. If the condition that one knows$^{\omega}$ p is luminous in the attenuated sense, why does the original argument not generalize to this case?

Knowing$^{\omega}$ may fail the gradualness requirement. Although someone can gain or lose knowledge$^{\omega}$, the change may necessarily be sudden. After all, it is the change from finitely many iterations of knowledge to infinitely many or vice versa; how could it be gradual? If knowing$^{\omega}$ does fail the gradualness requirement, it will be a hard state to enter or leave: how is one to jump instantaneously from the finite to the infinite or back again? The kind of common knowledge that we are supposed to have of conventions is usually defined in a way that requires us to know$^{\omega}$. For example, if John knows that Jane knows that John knows that Jane knows that John knows p, then John knows that John knows that John knows p, if he is sufficiently reflective. Common knowledge would therefore be a convenient idealization, like a frictionless plane.

The convenience need not be confined to the theoretician. Perhaps some everyday practices of communication and decision-making depend on a pretence that we have common knowledge. That hardly comes as a surprise, for infinitely many of the propositions involved in common knowledge are too complex for humans to be psychologically capable of entertaining them. The present point is that the obstacles to entertaining them are not the only obstacles to knowing them.[2]

5.3 CLOSE POSSIBILITIES

A reliability condition on knowledge was implicit in the argument of section 5.1 and explicit in sections 4.3 and 4.4. We have seen that such a condition generates an obstacle to iterating knowledge. We can better understand the nature of the obstacle by considering reliability in the more general context of a family of related notions such as safety, stability, and robustness.

Imagine a ball at the bottom of a hole, and another balanced on the tip of a cone. Both are in equilibrium, but the equilibrium is stable in the former case, unstable in the latter. A slight breath of wind would blow the second ball off; the first ball is harder to shift. The second ball is in danger of falling; the first ball is safe. Although neither ball did in fact fall, the second could easily have fallen; the first could not. The stable equilibrium is robust; the unstable equilibrium, fragile.

Reliability and unreliability, stability and instability, safety and danger, robustness and fragility are modal states. They concern what could easily have happened. They depend on what happens under small variations in the initial conditions. If determinism holds, it follows from the initial conditions and the laws of nature that neither ball falls. But it does not follow that both balls were in stable equilibrium, safe from falling, for the initial conditions themselves could easily have been slightly different. There is a danger in a given case that an event of type E will occur (for example, that the ball will fall) if and only if in some sufficiently similar case an event of type E does occur. The danger is slight if E occurs in very few sufficiently similar cases, but that is not the same as a distant danger, which occurs only in insufficiently similar

[2] Fagin, Halpern, Moses, and Vardi 1995: 395–422 discuss some weakenings of the notion of common knowledge. In general, forms of almost-common knowledge do not imply almost the same behaviour as common knowledge itself, which raises difficult problems beyond the scope of this book. Shin and Williamson 1996 discuss common belief in a context of inexact knowledge.

cases. The relevant similarity is in the initial conditions, not in the final outcome (with the laws presumably held fixed). 'Initial' here refers to the time of the case, not to the beginning of the universe; I may be safe once I have caught the last flight out of the besieged city, even though I could easily have been a few minutes late and missed the flight, in which case I should now have been in danger. Safety and danger are highly contingent and temporary matters. Just how similar the case must be to one in which an event of type E occurs for the term 'danger' to apply depends on the context in which the term is being used.[3]

Reliability resembles safety, stability, and robustness. These terms can all be understood in several ways, of course. For present purposes, we are interested in a notion of reliability on which, in given circumstances, something happens reliably if and only if it is not in danger of not happening. That is, it happens reliably in a case α if and only if it happens (reliably or not) in every case similar enough to α. In particular, one avoids false belief reliably in α if and only if one avoids false belief in every case similar enough to α. When the danger is a matter of degree, reliability involves a trade-off between the degree to which the danger is realized and the closeness of the case in which it is realized. A very high degree of realization in a not very close case and a lower degree of realization in a closer case both make for unreliability. The argument of section 4.3 involved such a trade-off, the closeness of case α_{i+1} to case α_i compensating for the slightly lower degree of belief in α_{i+1}.

