



Epistemic Dependence

John Hardwig

The Journal of Philosophy, Vol. 82, No. 7 (Jul., 1985), 335-349.

Stable URL:

<http://links.jstor.org/sici?sici=0022-362X%28198507%2982%3A7%3C335%3AED%3E2.0.CO%3B2-E>

The Journal of Philosophy is currently published by Journal of Philosophy, Inc..

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. 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/journals/jphil.html>.

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.

For more information on JSTOR contact jstor-info@umich.edu.

©2003 JSTOR

THE JOURNAL OF PHILOSOPHY

VOLUME LXXXII, NO. 7, JULY 1985

EPISTEMIC DEPENDENCE*

I find myself believing all sorts of things for which I do not possess evidence: that smoking cigarettes causes lung cancer, that my car keeps stalling because the carburetor needs to be rebuilt, that mass media threaten democracy, that slums cause emotional disorders, that my irregular heart beat is premature ventricular contraction, that students' grades are not correlated with success in the nonacademic world, that nuclear power plants are not safe (enough) . . . The list of things I believe, though I have no evidence for the truth of them, is, if not infinite, virtually endless. And I am finite. Though I can readily imagine what I would have to do to obtain the evidence that would support any one of my beliefs, I cannot imagine being able to do this for *all* of my beliefs. I believe too much; there is too much relevant evidence (much of it available only after extensive, specialized training); intellect is too small and life too short.

What are we as epistemologists to say about all these beliefs? If I, without the available evidence, nevertheless believe a proposition, are my belief and I in that belief necessarily irrational or nonrational? Is my belief then *mere* belief (Plato's right opinion)? If not, why not? Are there other good reasons for believing propositions, reasons which do not reduce to having evidence for the truth of those propositions? What would these reasons look like?

In this paper I want to consider the idea of intellectual authority, particularly that of experts. I want to explore the "logic" or epis-

* In a paper about epistemic dependence, it is fitting that I acknowledge my own debts. I have benefited by helpful comments and criticisms of earlier versions of this paper by William R. Carter, by members of the philosophy departments at the University of Tennessee and East Tennessee State University, and by Mary Read English. My dependence on William Bugg, Professor of Physics at the University of Tennessee, for discussion of a central example will become evident.

temic structure of an appeal to intellectual authority and the way in which such an appeal constitutes justification for believing and knowing. I have divided the paper into three parts. In the first, I argue that one can have good reasons for believing a proposition if one has good reasons to believe that *others* have good reasons to believe it and that, consequently, there is a kind of good reason for believing which does not constitute evidence for the truth of the proposition. In the second, I urge that because the layman is the epistemic inferior of the expert (in matters in which the expert is expert), rationality sometimes consists in refusing to think for oneself. In the third, I apply the results of these considerations to the concept of knowledge and argue that the expert-layman relationship is essential to the scientific and scholarly pursuit of knowledge.

If I am correct, appeals to epistemic authority are essentially ingredient in much of our knowledge. Appeals to the authority of experts often provide justification for claims to know, as well as grounding rational belief. At the same time, however, the epistemic superiority of the expert to the layman implies rational authority over the layman, undermining the intellectual autonomy of the individual and forcing a reexamination of our notion of rationality. The epistemic individualism implicit in many of our epistemologies is thus called into question, with important implications for how we understand knowledge and the knower, as well as for our conception of rationality.

1

Restricting ourselves—here and throughout the paper—to belief in and knowledge of propositions for which there is evidence, let us suppose that there are good reasons to believe a proposition—that *p*. What kinds of things can be good reasons to believe that *p*? The usual answer to this question is in terms of evidence, “evidence” being defined roughly as anything that counts toward establishing the truth of *p* (i.e., sound arguments as well as factual information). There is evidence, then, for the truth of *p*, but it does not follow that everyone has or even can have this evidence.

Suppose that person *A* has good reasons—evidence—for believing that *p*, but a second person, *B*, does not. In *this* sense *B* has no (or insufficient) reasons to believe that *p*. However, suppose also that *B* has good reasons to believe that *A* has good reasons to believe that *p*. Does *B* then, *ipso facto*, have good reasons to believe that *p*? If so, *B*'s belief is epistemically grounded in an appeal to the authority of *A* and *A*'s belief. And, if we accept this, we will be able to explain how *B*'s belief can be more than mere belief; how it can,

indeed, be rational belief; and how *B* can be rational in his belief that *p*. And our problems will be solved . . . or only starting.

