



OXFORD JOURNALS
OXFORD UNIVERSITY PRESS

Mind Association

Does the Gettier Problem Rest on a Mistake?

Author(s): Richard L. Kirkham

Source: *Mind*, New Series, Vol. 93, No. 372 (Oct., 1984), pp. 501-513

Published by: Oxford University Press on behalf of the Mind Association

Stable URL: <http://www.jstor.org/stable/2254258>

Accessed: 24/09/2008 08:54

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=oup>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We work with the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact support@jstor.org.



Oxford University Press and Mind Association are collaborating with JSTOR to digitize, preserve and extend access to *Mind*.

<http://www.jstor.org>

Does the Gettier Problem Rest on a Mistake?

RICHARD L. KIRKHAM

Attempts to resolve the Gettier Problem¹ rest on the mistaken assumption that an analysis of knowledge can be found which is (a) generous enough to include as items of knowledge all, or most, of those beliefs we commonly regard as knowledge,² and (b) rigorous enough to exclude from the class of knowledge any beliefs held in real or hypothetical cases which we would agree on reflection are situations where the epistemic agent does not know the belief in question. In actual practice there has been an inconsistency in the methods used to determine whether a particular analysis meets both criteria: We are asked to judge whether an analysis is too exclusive (criterion (a)) by reference to our ordinary, unreflective use of the word 'know'; but we are to judge whether it is too inclusive (criterion (b)) by reference to hypothetical situations (sometimes very elaborate) which force us to be acutely attentive to our linguistic and conceptual intuitions about 'know'. (See for example the introduction to Pappas and Swain, *ibid.*) The inconsistency is a reflection of what has gone wrong in the literature on the Gettier Problem.

A Gettier type counterexample is used to show that a proposed analysis of knowledge is too inclusive. Such counterexamples are hypothetical situations in which (1) all of the conditions for knowledge specified in the analysis are met, but (2) the epistemic agent does not have knowledge because the conditions have been met only by dumb luck, by accident, by coincidence, or by some means we intuitively regard as illegitimate. A typical paper on the Gettier Problem begins with a counterexample to the most recent state-of-the-art analysis of knowledge. The author then goes on to propose a new, more demanding analysis which cannot be falsified by his counterexample (or any previous Gettier counterexample). The new analysis becomes the state-of-the-art for a year or two

1 Edmund L. Gettier, 'Is Justified True Belief Knowledge?', *Analysis*, XXIII, 6 (June 1963), pp. 121–123.

2 George S. Pappas and Marshall Swain, *Essays on Knowledge and Justification* (Ithaca: Cornell U. Press, 1978), Introduction, p. 21.

until, invariably, someone finds a counterexample to it. It is this process which has produced the jungle of Gettier literature. The method implied by the process is backward. I believe that eventually (perhaps after many years) this process would yield the conclusion that the *only* analyses which are immune to all Gettier type counterexamples are those with very powerful sceptical implications, that is, the only analyses which can meet criterion (b) above are those which *cannot* meet criterion (a). I propose a shortcut to the same conclusion: (1) I shall suggest a very demanding analysis of knowledge (without concern for the moment for its sceptical implications). (2) I shall show that the analysis is immune to any and all possible Gettier type counterexamples. (3) I shall show that *no possible analysis* which is less demanding (however slightly) than mine can be immune to all Gettier counterexamples. (4) Since my analysis is too demanding to meet criterion (a) and no analysis less demanding can meet criterion (b), I shall conclude that no analysis can meet both criteria. Finally, I shall discuss the implications of this conclusion.

I suggest the following as an analysis of 'Smith knows that p':
Smith knows that p iff

- (i) p is true,
- (ii) Smith believes that p, and
- (iii) *either* p is self-evident for Smith *or* Smith has validly deduced p from ultimate premises which are all self-evident for Smith.

I am using 'self-evident' in the following sense: p is self-evident for Smith iff

- (i) if Smith believes p, then p is true and,
- (ii) Smith would point to the fact that p cannot fail to be true whenever he believes it as (one of) his rational reason(s) for believing it.

Given the analyses of 'knowledge' and 'self-evident', very few propositions could be known. Among the few which could be are: 'I believe something', 'I think my name is Smith' (not 'my name is Smith'), 'I am in pain', and 'Somebody believes something' (deduced from 'I believe something'). Necessary truths are also potential items of knowledge, since, if they are always true, then they are true whenever Smith believes them. A case could be made that nothing would ever qualify as knowledge on the above two

analyses, but we are not concerned (for the moment) with the sceptical implications of my analysis of knowledge.

