

Journal of Philosophy, Inc.

Knowledge: Undefeated Justified True Belief

Author(s): Keith Lehrer and Thomas Paxson, Jr.

Source: *The Journal of Philosophy*, Vol. 66, No. 8 (Apr. 24, 1969), pp. 225-237

Published by: Journal of Philosophy, Inc.

Stable URL: <http://www.jstor.org/stable/2024435>

Accessed: 08/01/2009 17:26

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=jphil>.

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.



Journal of Philosophy, Inc. is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Philosophy*.

<http://www.jstor.org>

THE JOURNAL OF PHILOSOPHY

VOLUME LXVI, NO. 8, APRIL 24, 1969

KNOWLEDGE: UNDEFEATED JUSTIFIED TRUE BELIEF

IF a man knows that a statement is true even though there is no other statement that justifies his belief, then his knowledge is basic. Basic knowledge is completely justified true belief. On the other hand, if a man knows that a statement is true because there is some other statement that justifies his belief, then his knowledge is nonbasic. Nonbasic knowledge requires something in addition to completely justified true belief; for, though a statement completely justifies a man in his belief, there may be some true statement that *defeats* his justification. So, we must add the condition that his justification is not defeated. Nonbasic knowledge is undefeated justified true belief. These analyses will be elaborated below and subsequently defended against various alternative analyses.¹

I

We propose the following analysis of basic knowledge: *S* has basic knowledge that *h* if and only if (i) *h* is true, (ii) *S* believes that *h*, (iii) *S* is completely justified in believing that *h*, and (iv) the satisfaction of condition (iii) does not depend on any evidence *p* justifying *S* in believing that *h*. The third condition is used in such a way that it entails neither the second condition nor the first. A person can be completely justified in believing that *h*, even though, irrationally, he does not; and a person can be completely justified in

¹ This analysis of knowledge is a modification of an earlier analysis proposed by Keith Lehrer, "Knowledge, Truth, and Evidence," *Analysis*, xxv.5, 107 (April 1965): 168-175. It is intended to cope with objections to that article raised by Gilbert H. Harman in "Lehrer on Knowledge," this JOURNAL, LXIII, 9 (April 28, 1966): 241-247, and by Alvin Goldman, Brian Skyrms, and others. Criticisms of various alternative analyses of knowledge are given in Lehrer's earlier article, and the reader is referred to that article; such discussion will not be repeated here. The distinction between basic and nonbasic knowledge that is elaborated here was suggested by Arthur Danto in "Freedom and Forebearance," in *Freedom and Determinism* (New York: Random House, 1965), pp. 45-63.

believing that *h*, even though, unfortunately, he is mistaken.² Furthermore, the third condition does not entail that there is any statement or belief that justifies *S* in believing that *h*. The analysis, then, is in keeping with the characterization of basic knowledge given above. In basic knowledge, *S* is completely justified in believing that *h* even if it is not the case that there is any statement or belief that justifies his believing that *h*.

There are cases in which a person has some, perhaps mysterious, way of being right about matters of a certain sort with such consistency that philosophers and others have said that the person knows whereof he speaks. Consider, for example, the crystal-ball-gazing gypsy who is almost always right in his predictions of specific events. Peter Unger suggests a special case of this.³ His gypsy is always right, but has no evidence to this effect and, in fact, believes that he is usually wrong. With respect to each specific prediction, however, the gypsy impulsively believes it to be true (as indeed it is). Whether or not the predictive beliefs of the ordinary gypsy and Unger's gypsy are cases of knowledge depends, we contend, on whether they are cases of basic knowledge. This in turn depends on whether the gypsies are completely justified in their beliefs. It is plausible to suggest that these are cases of knowledge, but this is only because it is also plausible to think that the gypsies in question have some way of being right that completely justifies their prognostications. We neither affirm nor deny that these are cases of knowledge, but maintain that, if they are cases of knowledge, then they are cases of *basic* knowledge.

It is consistent with our analysis of knowledge to admit that a man knows something even though no statement constitutes evidence that completely justifies his believing it. Philosophers have suggested that certain memory and perceptual beliefs are completely justified in the absence of such evidential statements. We choose to remain agnostic with respect to any claim of this sort, but such proposals are not excluded by our analysis.