Only starting because we are then faced with the prospect, not formerly considered by epistemologists, of a very odd kind of good reason for belief: a reason that does *not* constitute evidence for the truth of *p*. For *B*'s reasons for believing that *p* are not evidence for the truth of *p*. We can see this by noting two things. (1) Although *A*'s evidence counts toward establishing the truth of *p*, the case for *p* is not stronger after *B* discovers that *A* has this evidence than it was before *B* found out about *A* and *A*'s reasons. (2) The chain of appeals to authority must end somewhere, and, if the whole chain of appeals is to be epistemically sound, it must end with someone who possesses the necessary evidence, since truth claims cannot be established by an appeal to authority, nor by investigating what other people believe about them.¹

But *B* must have some good reasons to support his belief that *p*, or that belief would be *mere* belief (Plato's right opinion, again). *B* does have good reasons, all right; in fact, he has evidence. But his evidence doesn't count toward establishing the truth of *p*; it counts only toward establishing that *A* (unlike *B* himself) "knows what he's talking about" when he says that *p*. How can *B* have good reasons to believe that *A* has good reasons to believe that *p* when *B* does not himself have evidence that *p*? It's easy—*B* has good reasons to believe that *A* has conducted the inquiry necessary to have evidence for believing that *p*.

If the necessary inquiry is simple enough, *B*'s belief that *p* can be grounded in *A*'s despite the fact that we would not call *A* an expert. For example, if the service station attendant who checks my oil informs me that it's OK, I would believe him, but I would not call him an expert. However, the more epistemologically interesting cases are those in which expertise is involved—cases in which *B* has

¹ It might still seem that if *B* has good reasons to believe that *A* has good reasons to believe that *p*, then *B* has *evidence* that *p*. The dispute between me and someone inclined to press this objection would turn on delicate epistemological issues involved in clarifying the concept of evidence. But I would argue that *B* does *not* have evidence that *p*, urging, in addition to the arguments presented in the body of this paper, the following. (1) Evidence that *p* counts against evidence that not *p*. But consider a case of conflicting experts: *A* who has evidence that *p*, and *C*, who has evidence that not *p*. In such a case if *B* believes that *p* only because he believes that *A* has good reasons to believe that *p*, *B*'s reasons do not count against *C*'s; only *A*'s do. (2) It would be possible to construct cases in which *B* has good reason to believe that *A* has good reason to believe that *p* even where we would agree that there is *no* evidence for *p*. (More on such cases in sec. II below.) But regardless of how this dispute about "evidence" is to be resolved, I would observe that *B*'s reasons are logically dependent on *A*'s. Most of the rest of the points in this paper will follow if this is granted.

good reason to believe that *A* is an expert about whether or not *p*, in consequence of inquiry that has been sustained, prolonged, and systematic.²

The layman's appeal to the intellectual authority of the expert, his epistemic dependence on the expert, and his intellectual inferiority to the expert (in matters on which the expert is expert) are all expressed by the formula with which we have been working: *B* has good reasons to believe that *A* has good reasons to believe that *p*. But the layman's epistemic inferiority and dependence can be even more radical—in many such cases, extensive training and special competence may be necessary before *B* could conduct the necessary inquiry. And, lacking this training and competence, *B* may not be able to understand *A*'s reasons, or, even if he does understand them, he may not be able to appreciate why they are good reasons.

Michael Polanyi and Harry Prosch³ put the first part of this point dramatically, taking their examples from the physical sciences:

The popular conception of science says that science is a collection of observable facts that anybody can verify for himself. We have seen this is not true in the case of expert knowledge, like that needed in diagnosing a disease. Moreover, it is not true in the physical sciences. In the first place, for instance, a layman cannot possibly get hold of the equipment for testing a statement of fact in astronomy or in chemistry. Even supposing that he could somehow get the use of an observatory or a chemical laboratory, he would not know how to use the instruments he found there and might very well damage them beyond repair before he had ever made a single observation; and if he should succeed in carrying out an observation to check up on a statement of science and found a result that contradicted it, he could rightly assume that he

² I assume that we can all agree that there are experts, but I have not attempted in this paper to offer a precise definition of 'expert' or to delineate the range of possible expertise (beyond the introductory proviso that this paper is restricted to belief in and knowledge of propositions for which there is evidence). If the theses of this paper are correct, however, it will become crucial for epistemologists to argue about the definition of 'expert' and the range of actual and possible expertise.