To show that the above analysis of knowledge is immune from all possible Gettier type counterexamples it is only necessary to point out that the three conditions for knowledge cannot possibly be fulfilled by accident. In the case of a belief which is held because it is self-evident, how could it be self-evident by accident? And in the case of beliefs held because they are deduced from self-evident premises, how could the premises be accidentally self-evident or how could the belief be 'accidentally deduced'? It does not seem to make sense to talk of 'accidental self-evidence' or an 'accidental deduction'. What could such phrases mean? I take further indication, albeit not proof, that my analysis is immune to Gettier type counterexamples from the fact that none of the dozens of Gettier counterexamples in the literature would apply to it. Even if a Gettier counterexample can be found to my analysis, that fact would not refute what I am attempting to show in this article; for if an even more demanding analysis is required for immunity to all Gettier counterexamples, then my claim that no analysis can fulfil *both* criteria (a) and (b) above is made all the stronger.

The next step is to show that no analysis less demanding than mine can possibly be immune to all Gettier counterexamples. In sum the kind of justification needed for knowledge must have two characteristics: (1) it must begin with self-evident premises and (2) these premises must necessitate the truth of the proposition which is to be known. (And, of course, it is just these two characteristics which give my analysis its sceptical implications.) I shall argue that any analysis of knowledge which does not require a person to be justified in a manner that at least includes these two characteristics will be vulnerable to one or another of three Gettier type counterexamples. And my argument will be such that it will not matter what other conditions of knowledge are included in a proposed analysis of knowledge: It still will not avoid all three of the Gettier type counterexamples I shall be using, unless it also includes a justification condition which specifies that justification must have these two characteristics. Thus, my argument will also apply to the sort of Gettier solutions which have been called 'externalist' or 'non-evidential' (i.e. those which do not require that the would-be knower have any evidence at all). The only exceptions will be those analyses which imply scepticism every bit as strongly as mine does. Thus, the only solutions to the Gettier Problem which will be

invulnerable to all three of my Gettier counterexamples will be those which give no comfort to non-sceptics anyway. I shall not, of course, presuppose that my sense of knowledge is correct. Rather, I shall show that every possible weaker analysis can be falsified by one or another of three Gettier type counterexamples. And, thus, I shall conclude that no analysis weaker than mine is immune to all Gettier like counterexamples.

Before beginning the argument proper, let me point out that no analysis of knowledge will succeed if it does not require, in addition to whatever else it requires, that the would-be knower have some self-conscious evidence for his belief. For, if a would-be knower has a true belief, acquired *by whatever non-evidential process you wish*, but he has no self-conscious evidence at all for believing it, and, hence, no more *rational* reason for believing it than he does for many false beliefs he has acquired by whim or indoctrination or whatever, then intuitively we would say that he does not know. Thus, we can eliminate all *pure* non-evidentialist or externalist analyses of knowledge. What remain to be considered are the purely evidentialist analyses and the mixed analyses which contain a justification condition and a fourth, non-evidentialist condition for knowledge.

Let us consider first an analysis of knowledge which does have a justification condition but which does *not* require that justification begin with self-evident premises. For any analysis of this sort, it will always be possible to construct a Gettier type counterexample based on the idea that the ultimate premises are believed only by dumb luck, by accident, or by some other illegitimate means. Consider Smith who is allergic to cheese. His idiot cousin Ernie tells him that the moon is made of green cheese. Smith believes Ernie and deduces that if he (Smith) eats two pounds of the moon, he will get sick. And let us stipulate that it is true that anyone who eats two pounds of the moon will get sick. Hence, Smith has a justified true belief, but Smith does not know.