On a topological conception, a point x counts as safely in a region R if and only if x is in the interior of R. If R is a region in a metric space defined by some real-valued measure of distance, x is in the interior of R if and only if for at least one positive real number c, every point whose distance from x is less than c belongs to R. More generally, x belongs to the interior of R if and only if x belongs to some open subset of R. There is no difficulty in iterating safety on this conception, for the interior of the interior of R is just the interior of R. Thus x is safely safely in R—that is, safely in the region that contains all and only the points that are safely in R—if and only if x is safely in R. For if x is safely in R, then, for some non-zero distance c, every point less than c from x is in R, so every point less than $c/2$ from a point less than $c/2$ from x is in R, so every point less than $c/2$ from x is safely in R, so x is safely safely in R. On a corresponding conception of stability, a ball balanced in an indentation on the tip of the cone is in stable equilibrium, no matter how small and shallow the indentation.

For most practical purposes, the topological conception is not the

[3] See Sainsbury 1997 and Peacocke 1999: 310–28 for more discussion of the notion of easy possibility. It is applied in Williamson 1994b: 226–30.

one we need. The indentation must be of a certain size and depth for the ball not to be blown off by prevalent light breezes. To be safe on the top of a cliff, a young child must be at least three feet from the edge; it is not enough to be some positive distance or other, no matter how small, from the edge. Naturally, features of the context may contribute to fixing the margin for something to count as 'safe': for example, the severity of the consequences if one succumbs. Suppose that in some context a point is safely in a region if and only if every point less than three feet away is in the region. Then a point can be safely in a region R without being safely safely in R, for if the nearest point to x not in R is four feet away, x is safely in R but only two feet from a point two feet from a point not in R, so x is two feet from a point not safely in R, so x is not safely safely in R. The notion of what could easily happen behaves like the dual of safety; 'It could easily have been F' is close to 'It was not safely not F'. If it could easily have happened that an event of type E could easily have happened, it does not follow that an event of type E could easily have happened. For example, if exactly i humans were now alive, then it would be the case that it could easily have happened that exactly $i+1$ humans were now alive, but for some sufficiently large number k it would not be the case that it could easily have happened that exactly $i+k$ humans were now alive. If the actual number is i, then it could easily have happened that it could easily have happened . . . [k times] . . . that exactly $i+k$ humans were alive now, but it could not easily have happened that exactly $i+k$ humans were alive now. Thus iterations of 'it could easily have happened that' do not collapse.

The failures of knowledge to iterate observed in sections 5.1 and 5.2 are closely related to the failure of safety and reliability to iterate. One can be safe without being safely safe. In particular, one can be safe from error without being safely safe from error. One can be reliable without being reliably reliable. Since knowledge requires reliability, it is hardly surprising that one can know without knowing that one knows.

Safety is hard to iterate. For each natural number k, we can define x to be safely^k in R if and only if x is safely safely . . . [k times] . . . in R; x is safely$^\omega$ in R if and only if x is safely^k in R for every natural number k. Suppose that for some fixed non-zero distance c, a point is safely in a region if and only if every point less than c from the point is in the region. In n-dimensional Euclidean space, any two points are linked by a finite sequence of intermediate points each less than c from the next. Thus, unless R is the whole space, no point is safely$^\omega$ in R. A luminous condition resembles a region every point in which is safely in it; consequently, every point in such a region is safely$^\omega$ in it. In this instance, the only such regions of Euclidean space are the whole space and the null region. Similarly, we might think of a formula A as luminous in a system of epistemic logic if and only if A ⊃ KA is a theorem. The analogous feature would then be that A ⊃ KA is a theorem only if either A is a theorem (A corresponds to a region that is the whole space) or its negation is a theorem (A corresponds to the null region). Some natural systems have that property (see Appendix 2 and Williamson 1992a).

If R is the complement in full Euclidean space of a non-null bounded region (a sphere, for example), then for every natural number k some points are safely^k in R, even though no point is safely$^\omega$ in R. But if R itself is a bounded region, then for some natural number k no point is even safely^k in R.