But one observation about my use of 'expert' is in order: it does not presuppose or entail the truth of the expert's views. If one defines 'expert' in terms of the *truth* of his views (as Plato's Gorgias and Thrasymachus do), it is often impossible in principle to say who is an expert—even if one is an expert oneself!—since it is often impossible to say whose view is coincident with the truth. But I submit that it is not similarly impossible to say what constitutes sustained, relevant inquiry and to ascertain who is engaged in it (though there will sometimes be very real problems in making this judgment). And whenever sustained inquiry is both necessary for and efficacious with respect to determining whether or not *p*, the expert's views are less likely to be mistaken and likely to be less mistaken than an inexpert opinion. Thus, in my use of 'expert', the connection between truth and the views of the expert is not completely severed, though that connection is neither necessary nor simple.

³ *Meaning* (Chicago: University Press, 1977).

had made a mistake, as students do in a laboratory when they are learning to use its equipment (184/5).

Moreover, the training and resultant competence to conduct the required inquiry are often accessible only to those with certain talents and abilities. Consequently, *B* might not *ever* be able to obtain the evidence that supports his belief that *p*. If my own desperate and losing struggle with freshman calculus is a reliable indicator, I might *never* be able to obtain the evidence for my belief that relativity physics is correct, no matter how much time and effort I devoted to the enterprise. I may simply lack the mathematical ability to possess that evidence.

But extensive training and special competence may be necessary before one can assess or even understand the expert's reasons for believing that *p*. Although I might be able to understand studies about the impact of mass media on voters, I am not competent to assess the merits of those studies, unversed as I am in the issues surrounding social science research methods. And, lacking the requisite mathematical training and ability, I cannot even read the books and articles that support my belief that relativity physics is correct.

If, then, layman *B* (1) has not performed the inquiry that would provide the evidence for his belief that *p*, (2) is not competent, and perhaps could not even become competent, to perform that inquiry, (3) is not able to assess the merits of the evidence provided by expert *A*'s inquiry, and (4) may not even be able to understand the evidence and how it supports *A*'s belief that *p*, can *B* nonetheless have good reasons to believe that *A* has good reasons to believe that *p*? I think he can. If so, should we conclude that *B*'s belief that *p* is rationally justified? I think we should, acknowledging that *B*'s belief stands on better epistemic ground than other beliefs which we would call simply irrational or nonrational.

Many epistemologists may be tempted to reject this conclusion because it is so divergent from our received view about the nature of rational belief. But I think we *must* say that *B*'s belief is rationally justified—even if he does not know or understand what *A*'s reasons are—if we do not wish to be forced to conclude that a very large percentage of beliefs in any complex culture are simply and unavoidably irrational or nonrational. For, in such cultures, more is known that is relevant to the truth of one's beliefs than anyone could know by himself. And surely it would be paradoxical for epistemologists to maintain that the more that is known in a culture, the less rational the beliefs of individuals in that culture.

II

Nevertheless, acceptance of epistemological individualism dies

hard. It may well resurface in the guise of a suggestion about the appropriate stance of a responsible and rational layman in relation to the expert. If I am not presently in a position to know what the expert's good reasons for believing that p are or to understand why these are good reasons, I am obviously in no position to check the accuracy of what he tells me. What stance should I then take? A plausible and tempting suggestion is that, if I think I have the required ability, I should become informed so that I can assess the reliability of the expert's reports and thus escape my dependence on him and regain my intellectual autonomy.

The idea behind this suggestion lies at the heart of one model of what it means to be an intellectually responsible and rational person, a model which is nicely captured by Kant's statement that one of the three basic rules or maxims for avoiding error in thinking is to "think *for oneself*."⁴ This is, I think, an extremely pervasive model of rationality—it underlies Descartes's methodological doubt; it is implicit in most epistemologies; it colors the way we have thought about knowledge. On this view, the very core of rationality consists in preserving and adhering to one's own independent judgment; for how can one be sure one is being informed, not misinformed, if one suspends judgment?

But I submit that this model provides us with a romantic ideal which is thoroughly unrealistic and which, in practice, results in less rational belief and judgment. I could, indeed, escape epistemic dependence on *some* experts; perhaps, if I am talented enough, I could escape reliance on any given expert. I can and do choose where to establish my intellectual autonomy. But if I were to pursue epistemic autonomy across the board, I would succeed only in holding relatively uninformed, unreliable, crude, untested, and therefore *irrational* beliefs. If I would be rational, I can never avoid some epistemic dependence on experts, owing to the fact that I believe more than I can become fully informed about.

