THE IRREDUCIBILITY OF KNOWLEDGE

L. S. Carrier

I aim to show the impossibility of giving a reductive analysis of knowledge. My argument will depend upon characterizing the concept of knowledge as an epistemic concept. Other epistemic concepts besides knowing include perceiving and remembering such concepts introduce a propositional clause of the form 'that p'. Although I believe that my argument can also be extended to the epistemic concepts of perceiving and remembering, I shall concentrate exclusively on the concept of knowledge articulated in the sentence form 'A knows that p'. But first I must stipulate as to what constitutes an epistemic concept, and what constitutes a reductive analysis.

I

One mark of an epistemic concept is that the sentence expressing it is only a partial truth-function of the sentence embedded in the propositional clause. Its being true that p tells us nothing about the truth-value of 'A knows that p' (KAp); whereas, its being false that p does determine the falsity of KAp. The difference between a non-epistemic concept like negation and an epistemic concept like knowledge is shown by the following truth-tables:

(1)

p	~p
T	F
F	T

(2)

p	KAp
T	?
F	F

Having a 'mixed truth-table' like (2) is still only a necessary condition for being an epistemic concept, since KAp shares this same truth-table with the concept of logical ne-

cessity. So to distinguish 'A knows that p' from 'It is necessary that p' we must further require epistemic concepts to include a purely psychological component which, in the case of knowledge, is normally taken to be the 'propositional attitude' of *belief.*

Any *reductive* analysis of an epistemic concept must reduce this concept to other concepts,none of which is itself epistemic. But these other concepts must be non-epistemic in the different ways that reflect the logical difference between the truth component and the purely psychological component. Since KAp entails 'It is true that p' (Tp), the truth component is given by the identity function in (3) below. Truth is thus a non-epistemic concept by reason of its being, like negation, an *extensional* concept, or one that is completely truth-functional.

(3)

p	Tp
T	T
F	F

Since epistemic concepts include a purely psychological component, this component must also be non-epistemic. But purely psychological concepts like belief are non-epistemic by virtue of their being *intentional* concepts. That is, the truth-value of the sentence expressing such a psychological concept will be completely undetermined by the truth-value of the sentence embedded in the propositional clause. This feature is present in 'A believes that p' (BAp), as shown in (4) below, but it is also shared by sentences expressing wishes, hopes, fears, and all other non-epistemic propositional attitudes:

(4)

p	BAp
T	?
F	?

A reductive analysis of an epistemic concept, then, is one which would include at least one component that is extensional, at least one that is intentional, and none that is epistemic. The goal of a reductive analysis of knowledge is to split this concept into concepts which include those expressed by (3) and (4). But since mere true belief is not sufficient for knowledge, some other concept must also be included.

II

Since Plato's *Meno*, attempts have been made to formulate this additional concept in terms of A's having good reasons for his belief that p. But there is some controversy concerning whether of not the having of good reasons for belief is itself an intentional concept. Suppose that it is, with 'reasons' simply denoting other supporting beliefs that A has for his belief that p. Then (5) below gives the truth-functional relation between p and 'A has good reasons for believing that p' (RAp):

(5)

p	RAp
T	?
F	?

The concepts expressed by (3), (4), and (5) now provide a candidate for a reductive analysis of knowledge. But it is not a viable candidate. Counterexamples first raised by Bertrand Russell ([1]) and later popularized by Edmund Gettier ([2]) show that true belief with good reasons is still insufficient for knowledge. The counterexamples succeed because of the intentional interpretation of 'good reasons'. This interpretation

([1]) Bertrand Russell, *The Problems of Philosophy* (London: Oxford university Press, 1912), p. 132.

([2]) Edmund Gettier, 'Is Justified True Belief Knowledge ?', *Analysis*, 23 (1963), 121-123.

creates a logical gap between Tp and RAp, so that one can always introduce cases in which A's reasons for belief, although good ones, are irrelevant to the truth of what is believed.

For example, suppose that in July 1974 A came to the belief that Nixon would leave office prematurely, believing so for the good reason that the Watergate scandal would result in Nixon's removal. Though it is true that Nixon did leave office early, A did not know this in July if he failed to include in his reasons the belief that Nixon might resign rather than wait to be impeached and convicted. We refuse to attribute knowledge to people when their reasons for believing that p fail to explain why p is true. Since the possibility of such failure is a built-in feature of the intentional interpretation of 'good reasons' in (5), any reductive analysis of knowledge in terms of (3), (4), and (5) cannot succeed.