II

Not all knowledge that *p* is basic knowledge that *p*, because sometimes justifying evidence is essential. Consider the following analysis of nonbasic knowledge: (i) *h* is true, (ii) *S* believes that *h*, and

² Harman's criticism of Lehrer's earlier article rested on his interpreting Lehrer as saying that a person can be completely justified in believing something only if he does believe it. This interpretation leads to problems and is repudiated here.

³ "Experience and Factual Knowledge," this JOURNAL, LXIV, 5 (March 16, 1967): 152-173, esp. pp. 165-167; see also his "An Analysis of Factual Knowledge," *ibid.*, LXV, 6 (March 21, 1968): 157-170, esp. pp. 163-164.

(iii*) p completely justifies S in believing that h . In this analysis, p is that (statement) which makes S completely justified in believing that h . Note that (iii*), like (iii), does not entail (ii) or (i).

This analysis of nonbasic knowledge is, of course, defective. As Edmund Gettier has shown, there are examples in which some false statement p entails and hence completely justifies S in believing that h , and such that, though S correctly believes that h , his being correct is mostly a matter of luck.⁴ Consequently, S lacks knowledge, contrary to the above analysis. Other examples illustrate that the false statement which creates the difficulty need not entail h . Consider, for example, the case of the pyromaniac described by Skyrms.⁵ The pyromaniac has found that Sure-Fire matches have always ignited when struck. On the basis of this evidence, the pyromaniac is completely justified in believing that the match he now holds will ignite upon his striking it. However, unbeknownst to the pyromaniac, this match happens to contain impurities that raise its combustion temperature above that which can be produced by the friction. Imagine that a burst of Q-radiation ignites the match just as he strikes it. His belief that the match will ignite upon his striking it is true and completely justified by the evidence. But this is not a case of knowledge, because it is not the striking that will cause the match to ignite.

Roderick Chisholm has pointed out that justifications are defeasible.⁶ In the examples referred to above, there is some true statement that would defeat any justification of S for believing that h . In the case of the pyromaniac, his justification is defeated by the true statement that striking the match will not cause it to ignite. This defeats his justification for believing that the match will ignite upon his striking it.

Thus we propose the following analysis of nonbasic knowledge: S has nonbasic knowledge that h if and only if (i) h is true, (ii) S believes that h , and (iii) there is some statement p that completely justifies S in believing that h and no other statement defeats this justification. The question we must now answer is—what does it mean to say that a statement defeats a justification? Adopting a suggestion of Chisholm's, we might try the following: when p completely justifies S in believing that h , this justification is defeated by q if and only if (i) q is true, and (ii) the conjunction of p and q

⁴ "Is Justified True Belief Knowledge?", *Analysis*, xxxiii.6, 96 (June 1963): 121-123.

⁵ "The Explication of 'X knows that p ,'" this JOURNAL, LXIV, 12 (June 22, 1967): 373-389.

⁶ *Theory of Knowledge* (Englewood Cliffs, N.J.: Prentice-Hall, 1966), p. 48.

does not completely justify *S* in believing that *h*.⁷ This definition is strong enough to rule out the example of the pyromaniac as a case of knowledge. The statement that the striking of a match will *not* cause it to ignite, which is true, is such that when it is conjoined to any statement that completely justifies the pyromaniac in believing that the match will ignite, the resultant conjunction will fail to so justify him in that belief. Given this definition of defeasibility, the analysis of nonbasic knowledge would require that a man who has nonbasic knowledge that *h* must have some justification for his belief that is not defeated by any true statement.

However, this requirement is somewhat unrealistic. To see that the definition of defeasibility under consideration makes the analysis of nonbasic knowledge excessively restrictive, we need only notice that there can be true statements that are misleading. Suppose I see a man walk into the library and remove a book from the library by concealing it beneath his coat. Since I am sure the man is Tom Grabit, whom I have often seen before when he attended my classes, I report that I know that Tom Grabit has removed the book. However, suppose further that Mrs. Grabit, the mother of Tom, has averred that on the day in question Tom was not in the library, indeed, was thousands of miles away, and that Tom's identical twin brother, John Grabit, was in the library. Imagine, moreover, that I am entirely ignorant of the fact that Mrs. Grabit has said these things. The statement that she has said these things would defeat any justification I have for believing that Tom Grabit removed the book, according to our present definition of defeasibility. Thus, I could not be said to have nonbasic knowledge that Tom Grabit removed the book.