Once more, then: if I am not in a position to know what the expert's good reasons for believing that p are and why these are good reasons, what stance should I take in relation to the expert? If I do not know these things, I am also in no position to determine whether the person really is an expert. By asking the right questions, I might be able to spot a few quacks, phonies, or incompetents, but only the more obvious ones. For example, I may suspect that my doctor is incompetent, but generally I would have to know what doctors know in order to confirm or dispel my suspicion.

⁴ *Critique of Judgment*, J. H. Bernard, tr. (New York: Hafner, 1951), p. 136; Kant's emphasis. Kant repeats this statement in the *Anthropologie*, p. 118, and in the *Logik*, p. 371, both Cassirer editions (Berlin, 1932).

Thus, we must face the implications of the fact that laymen do not fully understand what constitutes good reasons in the domain of expert opinion.

Granted, I can check on a given expert and perhaps obtain a ranking of various experts⁵—by relying on other experts. If my doctor says that I should see a cardiologist, I can ask him and other physicians in the community about the local cardiologists. Or if I want to know the effects of mass media on voters, I can go to the political science department and ask who has done the best work in this area and whether there has been significant criticism of it. This checking and ranking of experts can be expressed by extending our formula and its implied chain of authority: *B* has good reasons for believing that *C* has good reasons for believing that *A* has good reasons for believing that *p*. However, by appealing to such a hierarchy of experts I have not regained my epistemic autonomy by avoiding reliance on experts—I have only extended and refined this reliance. Nor could I regain my epistemic autonomy in all cases without believing on the basis of relatively crude and untested reasons.

Granted also, if I do not know and have no way of finding out who the experts are, I will have no way to appeal to the chain of authority. I will then not know who has good reasons to believe that *p*, to whom to defer, or whose opinion (if any) will give me good reasons for believing that *p*. This sometimes happens, and, when it does, rational deference becomes impossible. But generally I can find someone whose opinion is more informed than mine and who can refer me to someone who is knowledgeable about whether or not *p*. And even if a layman, because of his relative inability to discriminate among experts, ends up appealing to a lesser instead of a greater expert, the lesser expert's opinion will still be better than the layman's.⁶

⁵In a series of recent articles, Keith Lehrer has explored the issues concerning the ranking of experts and the opinions of various experts and, consequently, the way to handle the problem of disagreement among experts, all with much more rigor and precision than I can muster here. Cf., e.g., "Social Information," *Monist*, I.X, 4 (October, 1977): 473-487, and also the articles Lehrer refers to in his footnotes to this article.

⁶Of course, a more detailed account of the whole issue of identifying relevant experts would have to distinguish among (1) *B* merely believing that *A* has good reasons to believe that *p*, (2) *B* having some reason to believe that *A* has good reasons to believe that *p*, and (3) *B* having good reasons to believe that *A* has good reasons to believe that *p*. And none of this resolves the often excruciating practical problem of identifying who the real or best experts are—e.g., what is the patient faced with conflicting medical opinions to do? But these are logically posterior issues and problems; the argument of this section of the paper is that in any case he should *not* make his own diagnosis, nor even read up some about his problem and then make his own diagnosis.

In terms of our formula, then, *B* could believe that *p* either because *B* has good reasons to believe that *A* has good reasons to believe that *p*, or because *B* has good reasons to believe that *C* has good reasons to believe that *A* has good reasons to believe that *p*. But, in either case, *B* cannot have sufficiently good reasons not to believe that *p* or to believe that not *p*. In other words, the layman cannot rationally refuse to defer to the views of the expert or experts he acknowledges. This does not mean that *B* can never successfully raise a devastating objection to believing that *p* or imagine an alternative to believing that *p*, but it does mean that only someone with *A*'s expertise can make an accurate assessment of the value and validity of the objection or alternative. Under cross-examination by the layman, the expert may admit the cogency of a given point, but *he* (and his fellow experts) must judge whether it is cogent and germane, since they are the only ones who fully understand what is involved in the methods, techniques, premises, and bases of the expert's training and inquiry and how these affect the resultant belief.

The layman can, in other words, propose criticisms and alternatives, but rationally he must allow the expert to dispose of them, for in a conversation with an expert (as opposed to a dialogue among equals⁷), the final court of rational appeal belongs solely to one party, by virtue of that party's greater competence for and commitment to inquiry into the relevant subject matter. The rational layman recognizes that his own judgment, uninformed by training and inquiry as it is, is *rationally inferior* to that of the expert (and the community of experts for whom the expert usually speaks) and consequently can always be *rationally* overruled. Recognizing that the highest court of rational appeal lies outside of himself, the layman may simply have to accept the fact that his objection is not a good one, even though it still seems good to him.