Euclidean space is not the only kind of space, of course. We should not assume without argument that the space of possibilities in which we are interested has a Euclidean structure. In principle, it might consist of several disconnected regions. Every point in one of those regions might be safely in it; consequently, every point in the region is safely$^\omega$ in it. We also cannot assume that the required margin for safety c is uniform throughout the space. Prevailing winds may be stronger in some areas than in others. If they have a prevailing direction, one may be more easily blown from x to y than from y to x. Suppose, for example, that the closer one comes to a fixed point z_0 the more conditions favour stability. We can imagine contexts in which the required margin for safety at each point is its distance from z_0. Thus, unless x is z_0 itself, any point y is easily accessible from x if and only if y is closer to x than z_0 is; x is safely in a region R if and only if every point easily accessible from x is in R. If we fix a margin for safety at z_0 too, every point has a margin for safety. But since z_0 is accessible from no point other than itself, every point in the region consisting of the whole space except for z_0 is safely in that region. Thus every point in that region is safely$^\omega$ in it. Formally, such examples model non-trivial luminous conditions. Chapter 4 indicates that such a model would not be an accurate representation of knowledge.

Suppose that one is in a position to know only if one is safe from error in the relevant respect. We might try to deduce that, if a condition C can obtain without safely obtaining, then C can obtain even if one is not in a position to know that C obtains, and therefore that C is not luminous. The idea would be that one is in a position to know that C obtains only if one is safe from error in believing that C obtains, which requires C to obtain safely. But that is too quick. To be safe from error in believing that C obtains is to be safe from falsely believing that C obtains. Thus in a case α one is safe from error in believing that C obtains if and only if there is no case close to α in which one falsely believes that C obtains. But even if in α one believes that C obtains and is safe from error in doing so, it does not follow that C obtains in every case close to α, for there may be cases close to α in which C does not obtain and one does not believe that it obtains. One can believe that C obtains and be safe from error in doing so even if C does not safely obtain, if whether one believes is sufficiently sensitive to whether C obtains. For example, one may be safe from error in believing that the child is not falling even though she is not safe from falling, if one is in a good position to see her but not to help her.

We need a further assumption to generate an argument against luminosity. If we combine the safety from error requirement on knowledge with limited discrimination in the belief-forming process and some plausible background assumptions, then we can deduce failures of luminosity. That is not intended to formalize the anti-luminosity argument of Chapter 4, which depends on applying reliability considerations in a subtler way to degrees of confidence. The argument below models those considerations under highly simplified assumptions, which permit us to restrict our attention to the binary contrast between believing and not believing. It explains how the model falsifies luminosity and verifies a margin for error principle.

Suppose that for some parameter v, such as the height of the tree, for every case α, whether the condition C obtains in α depends only on the value $v(\alpha)$ of v in α. For example, C might be the condition that the tree is at most fifty feet high. We may assume for simplicity that v takes non-negative real numbers as values. To be explicit:

(7) For all cases α and β, if $v(\alpha) = v(\beta)$ then C obtains in α if and only if C obtains in β.

In many examples, something like the following will hold, for some small positive real number c:

(8) For all cases α and non-negative real numbers u, if $|u - v(\alpha)| < c$ and in α one believes that C obtains then, for some case β close to α, $v(\beta) = u$ and in β one believes that C obtains.

Less formally: if one has the belief, then one could easily still have had it if the parameter had taken a given slightly different value. One's belief is not perfectly discriminating. As already noted, iterations of close possibility do not collapse, so (8) does not entail that, if one has the belief, then one could easily still have had it if the parameter had taken a very different value. If one believes that the tree is at most fifty feet high, then one could easily still have believed that if the tree had been an inch higher, but not if it had been one hundred feet higher.

Now assume a connection between knowledge and safety from error:

(9) For all cases α and β, if β is close to α and in α one knows that C obtains, then in β one does not falsely believe that C obtains.

In a more careful version of (9), we might qualify both 'know' and 'believe' by 'on a basis B'. Knowledge on one basis (for example, seeing an event) is quite consistent with false belief in a close case on a very different basis (for example, hearing about the event). We might also relativize (8) and (9) to a subclass of cases by restricting the quantifiers over cases to that subclass. The argument below will still go through if we modify the other propositions in the same way. For simplicity, we may ignore these complications.