Some would be tempted to suggest that the problem here is that Smith's ultimate premise is false. The fact that it is not self-evident, they would claim, has nothing to do with why we would say that he does not know. But variations of the idiot cousin Ernie case can be constructed in which Smith's ultimate premise is true. Suppose Ernie tells Smith one hundred things about the moon. One of them, by Ernie's dumb luck, happens to be true: The moon is too heavy for any mortal to lift. Smith believes all one hundred things Ernie

tells him and deduces something from each one. Some of his conclusions are false, some true. From the proposition that the moon is too heavy for any mortal to lift, Smith deduces that if he is mortal then he cannot lift the moon. Smith has another justified true belief, but he still does not have knowledge, since he believes the premise only because he heard Ernie say it and it is only dumb luck that what Ernie said is true. (Note that Smith does not believe that Ernie is always or even usually right. It is simply a fact that Smith believes what Ernie tells him. Smith has never entertained, much less believes, the proposition 'Ernie is right'. So *Smith's deduction does not rely on any false belief*. It is true that Smith's behaviour in believing what Ernie tells him would normally be good evidence that he also believes the proposition 'Ernie is usually right'. But unless we are prepared to endorse behaviourism, the possibility that Smith does not have the latter belief is not eliminated; and, since this is my own hypothetical story anyway, I can stipulate that Smith does not have the latter belief.)

It might be thought that all we must do to handle cases of this sort is to strengthen the justification condition by requiring that the ultimate premises be justified in some appropriate way. But this will not do, for, if the ultimate premises are so justified, then they are not the ultimate premises. They are, in fact, justified with reference to some other, more ultimate, premises. So then it will be possible to construct an 'idiot cousin Ernie' counterexample based on the idea that these other, more ultimate, premises are arrived at by luck, or accident, or some such. The principle behind this is that if an ultimate premise truly is ultimate, then either it is self-evidenced in some way or it is not justified at all. Thus, the idiot cousin Ernie case shows not just that some or another kind of justification is needed for the ultimate premises, it shows precisely that they must be *self-evident*, for there is no other kind of justification for *ultimate* premises.

Perhaps the problem of the ultimate premises can be solved by introducing non-evidential requirements to ensure that the ultimate premises are not believed just by luck or by accident. Perhaps, taking a cue from Goldman,¹ we could say that 'the fact that p' must be causally connected to Smith's belief in the ultimate premises from which he deduces his 'belief that p'. However,

1 Alvin I. Goldman, 'A Causal Theory of Knowing,' *Journal of Philosophy*, LXIV. 12 (June 1967), pp. 355-372, reprinted in Pappas and Swain, eds. *ibid.* pp. 67-86, *passim*.

Goldman realizes that the would-be knower must be able to *correctly* reconstruct the causal chain. Moreover, Goldman is aware that the would-be knower might be able to reconstruct the causal chain just by a series of lucky guesses:

An additional requirement for knowledge based on inference is that the knower's inferences be warranted. That is, the proposition on which he bases his belief of P must genuinely confirm P. . . . Reconstructing a causal chain merely by lucky guesses does not yield knowledge.¹

But what about these propositions themselves which are a part of his reconstruction or on which his reconstruction is based? These have now become the new ultimate premises. How will we ensure that they are not believed just by luck? If we say that these, too, must be the result of a causal chain, then this chain will also have to be correctly reconstructed, and this reconstruction will have to be warranted by still more ultimate premises. And so on *ad infinitum*.

The situation cannot be improved by insisting that Smith's ultimate premises must be arrived at by a reliable method. Suppose he has three methods for attaining new beliefs. One of them has always led him to believe true things while the other two have always led him to believe false things. But suppose he is not aware of the difference in reliability of the three methods. He has always confidently believed every conclusion of all three methods. Wondering whether or not the moon is too heavy for a mortal to lift, Smith chooses *at random* one of the methods and, as it happens, it is the method which has always been reliable. The method leads him to believe that the moon is too heavy and, hence, that if he is mortal he cannot lift it. Yet, intuitively we would say that Smith does not have knowledge, because it is only by dumb luck that he used a reliable method. Would it help to insist that he be warranted in thinking his method is reliable? But what about the premises on which this latter warrant is based? These are now the ultimate premises of his belief that if he is mortal he cannot lift the moon. Why does he believe *them*? Because his idiot cousin Ernie told him to?

I am unable to conceive any manner in which we might try to ensure that the ultimate premises are not believed just by luck which would not simply create new, more ultimate, premises.

1 Goldman, *ibid.*, reprinted in Pappas and Swain, *ibid.* p. 75.

Hence, I conclude that some Gettier type counterexample based on the idea that the ultimate premises are believed by dumb luck, or by accident, or by coincidence, will work against any analysis of knowledge which does not specify that the would-be knower's ultimate premises must be self-evident.