The preceding might seem acceptable until we finish the story by adding that Mrs. Grabit is a compulsive and pathological liar, that John Grabit is a fiction of her demented mind, and that Tom Grabit took the book as I believed. Once this is added, it should be apparent that I did know that Tom Grabit removed the book, and, since the knowledge must be nonbasic, I must have nonbasic knowledge of that fact. Consequently, the definition of defeasibility must be amended. The fact that Mrs. Grabit said what she did should not be allowed to defeat any justification I have for believing that Tom Grabit removed the book, because I neither entertained any beliefs concerning Mrs. Grabit nor would I have been justified in

⁷ Chisholm, "The Ethics of Requirement," *American Philosophical Quarterly*, 1, 2 (April 1964): 147-153. This definition of defeasibility would make our analysis of nonbasic knowledge very similar to one Harman derives from Lehrer's analysis and also one proposed by Marshall Swain in "The Analysis of Non-Basic Knowledge" (unpublished).

doing so. More specifically, my justification does not depend on my being completely justified in believing that Mrs. Grabit did *not* say the things in question.

To understand how the definition of defeasibility must be amended to deal with the preceding example, let us consider an example from the literature in which a justification deserves to be defeated. Suppose that I have excellent evidence that completely justifies my believing that a student in my class, Mr. Nogot, owns a Ford, the evidence consisting in my having seen him driving it, hearing him say he owns it, and so forth. Since Mr. Nogot is a student in my class who owns a Ford, someone in my class owns a Ford, and, consequently, I am completely justified in believing that someone in my class owns a Ford. Imagine that, contrary to the evidence, Mr. Nogot does not own a Ford, that I have been deceived, but that unknown to me Mr. Havit, who is also in my class, does own a Ford. Though I have a completely justified true belief, I do not know that someone in my class owns a Ford. The reason is that my sole justification for believing that someone in my class does own a Ford is and should be defeated by the true statement that Mr. Nogot does not own a Ford.

In the case of Tom Grabit, the true statement that Mrs. Grabit said Tom was not in the library and so forth, should not be allowed to defeat my justification for believing that Tom removed the book, whereas in the case of Mr. Nogot, the true statement that Mr. Nogot does not own a Ford, should defeat my justification for believing that someone in my class owns a Ford. Why should one true statement but not the other be allowed to defeat my justification? The answer is that in one case my justification depends on my being completely justified in believing the true statement to be false while in the other it does not. My justification for believing that Tom removed the book does not depend on my being completely justified in believing it to be false that Mrs. Grabit said Tom was not in the library and so forth. But my justification for believing that someone in my class owns a Ford does depend on my being completely justified in believing it to be false that Mr. Nogot does not own a Ford. Thus, a defeating statement must be one which, though true, is such that the subject is completely justified in believing it to be false.⁸

⁸ In Skyrms' example of the pyromaniac cited earlier, the defeating statement is not one which the pyromaniac need believe; Skyrms suggests that the pyromaniac neither believes nor disbelieves that striking the match will cause it to ignite. Nevertheless, the pyromaniac would be completely justified in believing that striking the Sure-Fire match will cause it to ignite. Hence the statement that striking the match will *not* cause it to light is defeating.

The following definition of defeasibility incorporates this proposal: when p completely justifies S in believing that h , this justification is defeated by q if and only if (i) q is true, (ii) S is completely justified in believing q to be false, and (iii) the conjunction of p and q does not completely justify S in believing that h .