There are, of course, a whole series of ad hominem arguments that permit a layman rationally to refuse to defer to the expert's opinion. The layman can assert that the expert is not a disinterested, neutral witness; that his interest in the outcome of the discussion prejudices his testimony. Or that he is not operating in good faith—that he is lying, for example, or refusing to acknowledge a mistake in his views because to do so would tend to undermine his claim to special competence. Or that he is covering for his peers or knuckling under to social pressure from others in his field, etc., etc. Such ad hominem arguments are not always fallacious, and they sometimes do ground the

⁷I have attempted to explicate the logic of dialogue among presumed epistemic equals in the area of moral reasoning in my article "The Achievement of Moral Rationality," *Philosophy and Rhetoric*, vi, 3 (Summer 1973): 171-185.

rational refusal to defer to the statements of experts. But one interesting feature of such ad hominem is that they seem and perhaps are much more admissible, important, and damning in a layman's discussions with experts than they are in dialogues among peers. It doesn't matter so much if one's peers are biased or operating in bad faith; they will be found out. The merits of their arguments can be tested and evaluated rather than just accepted.

With the exception—often an important exception—of such ad hominem, I see no way to avoid the conclusion I have proposed above: that the rational layman will recognize that, in matters about which there is good reason to believe that there is expert opinion, he ought (methodologically) not to make up his own mind. His stance on these matters will—if he is rational—usually be rational deference to the epistemic authority of the expert.

If it is objected that, in cases of divided expert opinion, the layman will have no method for deciding whether or not to believe that p , this is granted.⁸ But in such cases the rational layman, recognizing that his own relatively casual and crude inquiry is not competent to resolve issues that even the sustained inquiry of experts cannot resolve, will also recognize that he is confronted with a situation in which he must either suspend belief or—if this is impossible or undesirable—arrive at belief on some admittedly nonrational basis. And if it is objected that layman B can have good reasons to believe that p even when p is false and even when expert A does not have good reasons to believe that p , this is also granted. For B will sometimes be misled by phony or mistaken claims to expertise, despite a careful attempt to ascertain that A is indeed an expert about p ; and, moreover, there is simply no guarantee that the views of even the best present experts are coincident with the final truth.

The conclusion that it is sometimes *irrational* to think for oneself—that *rationality* sometimes consists in deferring to epistemic authority and, consequently, in passively and uncritically accepting what we are given to believe—will strike those wedded to epistemic individualism as odd and unacceptable, for it undermines their paradigm of rationality. To others, it may seem too obvious for such belaboring. But in either case, I submit, we should recast our epistemologies and our accounts of rationality to make them congruent with this important fact of modern life.

⁸ If it is possible to rank experts in the ways Lehrer (*op. cit.*) has explored or in some other way, the layman can of course resolve the dilemma posed by divided expert opinion by deferring to the *best* expert opinion. However, there will still be cases in which even the best experts will disagree.

III

Although the preceding discussion is obviously relevant to the big word of epistemology—*know*—I have so far astutely avoided using it. But the relevance of the discussion is clear, given the standard analysis of “*A* knows that *p*” in terms of (1) *A* believes that *p*, (2) *A* has good reasons to believe that *p*, and (3) that *p* is true. The third condition is standardly taken to be the kicker, and it threatens to render the whole analysis inapplicable to knowledge; for *A* can have good reasons to believe that *p* even if *p* is false and *B* can have good reasons to believe that *A* has good reasons to believe that *p*—again, even if *p* is false. However, this third condition is not my primary concern, for I would argue for a fallibilist conception of knowledge.

I wish, rather, to focus our attention on the second, more neglected condition in the above analysis of “*A* knows that *p*.” It seems plausible that both *A* and *B* must have better or more complete reasons in order to *know* that *p* than are necessary merely to have good reasons to *believe* that *p*; for some beliefs, though rational, would not be well founded enough to qualify as knowledge (even on a fallibilist conception of knowledge). Thus it seems reasonable to hold that there is a progression from (1) believing that *p* (mere belief or right opinion), to (2) having good reasons to believe that *p* (rational belief), to (3) knowing that *p*.