We must also articulate a connection between knowing and being in a position to know. One is in a position to know something determined by the value of a parameter only if one can know without changing the value of the parameter:

(10) For all cases α, if in α one is in a position to know that C obtains then, for some case β, $v(\alpha) = v(\beta)$ and in β one knows that C obtains.

Statement (10) can be understood as a stipulation about the meaning of 'in a position to know'.

Finally, we assume that knowledge implies belief:

(11) For all cases α, if in α one knows that C obtains then in α one believes that C obtains.

From (7)–(11) and the assumption (L) that C is a luminous condition, we can deduce this:

(12) For all cases α and β, if $|v(\alpha)-v(\beta)|<c$ then C obtains in α if and only if C obtains in β.

For suppose that C obtains in α and $|v(\alpha)-v(\beta)|<c$. By (L), in α one is in a position to know that C obtains. By (10), for some case α^*, $v(\alpha) = v(\alpha^*)$ and in α^* one knows that C obtains. Thus $|v(\alpha^*)-v(\beta)|<c$ and, by (11), in α^* one believes that C obtains. Consequently, by (8), for some case β^* close to α^*, $v(\beta^*) = v(\beta)$ and in β^* one believes that C obtains. Since β^* is close to α^* and in α^* one knows that C obtains, by (9) in β^* one does not falsely believe that C obtains. Therefore, C obtains in β^*. Since $v(\beta^*) = v(\beta)$, C obtains in β by (7). This shows that if $|v(\alpha)-v(\beta)|<c$ then C obtains in α only if C obtains in β. The converse is similar.

Statement (12) is a disastrous conclusion if the parameter v can vary continuously in this sense:

(13) For all non-negative real numbers u, for some case α, $v(\alpha) = u$.

For (12) and (13) entail:

(14) For all cases α and β, C obtains in α if and only if C obtains in β.

For any real number can be reached from any other in a series of arbitrarily short steps; there will be a sequence of non-negative real numbers $u_0, \ldots, u_n$ such that $u_0 = v(\alpha)$, $u_n = v(\beta)$, and, for all i, $(0 \le i < n)$, $|u_i - u_{i+1}| < c$. By (13), there is a corresponding sequence of cases $\alpha_0, \ldots, \alpha_n$ such that $v(\alpha_i) = u_i$ for all i $(0 \le i \le n)$, where $\alpha_0 = \alpha$ and $\alpha_n = \beta$. Consequently, for all i $(0 \le i < n)$, $|v(\alpha_i) - v(\alpha_{i+1})| < c$, so, by (12), C obtains in α_i if and only if C obtains in α_{i+1}. By the transitivity of the biconditional, C obtains in α if and only if C obtains in β. Thus C obtains in all cases or in none; it is trivial. Contrapositively, if C is not trivial and the assumptions (7)–(11) and (13) hold, then C is not luminous.

If we like, we can replace the assumption (13) that the parameter v varies continuously by the weaker assumption that v varies in an approximately continuous way, in the sense that for every non-negative real number u there is a case α such that $|u - v(\alpha)| < c/3$.

When we drop the luminosity assumption (L), we can still deduce this consequence from (7)–(11):

(15) For all cases α and β, if $|v(\alpha) - v(\beta)| < c$ and in α one is in a position to know that C obtains then C obtains in β.

The argument for (15) is like the argument for (12), but without the initial application of (L). Statement (15) is a *margin for error* principle: one knows that a condition obtains only if it obtains in all cases in which the relevant parameter differs at most slightly in value. The disastrous conclusion (14) that C is trivial follows easily from (13), (15), and (L). Since (13) or a suitable weakening of it is usually uncontentious, the margin for error principle usually blocks luminosity.

The margin for error c may depend on the condition C. However, if conditions $C_0, \ldots, C_k$ satisfy (15) with respect to margins for error $c_0, \ldots, c_k$ respectively (for the same parameter v), then of course $C_0, \ldots, C_k$ all satisfy (15) with respect to the minimum of $c_0, \ldots, c_k$. But an infinite class of conditions each with a positive margin for error might not have a common positive margin for error, for the greatest lower bound of their individual margins for error might be 0.