This definition of defeasibility, though basically correct, requires one last modification to meet a technical problem. Suppose that there is some statement h of which S has nonbasic knowledge. Let us again consider the example in which I know that Tom Grabit removed the book. Now imagine that there is some true statement which is completely irrelevant to this knowledge and which I happen to be completely justified in believing to be false, for example, the statement that I was born in St. Paul. Since I am completely justified in believing it to be false that I was born in St. Paul, I am also completely justified in believing to be false the conjunctive statement that I was born in St. Paul and that q , whatever q is, because I am completely justified in believing any conjunction to be false if I am completely justified in believing a conjunct of it to be false. Therefore, I am completely justified in believing to be false the conjunctive statement that I was born in St. Paul and Mrs. Grabit said that Tom Grabit was not in the library and so forth. Moreover, this conjunctive statement is true, and is such that, when it is conjoined in turn to any evidential statement that justifies me in believing that Tom Grabit removed the book, the resultant extended conjunction will not completely justify me in believing that Tom Grabit removed the book. Hence, any such justification will be defeated.⁹ Once again, it turns out that I do not have nonbasic knowledge of the fact that Tom is the culprit.

In a logical nut, the problem is that the current definition of defeasibility reduces to the preceding one. Suppose there is a true statement q such that, for any p that completely justifies S in believing h , the conjunction of p and q does not completely justify me in believing that h . Moreover, suppose that I am not completely justified in believing q to be false, so that, given our current definition of defeasibility, q does not count as defeating. Nevertheless, if there is any true statement r , irrelevant to both p and q , which I am completely justified in believing to be false, then we can indirectly use q to defeat my justification for believing h . For I shall be completely justified in believing the conjunction of r and q to be false, though in fact it is true, because I am completely justified in believing r to be false. If the conjunction of q and p does not

⁹ A similar objection to Lehrer's earlier analysis is raised by Harman, p. 243.

completely justify me in believing that h , then, given the irrelevance of r , neither would the conjunction of r , q and p justify me in believing that h . Hence, my justifications for believing h would be defeated by the conjunction r and q on the current definition of defeasibility as surely as they were by q alone on the preceding definition.

The defect is not difficult to repair. Though S is completely justified in believing the conjunction of r and q to be false, one consequence of the conjunction, q , undermines my justification but is not something I am completely justified in believing to be false, while another consequence, r , is one that I am completely justified in believing to be false but is irrelevant to my justification. To return to our example, I am completely justified in believing to be false the conjunctive statement that I was born in St. Paul and that Mrs. Grabit said that Tom was not in the library and so forth. One consequence of this conjunction, that Mrs. Grabit said that Tom was not in the library and so forth, undermines my justification but is not something I am completely justified in believing to be false, while the other consequence, that I was born in St. Paul, is something I am completely justified in believing to be false but is irrelevant to my justification. The needed restriction is that those consequences of a defeating statement which undermine a justification must themselves be statements that the subject is completely justified in believing to be false.

We propose the following definition of defeasibility: if p completely justifies S in believing that h , then this justification is defeated by q if and only if (i) q is true, (ii) the conjunction of p and q does not completely justify S in believing that h , (iii) S is completely justified in believing q to be false, and (iv) if c is a logical consequence of q such that the conjunction of c and p does not completely justify S in believing that h , then S is completely justified in believing c to be false.

With this definition of defeasibility, we complete our analysis of nonbasic knowledge. We have defined nonbasic knowledge as true belief for which some statement provides a complete and undefeated justification. We previously defined basic knowledge as true belief for which there was complete justification that did not depend on any justifying statement. We define as knowledge anything that is either basic or nonbasic knowledge. Thus, S knows that h if and only if S has either basic or nonbasic knowledge that h . Having completed our analysis, we shall compare it with other goods in the epistemic marketplace to demonstrate the superiority of our ware.

III

The analysis offered above resembles two recent analyses formulated by Brian Skyrms and R. M. Chisholm. Both philosophers distinguish between basic and nonbasic knowledge, and both analyze knowledge in terms of justification. Moreover, these analyses are sufficiently restrictive so as to avoid yielding the result that a person has nonbasic knowledge when his justification is defeated by some false statement. However, we shall argue that both of these analyses are excessively restrictive and consequently lead to skeptical conclusions that are unwarranted.