So the search for an analysis which is immune to all Gettier type counterexamples, but which does not have the apparent sceptical implications of my sense, must now focus on those analyses of knowledge which allow that the inference from the premises to the belief in question need not be deductive, that is, the truth of the belief need not be necessitated by the evidence in its favour. It is my contention that any justification condition of this sort will be vulnerable to at least one of two kinds of Gettier like counterexamples. The first of these two kinds, called 'the penguin case', is based on the idea that a method of inference which consistently leads to false belief does not yield knowledge even in the occasional case when, by luck, it produces a true belief. The second kind, called 'the cube case', is based on the idea that, however reliable a method of inference usually is, it does not yield knowledge in the occasional case when it misfires (but a true belief is produced anyway purely by coincidence).

First imagine as strong a method of inference as you can short of deduction. Make it inductive, or inference-to-the-best-explanation, or a reconstruction of a causal chain, or whatever you wish. Now strengthen the method of inference as you like with requirements of indefeasibility and not reasoning through a false premise. In short, your assignment here is to devise the most rigorous method of inference you can, subject only to the constraint that the conclusion of such an inference is not necessitated by the ultimate premises. Call this 'Method of Inference A'. If you prefer, do not think of Method A as a purely inferential process: Think of it only as what Nozick calls a 'way of coming to believe'¹ propositions which has some non-evidential elements. The only constraint is that the use of the method may not logically necessitate the truth of the resulting beliefs.

Smith has gone to Antarctica to study penguins. Examining first some question about the fur of penguins, Smith applies Method A and comes up with a false conclusion. This is possible because in

1 Robert Nozick, *Philosophical Explanations* (Cambridge, Mass.: Harvard University Press, 1981), p. 179.

using Method A the premises do not, nor does the method, necessitate the truth of the conclusion. For a method with this characteristic it is *always* possible that something has gone wrong, even when one begins with true premises. Smith does not know his conclusion is false, but we in our ideal observer position do know that it is false. (All Gettier counterexamples put the reader in the position of an ideal observer who knows all the relevant facts about the would-be knower's evidence, how he acquired it, and whether or not his belief is true.) We know that he does not know his conclusion because the truth condition for knowledge has not been fulfilled. Smith applies Method A to another matter about penguins and, again, it issues in a false belief. He uses it again and again, always believing the conclusions, but (as we the ideal observers know) something always goes wrong. His resulting beliefs are always false. Finally, on the one hundred and first application, Method A issues in a true belief. At last Smith has a justified true belief (justified via Method A at any rate), but he does not have knowledge. He just got lucky; Method A finally produced a true belief.

The above counterexample will work against *any* analysis of knowledge that includes a justification condition (or 'way of coming to believe') which does *not* require that, to be justified sufficiently to be known, the truth of a belief must be *necessitated* by the evidence in its favour. Because, for *any* method which does not have this characteristic, it is logically possible that it could issue in a false conclusion, not just once but a hundred times in a row.

One proposed solution to the Gettier Problem which would be vulnerable to a Gettier type counterexample like the penguin story is Robert Nozick's. He proposes the following conditions for knowledge:

- (1) P is true.
- (2) S believes, via method or way of coming to believe M, that P.
- (3) If P weren't true and S were to use M to arrive at a belief whether (or not) P, then S wouldn't believe, via M, that P.
- (4) If P were true and S were to use M to arrive at a belief whether (or not) P, then S would believe, via M, that P.¹

Nozick says that a method which meets the latter two conditions is

1 Ibid.

'sensitive' to the truth value of P. But note that Nozick's analysis of 'knowing that P' makes the applicability of M depend solely on its sensitivity to the truth value of P alone. There is no requirement that M have a general reliability. Hence, Smith (perhaps thinking that M does have a general reliability) might use it one hundred times producing false beliefs each time. But if, by luck, Smith applies M, the one hundred and first time, to a proposition P such that M is sensitive to the truth value of P, then, on Nozick's view, we would have to say, counter-intuitively, that Smith knows P; for, with respect to P at least, all Nozick's conditions for knowledge are met. Moreover, the penguin case would be a counterexample to Nozick's analysis whether or not we think of M as a method of inference or only a 'way of coming to believe'.

What the penguin case suggests is that however rigorous a method of inference or 'way of coming to believe' is and however 'sensitive' to the truth value of a given proposition P it is, it will still not provide sufficient justification for knowing P in any given case *unless it has proven at least generally reliable in past instances*.