The argument of this section does not justify us in believing that every condition satisfies a principle like (15). The argument for (15) depends on the premise (8), that one's belief is not perfectly discriminating with respect to the underlying parameter. That assumption is not obvious, especially if the underlying parameter itself constitutively depends on one's belief, as some philosophers postulate for phenomena that they would classify as response-dependent. For example, they hold that the intensity of one's pain constitutively depends on one's beliefs about the intensity of one's pain. Such cases require the subtler argument of Chapter 4. Nevertheless, the assumptions (7)–(11) and (13) are plausible in a wide range of cases; they explain margins for error and the failure of luminosity. In particular, if the condition that one knows (or that one knows^k) that C obtains satisfies anything like (15) in place of C—naturally, with a parameter v that encodes enough about the case to determine whether one knows—then one will expect just the kind of difficulty in iterating knowledge that sections 5.1 and 5.2 observed. In particular, the crucial premises (1_i) and (1_i^k) simply attribute to Mr Magoo knowledge of a contraposed instance of the margin for error principle (15). Every iteration requires a further margin.

5.4 POINT ESTIMATES

I might reach my belief about the height of a tree by estimating its height and then applying an upper bound on the inaccuracy of my estimate.[4] For example, I estimate that the tree is 55 feet high, and come to believe that it is between 50 and 60 feet high, on the grounds that in these circumstances my estimate will not be out by more than 5 feet. In effect, I deduce (18) from the premises (16) and (17):

(16) I estimated that the tree is fifty-five feet high.

(17) My estimate of the height of the tree differs from the height of the tree by at most five feet.

(18) The tree is between fifty and sixty feet high.

Since I reached the conclusion (18) by inference from (16) and (17), I know (18) if and only if I know (16) and (17). Suppose that in these circumstances my estimates are never out by more than five feet, but are sometimes out by as much as five feet. Thus I might estimate that the

[4] Sections 5.4 and 5.5 answer points raised by Peter Mott 1998 and others. Williamson 2000b considers some further details of Mott's arguments.

tree is fifty-five feet high when it is in fact fifty feet high. In that case, it may appear, I can know (18) without satisfying any principle like (15). If the tree were even slightly less tall, my belief would be false.

The objection assumes that I can know that my estimate was out by at most five feet when it was in fact out by exactly five feet. That is in effect to assume that I need no margin for error in my knowledge of the accuracy of my own estimates. But my belief about my own accuracy has no more exact basis than my perceptual beliefs. If I were further away from the tree, or the light were worse, my estimate could be out by more than five feet. My judgement that my estimate of the height of this tree is out by at most five feet depends on my perceptual beliefs about my distance from the tree and the quality of the light. If I believe that my estimate is out by at most five feet, when in fact it is out by exactly five feet, then I could easily have formed that belief in slightly different circumstances in which my estimate was out by slightly more than five feet. Certainly the objector has not shown that one can know in the envisaged circumstances that one's estimate is out by at most five feet when in fact it is out by exactly five feet. There is almost no limit to how far out my estimates can be on a really bad day.

If my estimate is more than five feet out, I cannot know that it is at most five feet out, simply because knowledge is factive. A margin for error principle exhibits a further way in which my knowledge of the accuracy of my estimate depends on the accuracy of that estimate.

Naturally, we can imagine situations in which one knows exactly how far out one's estimate can be, just as we can imagine situations in which one knows exactly how tall the tree is. But those situations involve ways of knowing quite different from those we actually employ. The objector has done nothing to show that our actual methods enable us to dispense with margins for error. When C is the condition that one's estimate is out by at most five feet, the premises of the argument for (15) remain plausible. If one's knowledge of upper bounds on the inaccuracy of one's estimate of the height of the tree satisfies margin for error principles, then one's derivative knowledge of the height of the tree will satisfy a corresponding margin for error principle.

5.5 ITERATED INTERPERSONAL KNOWLEDGE

Iterating knowledge is hard, whether it is knowing about one's own knowledge or knowing about another's. Do margin for error principles make it too hard?