Skyrms says that a man has nonbasic knowledge that p if and only if he has either derivative or nonderivative knowledge that p . He analyzes the latter two kinds of knowledge as follows:

Derivative Knowledge: X has derivative knowledge that p if and only if there is a statement ' e ' such that:

- (i) X knows that e
- (ii) X knows that ' e ' entails ' p '
- (iii) X believes that p on the basis of the knowledge referred to in (i) and (ii)

Nonderivative Knowledge: X has nonderivative knowledge that p if and only if there is a statement ' e ' such that:

- (i) X knows that e
- (ii) X knows that ' e ' is good evidence for ' p '
- (iii) X believes that p on the basis of the knowledge referred to in (i) and (ii)
- (iv) ' p ' is true
- (v) There is no statement ' q ' (other than ' p ') such that:
 - (a) X knows that ' e ' is good evidence of ' q '
 - (b) X knows that ' q ' entails ' p '
 - (c) X believes that ' p ' on the basis of the knowledge referred to in (a) and (b) (*op. cit.*, 381)

Later in his paper, Skyrms points out a defect in his analysis of nonderivative knowledge, namely, that the words, 'There is a statement ' e ' such that . . .' must be replaced by some such expression as 'There is some statement ' e ' consisting of the total evidence of X relevant to p such that . . .' or else the analysis will lead to trouble (387).

We shall now show why this analysis is unsatisfactory. According to Skyrms, a man who knows that a disjunction is true without knowing any specific disjunct to be true, has nonderivative knowledge of the disjunction (380). Indeed, his analysis of nonderivative knowledge is simply a generalization of his analysis of knowledge with respect to such disjunctions. But his analysis is overrestrictive

in the case of our knowledge of disjunctions. Suppose I know that a business acquaintance of mine, Mr. Romeo, arrived in Rochester from Atlanta on either one of two flights, either AA 107 or AA 204. My evidence is that these are the only two flights into Rochester from Atlanta, that Mr. Romeo telephoned earlier from Atlanta to say he would be arriving on one of these two flights, that he is now in Rochester, and that no other flight to Rochester or nearby would enable Mr. Romeo to be in Rochester at the present time. On the basis of this evidence, I may on Skyrms' analysis be said to have nonderivative knowledge that Mr. Romeo arrived on either AA 107 or AA 204. So far so good.

However, suppose that we add to my evidence that, when I meet Mr. Romeo at the airport shortly after the arrival of AA 204 (the later flight), he tells me that he just arrived on AA 204. By Skyrms' analysis I now *lack* nonderivative knowledge that Mr. Romeo arrived on either AA 107 or AA 204. The reason is that condition (v) of his analysis of nonderivative knowledge is no longer satisfied with respect to that disjunction. I now have good evidence that Mr. Romeo arrived on AA 204, and I believe that disjunction on the basis of my knowledge that this evidence is good evidence for the statement "Mr. Romeo arrived on AA 204" and this statement entails "Mr. Romeo arrived on either AA 107 and AA 204." Thus, there is a statement '*q*' that satisfies condition (a), (b), and (c) under (v) where '*p*' is the disjunction.

The consequence that I now lack nonderivative knowledge that Mr. Romeo arrived on either AA 107 or AA 204 would not be fatal if it could be argued that I have derivative knowledge of that disjunction because I know that Mr. Romeo arrived on AA 204. But there is an unmentioned twist of romance in our tale. In fact, Mr. Romeo arrived on the earlier flight, AA 107, and, having entertained his secret love, deceitfully told me he arrived on the later flight. Thus, I do not know that Mr. Romeo arrived on AA 204, because he did not so arrive. By Skyrms' analyses, I have neither derivative nor nonderivative knowledge that Mr. Romeo arrived on either AA 107 or AA 204, and, therefore, I lack nonbasic knowledge of that disjunction. So, as Skyrms would have it, I do not know that Mr. Romeo arrived on either of those flights. However, although there is much of interest that I do not know in this case, I surely do know, on the basis of my original evidence which I may yet brandish with epistemic righteousness, that Mr. Romeo must have arrived on either AA 107 or AA 204. He did so arrive, and my evidence completely justifies me in believing that he did, regardless of the fact that Mr. Romeo spoke with a crooked tongue. Since I

do have knowledge of the disjunction, Skyrms' analyses must be rejected.