Let us suppose, then, an analysis of knowledge which requires that the belief in question be justified by some Method B such that Method B is just like Method A, except for the additional characteristic that it has always been reliable in the past.

But there is a Gettier type counterexample to this analysis as well. Suppose Jones has habitually used Method B all his life and it has never failed him. Jones walks into a room and Method B leads him to believe (or, if you prefer, to conclude), *correctly*, that there is a cube in the room and the cube is painted red. But suppose that, unknown to Jones there is a red light-bulb in the room's light-socket. If the cube had been painted white he would have believed it was painted red anyway. Although Method B has always been reliable in the past, it has finally misfired in this case: It is only by luck that the belief it produces is true. I think we would say, intuitively, that Jones does not know the cube is painted red.

A variation has Jones entering a room and correctly concluding that there is a white cube in the room. The light-bulb is white, but *if* the light-bulb had been red, then he would have falsely concluded, via Method B, that the cube was red. It is only lucky that the circumstances are such that Method B leads him to a true belief. Nozick, among others, has produced Gettier type counterexamples which parallel these two, but which are not cases of illusion (illusion luckily avoided, I should say); rather, they are based on some other

sort of error.¹ (Nozick, of course, would say that the problem here is that Method B is not 'sensitive' to the truth value of the propositions in these two cases. I shall discuss this suggestion below).

A misfiring of Method B is possible because the truth of the premises does not necessitate the truth of the belief. (Putting the same point another way, if Method B is thought of only as a 'way of coming to believe,' then the use of the method does not guarantee the truth of the resulting belief.)

Now it might be objected that if Method B were really as rigorous as it supposed to be, then it will include requirements to the effect that Jones must look to see what colour the light-bulb is. But this would miss the point of the cube stories. Even if it is true that, for any given sort of illusion, we could create a Method B such that the possibility of *that* sort of illusion is eliminated, it is not the case that Method B could eliminate every sort of illusion or error; for if Method B involved gathering enough premises to eliminate all possibility of error, then the truth of the premises would necessitate the truth of the belief produced. Which would make the analysis of knowledge of which it is a part every bit as demanding as mine.² Remember, the point here is to find some justification condition which is at least *slightly weaker* than mine but which is still immune to all Gettier type counterexamples. We are testing whether or not a justification condition which does *not* require that the premises necessitate the conclusion can have the needed immunity. We have already seen that some clause requiring reliability must be a part of such a justification condition if it is to be immune to the penguin case. We have now seen that, *if this slightly weaker justification condition really is weaker*, then it cannot eliminate every possibility of error, thus, it will always be possible that error is avoided only by dumb luck and this means that the slightly weaker justification condition will be subject to some Gettier type counterexample like the cube case.

1 Ibid. p. 175, see also Ernest Sosa, 'Propositional Knowledge', *Philosophical Studies*, XX, 3 (April 1969), pp. 39 ff.; and Keith Lehrer and Thomas Paxson, Jr., 'Knowledge: Undefeated Justified True Belief', *Journal of Philosophy*, LXVI, (1969), pp. 236 ff.

2 Some might suggest that if Method B had an appropriate defeasibility clause (see Lehrer and Paxson; *ibid.*), then it would prevent the cube illusions. But here, again, if the defeasibility clause is strong enough to eliminate every possibility of error, then Method B is not weaker than my kind of justification anyway; for the use of the method would necessitate the truth of the belief produced.

The red cube case shows that Method B cannot assure that Jones would disbelieve the proposition that the cube is red even if it were false. And the white cube case shows that Method B cannot assure that Jones would always believe that the cube is white whenever it is white. It was for these reasons that Nozick proposed his third and fourth conditions for knowledge, quoted above. Perhaps, then, Nozick's two conditions (interpreted as strict implications for the time being instead of as subjunctive conditionals the way Nozick actually interpreted them) along with a clause about reliability will give Method B the needed immunity from Gettier type counterexamples. I think, in fact, that such a method or 'way of coming to believe' would have this immunity, but it would have it only because it would not really be weaker than my justification condition; for to meet Nozick's third and fourth conditions (interpreted as strict implications) the method would have to eliminate all possibility of error. And, as I pointed out above, the point here is to find a justification condition which is *weaker* than mine and which is also immune to the Gettier examples.