One way to seal this opening against counterexamples is to introduce another concept of good reasons (good reasons). Under this interpretation, it will be false to say that A has good reasons^2 for his belief that p if p is false. The truth-functional relation between p and 'A has good reasons^2 for believing that p' (R^2Ap) can then be expressed by:

(6)

p	R^2Ap
T	?
F	F

Now A's having good reasons^2 for believing that p is simply for A to be possessed of a set of other beliefs S such that if A has S, then it is true that p. But we can now abstract S itself from its truh-functional connection to p, just as we abstracted A's belief that p from his knowledge that p; so S becomes a purely psychological component of R^2Ap. But then (6) expresses an epistemic concept, just as (2) does; and thus there can be no

reductive analysis of knowledge in terms of true belief for which one has good reasons[2].

Attempts to provide a reductive analysis of knowledge by adding a fourth concept to those expressed by (3), (4), and (5) fare no better. For example, suppose we require that A's good reasons be grounded in what is 'directly evident' to A, and suppose that the directly evident is constituted by A's conscious experience. We must then ask whether the concept of the directly evident is itself an intentional concept. If it is, then a gap is created for the same sort of counterexample that defeated the analysis couched in terms of (3), (4), and (5) alone. As an instance, consider that A's reasons for believing that a duck is swimming in the lake are grounded in his *seeming to see* one there. But if what A seems to see is really a piece of debris, and not a duck, then even if there is a duck swimming out there someplace, A does not know it. We can, of course, bar the way to counterexamples like this by making what is directly evident provide at least a truth-functional guarantee for the truth of what is believed. But if we do this, then we are forced to construe the concept of the directly evident as an epistemic concept, like good reasons[2]; so any analysis which adds such a concept to those expressed by (3), (4), and (5) would not constitute a reductive analysis.

III

It has recently been suggested that we drop the requirement of good reasons (of either sort) in favor of either a 'nomic' account of knowledge, or one which is explicitly causal. The former type of account is defended by D. M. Armstrong[3] and by Peter Unger[4]; whereas, the latter type is espoused by Alvin Goldman.[5] For A to know (non-inferentially) that

(3) D.M. Armstrong, *A Materialist Theory of the Mind* (London: Routledge & Kegan Paul, 1968), p. 189.

(4) Peter Unger, 'An Analysis of Factual Knowledge', *Journal of Philosophy*, 65 (1968), 157-170.

(5) Alvin Goldman, 'A Causal Theory of Knowing', *Journal of Philosophy*, 64 (1967), 357-372.

p, Armstrong requires that p be 'empirically necessary' for A's belief that p; whereas, Unger requires that it be 'not at all accidental' that A's belief is true. Goldman, on the other hand, demands that p be 'causally connected in an appropriate way' with A's believing that p. Despite differences in detail, all of these accounts share the same logical basis; for each of the analyses exhibits a common relationship between p and A's belief that p, *viz.* that if A knows that p, then, had p been false, A would not have been come to believe that p. So it makes no difference whether we take B*Ap in (7) below to mean 'A's belief that p is empirically sufficient for p', or 'A's belief that p is not at all accidental', or 'A's belief that p was caused in the appropriate way by p'. We still get the following truth-functional relationship between p and B*Ap:

(7)

p	B*Ap
T	?
F	F

But since there is a purely psychological element present in B*Ap, *viz.* A's belief that p, the context created by (7) is no less epistemic than those exhibited in (2) and (6). So no nomic or causal analysis of knowledge in terms of (3), (4), and (7) can qualify as a reductive analysis. This is not to say that such accounts fail to say something correct about knowledge. For each of the analyses exhibits a common logical relationship for those seeking a reductive analysis, to be trivial; for we are just offered another form of words for the same epistemic concept.