Chisholm's analysis of knowledge is very similar to ours except for the condition intended to deal with situations in which, though a man has completely justified true belief, his justification is undermined by some false statement. In the sort of cases we have been considering, Chisholm's analysis requires, among other conditions, that if a person knows that h , then there is a proposition p such that p justifies h but p does not justify any false statement.¹⁰ However, it seems reasonable to suppose that every statement, whatever epistemic virtues it might have, completely justifies at least one false statement. This supposition is supported by the fact that justification in Chisholm's system need not be deductive justification. Any nondeductive justification may fail to be truth-preserving; that is, the conclusion may be false though the premise be true. Thus, though our analysis is in a number of ways indebted to Chisholm's proposals, the foregoing argument is our reason for concluding that Chisholm's analysis would lead to some form of skepticism, that is, to the conclusion that people do not know some things they would generally and reasonably be said to know.

IV

Having indicated our reasons for rejecting those analyses which are most similar to our own, we shall now turn to some analyses that differ from ours in more fundamental ways. Peter Unger has analyzed knowledge as follows: For any sentential value of p , (at a time t) a man knows that p if and only if (at t) it is not at all accidental that the man is right about its being the case that p .¹¹ Unger nowhere rules out the possibility that there are some cases in which it is not at all accidental that a man is right simply because he has justification for believing what he does. So it could be that any case that satisfies our conditions for knowledge would satisfy his as well. But there are cases that satisfy his analysis though they fail to satisfy ours.

Let us consider an example. A hologram, or laser photograph, when illuminated by laser light looks three-dimensional even with respect to parallax effects when the viewer shifts his position. Imagine that holography has been so perfected that a laser-illuminated hologram of an object can, under certain observational conditions, be indistinguishable from the real thing.¹² More particularly, suppose that a man, Mr. Promoter, seeking to demonstrate the remark-

¹⁰ Chisholm; see footnote at end of chap. 1, *Theory of Knowledge*, p. 23.

¹¹ "An Analysis of Factual Knowledge," p. 158.

¹² Cf. Alvin Goldman, "A Causal Theory of Knowing," this JOURNAL, LXIV, 12 (June 22, 1967): 357-372; p. 359.

able properties of laser photography, constructs a boxlike device which contains a vase, a laser photograph of the vase, and a laser source by which the photograph may be illuminated. The device is so constructed that Mr. Promoter by turning a knob may show a viewer the vase or the illuminated laser photograph of the vase, and the visual experience of the viewer when he sees the vase will be indistinguishable from his visual experience when confronted with the photograph. Of course, the very purpose of constructing the device is to arrange things so that people will be completely deceived by the photograph. Now suppose I walk up to the viewer, innocent as the fool who stoned the water to destroy his twin, and peer in at the illuminated photograph. Blissfully ignorant of the technical finesse being used to dupe me, I take what I see to be a vase. I believe that the box contains a vase. I am right, there is a vase in the box, and it is not at all accidental that I am right. For Mr. Promoter has constructed the device in such a way that, though I do not see the vase, I will believe quite correctly that there is one there. On Unger's analysis, I know that there is a vase in the box when I see the illuminated laser photograph.

However, it is perfectly apparent that I know nothing of the sort. Any justification I have for believing that there is a vase in the box is defeated by the fact that I do not see a vase in the box but merely a photograph of one. On our analysis it would follow that I do *not* know there is a vase in the box, and that result is the correct one.

Unger might object that it is to some extent accidental that I am right in thinking there is a vase in the box, because I might have had the same visual experiences even if there had been no vase in the box. Hence his analysis yields the same result as ours in this case. But this objection, if taken seriously, would lead us to reject Unger as a skeptic. To see why, imagine that, contrary to the preceding example, Mr. Promoter turns off his device for the day, leaving the knob set so that when I enter the room the vase is before my eyes. I could reach out and touch it if I wished, but good manners restrain me. Nevertheless, there is nothing between me and the vase; I see it and know that it is before my eyes in just the way that I see and know that countless other objects sit untouched before me. However, the statement—I might have had the same visual experience even if there had been no vase in the box—is true in this case as in the former one where I was deceived by the photograph. If this truth shows that my being right in the former case was to some extent accidental, then it would also show that my being right in the present case was to some extent accidental. Therefore, either Unger must agree that the truth of this statement fails to show that

my being right in the former case was accidental, in which case his analysis would yield the result that I know when in fact I am ignorant; or he must maintain that its truth shows that my being right in the present case was accidental, in which case his analysis yields the result that I am ignorant when in fact I know. Thus, his analysis is unsatisfactory.