Thus, I conclude that *no* analysis of knowledge which does not require that the would-be knower be justified such that his premises necessitate his belief is immune to all Gettier type counterexamples: All analyses which do not require that the ultimate premises be self-evident are, as I argued above, vulnerable to the 'idiot cousin Ernie' counterexample. Hence, despite its apparent sceptical implications, my analysis of knowledge is the weakest analysis which is immune to all Gettier like counterexamples.

Still some philosophers would rebel against the apparent sceptical implications of my justification condition. Nozick is one such philosopher and his reluctance to accept scepticism apparently motivated him to weaken the sense of his third and fourth conditions for knowledge. I have been treating these two conditionals as strict implications. In fact, Nozick calls them subjunctive conditionals and he does not intend for them to rule out every possibility of illusion (luckily avoided) nor, hence, every Gettier counterexample. He wants only to protect his analysis from those counterexamples which are actualized on possible worlds 'close' to this one.¹ He never makes clear exactly what possible worlds he has in mind, but at any rate he seems to realize that some Gettier type counterexamples will have force against his analysis because by

1 Nozick, *ibid.* p. 173.

limiting the applicability of his third and fourth conditions for knowledge to possible worlds 'close' to this one, he is not eliminating every possible situation in which an epistemic agent uses method M and yet avoids error only by dumb luck.

Other writers on the Gettier Problem have wanted to set out the necessary and sufficient conditions for knowledge. Nozick seems to realize that his analysis is not logically sufficient, so exactly what it is supposed to be is not at all clear. He says in one place that his analysis is not intended to clear up all cases where the applicability of the term 'know' is itself unclear at the preanalytic level.¹ But 'leaving some cases unclear' and 'yielding to falsifying counterexamples' are two different things. Even if the former is not a vice for an analysis of knowledge, the latter certainly is.²

The first implication of my arguments is that the assumption on which the Gettier Problem rests (described in the first sentence) is mistaken; for, since my analysis is too demanding to meet criterion (a) of the assumption and no analysis less demanding than mine can meet criterion (b), there is no analysis of knowledge which can meet both. In the face of these results we have several options: We could conclude that there are two senses of knowledge. One of them is at work in ordinary life when we say that we know all kinds of propositions including many which are not justified in the manner my analysis requires. The other is revealed when Gettier counterexamples force us to be acutely attentive to our conceptual intuitions about knowledge. On the other hand, we could simply conclude that the concept of knowledge is hopelessly confused and contradictory; hence, no analysis is possible. The option I prefer is to conclude that the Gettier counterexamples reveal the correct analysis of knowledge and, thus, if my arguments have been cogent, a very radical form of scepticism is correct. This conclusion entails that most of the knowledge claims we make in ordinary life are simply incorrect. My preferred choice is more plausible if we

¹ Ibid. p. 192.

² Fred Dretske has proposed that, in order to know, one's reasons, R must be such that $\sim \diamond (R \ \& \ \sim P)$ in the given situation (where P is the would-be knower's belief). Dretske intends that the reference to the given situation will put the same kind of limitations on the scope of the modal formula as Nozick achieves by calling his conditionals 'subjunctive'. In other words, Dretske's modal formula is not intended to rule out every possibility of error. It should be clear, then, that Dretske's formula will be vulnerable to a cube style counterexample as well as, of course, the penguin case. See Fred Dretske, 'Conclusive Reasons', *Australasian Journal of Philosophy*, XLIX, 1 (May 1971), pp. 1-22.

remember that a belief or proposition does not become less valuable merely because we can no longer apply the 'hurrah' word 'knowledge' to it. Only the discovery that it has less justification than we thought it had can cause it to lose epistemic value. My inductively well justified belief that there is a blue typewriter here is no less reliable or pragmatically useful (nor is it less justified) when I discover that I cannot say that I *know* there is a blue typewriter here. At least not when the discovery is only based on my finding that knowledge requires more justification than I thought it required (and more than I ever had for my belief in the blue typewriter). The discovery that I do not know the typewriter is here signals a reduced value for my belief in the typewriter only when the discovery is based on my finding that I have less justification for the belief than I thought I had. Hence, the kind of scepticism my analysis of knowledge implies is not the sort which should cause anyone discomfort.¹

DEPARTMENT OF PHILOSOPHY,
UNIVERSITY OF NOTRE DAME,
NOTRE DAME,
INDIANA, 46556,
U.S.A.

1 Cf. William W. Rozeboom, 'Why I Know So Much More Than You Do', *American Philosophical Quarterly*, IV. 4 (October 1967), pp. 281-290.