IV

Yet another attempt to provide an analysis of knowledge is made by the Cartesian (or by the skeptic[6]), in which 'A knows that p' is held to entail 'A is absolutely certain that p'. Part of what is commonly meant by 'absolute certainty' is that the only propositions known with such certainty are those about which one has *indubitable* beliefs. To say that A has an indubitable belief about a proposition is just to say A believes this proposition, and that it is logically necessary that if A believes such a proposition, then this proposition is true. The only propositions known by A, according to high standard set by this analysis, typically turn out to be about A's existence or A's own mental states. Let us take p* to go proxy for any sentence usually said to describe such a state (e.g., 'I exist', 'I am in pain', or 'I seem to see a red patch'). The following truth-table then yields the truth-functional relationship between p* and 'A has an indubitable belief that p*' (BAp*):

(9)

p*	BAp*
T	?
F	F

But we now see immediately from (8) that BAp* is an epistemic concept for the same reason that what (6) expressed was. So no reductive analysis of knowledge results from our making use of the concept of indubitable belief.

Suppose that we decide to proceed further with the Cartesian program and characterize absolute certainty, not only in terms of indubitable belief, but also in terms of states that are *self-intimating*. To say that A is in a self-intimating state is just to say that A's being in such a state gives a logical

(6) Peter UNGER has retreated to such a view in 'A Defense of Skepticism', *Philosophical Review*, 80 (1971), 198-219.

guarantee that A believes that he is in it. Again, according to the program, it turns out that A's self-intimating states will be limited to his mental states. Let us now take B^2Ap^* to abbreviate 'A has an indubitable belief that he is in the self-intimating state that p*'. The truth-functional relationship between p* and B^2Ap^* will then be expressed by:

(9)

p*	B^2Ap^*
T	T
F	F

Now (9) does not yield an epistemic concept but instead reveals a purely *extensional* one. For this very reason, however, B^2Ap^* cannot express a concept that qualifies as an appropriate *analysans* for our concept of knowledge. If knowledge is limited to indubitable beliefs about self-intimating states, then knowledge is no longer being viewed as the epistemic concept which we agreed that it was. The full Cartesian approach fails to analyze the epistemic concept of knowledge with which we began, as the difference between (2) and (9) demonstrates. Instead, this approach has replaced our epistemic concept with an entirely different concept: a concept of absolute certainty that creates a truth-functional equivalence between the act and the objects of knowing.

V

I can think of no other plausible account of knowledge which is not a variant of one or more of the accounts mentioned above. (7) Since none of these accounts yields a reductive

(7) One might think that *some* form of the 'good reasons' approach can still succeed, but I think that this is an illusion. My reason for saying this is that there seem to be only three sorts of connections between having

analysis, I conclude that no such analysis can be given for knowledge. If my argument has been sound, what remains puzzling is that the only epistemic concept clearly susceptible to a correct reductive analysis is that of *true belief*. (We might construct yet other analogues, such as 'truly seeming to remember' and 'truly seeming to perceive', but these locutions are perhaps odd enough so as not to express concepts that we might profitably use.) True belief splits cleanly and without remainder into an extensional and an intentional component, as witness (3) and (4).

Given this striking logical dissimilarity between the concept of knowledge and that of true belief, it might now seem odd to include true belief among our epistemic concepts at all. One quick way to exclude it would be to make irreducibility a further mark of any epistemic concept. But this sly tactic would have all the advantages of theft over honest toil. It would only serve to label a difficulty rather than to explain it, and there would then be no clear way to say just what an epistemic concept was. But retaining sharp boundaries for our epistemic concepts leaves us with the problem of explaining why only *some* of our epistemic concepts are irreducible. To this problem I must confess the lack of any solution. I pretend only to have shown that knowledge is irreducible, not to have provided the explanation for its irreducibility.

The positive lesson to be learned, though, is not that the failure to provide reductive analyses somehow leaves our concepts unclear, but rather that when we find reductive analyses to be inapplicable, we should have recourse to a more humble analytical goal. Though epistemic concepts like knowledge are

good reasons for believing that p and p itself: (i) logical entailment, (ii) nomic or causal connection, and (iii) some evidential connection that is neither (i) nor (ii). Under interpretations (i) and (ii) the notion of good reasons is epistemic, and so no reductive analysis can result. Under interpretation (iii) the notion of good reasons must be intentional, and the counterexamples previously mentioned again intrude.

impervious to reductive analysis, they are still subject to illuminating analyses of a non-reductive sort. I hope to have shown that it is to this sort of conceptual clarification that our analytical aims are best directed.

University of Miami L. S. Carrier