Finally, we wish to consider another kind of theory suggested by Alvin Goldman. His analysis is as follows: *S* knows that *p* if and only if the fact that *p* is causally connected in an "appropriate" way with *S*'s believing that *p* (*op. cit.*, 369). We wish to assert, in opposition to Goldman, that the causal etiology of belief may be utterly irrelevant to the question of what a man knows. Consider yet a third round between Mr. Promoter and me. This time I imagine that I enter as in the first example, where the photograph is illuminated, and become completely and thoroughly convinced that there is a vase in the box. Now imagine that Mr. Promoter, amused with his easy success, tells me that I am quite right in thinking there is a vase in the box, but he then goes on to show me how the device is constructed, removing parts and lecturing about lasers from smirk-twisted lips. With respect to the etiology of my belief that there was a vase in the box, it is possible that my belief was fixed from the time I first looked at the photograph and, moreover, was so firmly and unequivocally fixed that the subsequent revelations neither altered nor reinforced it. This belief is to be *causally explained* by my mistakenly believing that I was seeing a vase when I first entered the room and by the facts about the illuminated laser photograph that caused that erroneous belief. There is no "appropriate" causal connection between the fact that there is a vase in the box and my belief that *p*; so, according to Goldman's analysis, I did not know that there was a vase in the box.

There is something to recommend this result. When I first looked into the device, I did not see the vase, and, consequently, I did not *then* know that there was a vase in the box then. However, after Mr. Promoter's revelations, when I do really see the vase, I do then know that there is a vase in the box. This is not due to any change in the causal etiology of my belief that there is a vase in the box. So, according to Goldman, I still do not know. But Goldman is wrong. I do subsequently know that there is a vase in the box, not because of any change in the causal etiology of the belief, but because I then have some justification for the belief that I formerly lacked. The justification consists of what I learned from Mr. Promoter's demonstration about the box and its contents. In short,

there is no reason to suppose that all new evidence that a man could appeal to in order to justify a belief changes the causal etiology of that belief. And such evidence may make the difference between true belief and knowledge.

v

We have contended that our analysis of knowledge in terms of undefeated justified true belief has various advantages over competing analyses. Unlike some of our competitors, we do not presuppose any one theory of justification rather than another. Since current theories of justification are highly controversial, we have employed a notion of justification that is consistent with diverse theories on this subject. By so doing, we hope to have presented a satisfactory analysis of knowledge without waiting for the development of an equally satisfactory theory of justification.

Moreover, the problems that confront a theory of justification can be formulated in terms of the locutions we have introduced in our analysis. For example, Chisholm has maintained that some statements are self-justifying, and, in our terminology this amounts to answering affirmatively the question whether it is ever the case that some statement *h* completely justifies a person in believing that *h*.¹³ Some philosophers have affirmed that all justification must be either inductive or deductive; others have denied this and affirmed that there are other forms of justification as well. In our terminology, this question may be formulated as the question whether, when a statement *p* completely justifies a person in believing *h*, the justified statement must be deduced or induced from the justifying statement or whether there are other alternatives. Finally, philosophers have disagreed about the kind of statement that may justify a man in believing something: whether those statements must be known, or whether they need not be, whether they must include all of a man's evidence, or whether they might exclude some of his evidence, and so forth. We have avoided dogmatically assuming one or the other of these alternatives.

Nevertheless, it may be found that only one theory of justification is suitable to supplement our analysis. Our claim is that, on any satisfactory theory of justification, some knowledge must be undefeated completely justified true belief, and the rest is basic.

KEITH LEHRER

THOMAS PAXSON, JR.

University of Rochester

¹³ R. M. Chisholm and others, *Philosophy* (Englewood Cliffs, N.J.: Prentice-Hall, 1964), pp. 263-277.