What happens, then, if we substitute ‘know’ for ‘has good reasons to believe’ in our formula: i.e., *B* knows that *A* knows that *p*? Is it possible to *know* vicariously, as it were, or must *knowers* (as opposed to mere rational believers) stand on their own epistemic feet? I argued above that *B* can have good reasons to believe that *p* without having direct reasons or evidence for *p*. Is the same true for knowing? Or must *B* know that *p* before he can know that *A* knows that *p*, thus precluding an appeal to *A*’s knowledge as the basis and justification of his own claim to know? In other words, recalling the earlier distinction between having evidence that *p* and another kind of good reason to believe that *p*, must *B* possess the evidence for the truth of *p* in order to *know* that *p*? Or can knowledge, as well as rational belief, be based on an appeal to epistemic authority?

Suppose someone tells me something that is true without giving me evidence for its truth. Perhaps *A* tells me that laetrile does not cure cancer without giving me the studies that prove this, much less the concrete data on which those studies are based. But suppose I have good reasons to believe that *A* is an authority in the field of cancer research and so I believe what he tells me. Do I then *know*

that laetrile does not cure cancer, or have I achieved something much less than knowledge (perhaps only right opinion or rational belief)? If I then know, it is possible for one to know that p without possessing evidence for the truth of p . But that seems paradoxical or counterintuitive; for, in the cases we are considering, evidence is relevant to establishing knowledge, but we are asking whether it is possible to have this knowledge without the relevant evidence.

Even more paradoxical is the idea that B can know that p even though he doesn't understand that p . Suppose an eminent authority in particle physics tells me that a quark is a fundamental particle, and suppose this is true. But I don't even understand what that means, because I have no notion of what a quark is or what counts as a fundamental particle. However, I check up on the physicist, and, as a result, I *know* that he has unsurpassed credentials. Could I then be said to *know* that a quark is a fundamental particle, though I don't even understand what I know?

To sum up: should we say that B can (1) know that p by knowing that A knows that p , and (2) know this without first knowing that p ? Should we say this even if it implies that B can know that p without having evidence for p and perhaps without even understanding p ? Instead of attempting to answer these questions directly, I will argue that much of what we want to count as knowledge rests on the epistemic structure expressed by the formula, B knows that A knows that p .⁹ I will then offer two conclusions and leave it to the reader to decide which is more epistemologically palatable.

Scientists, researchers, and scholars are, sometimes at least, knowers, and all of these knowers stand on each other's shoulders in the way expressed by the formula: B knows that A knows that p . These knowers could not do their work without presupposing the validity of many other inquiries which they cannot (for reasons of competence as well as time) validate for themselves. Scientists, for example, simply do not repeat the experiments of other scientists unless the experiment is important and something seems fishy about it. It would, moreover, be impossible for anyone to get to the research front in, say, physics or psychology, if he relied only on the results of his own inquiry or insisted on assessing for himself the evidence behind all the beliefs he accepts in his field. Thus, if scientists, researchers, and scholars are knowers, the layman-expert

⁹ This strategy for approaching these issues will mean, of course, that it remains open to a courageous enough epistemologist to avoid my conclusion by embracing the view that the achievements of scientists, researchers, and scholars are not and could not be *knowledge* whenever these achievements are based on cooperative methodologies. This option does not seem very attractive to me, to say the least.

relationship is also present *within* the structure of knowledge, and the expert is an expert partially because he so often takes the role of the layman *within his own field*.

Moreover, research in many fields is increasingly done by teams rather than individuals. For example, it is not uncommon for the title of an article reporting experimental results in particle physics to look like this:

VOLUME 51, NUMBER 3

PHYSICAL REVIEW LETTERS

18 JULY 1983

Charm Photoproduction Cross Section at 20 GeV

K. Abe, T. C. Bacon, J. Ballam, L. Berny, A. V. Bevan, H. H. Bingham, J. E. Brau, K. Braune, D. Brick, W. M. Bugg, J. Butler, W. Cameron, J. T. Carroll, C. V. Cautis, J. S. Chima, H. O. Cohn, D. C. Colley, G. T. Condo, S. Dado, R. Diamond, P. J. Dornan, R. Erickson, T. Fieguth, R. C. Field, L. Fortney, B. Franek, N. Fujiwara, R. Gearhart, T. Glanzman, J. J. Goldberg, G. P. Gopal, A. T. Goshaw, E. S. Hafen, V. Hagopian, G. Hall, E. R. Hancock, T. Handler, H. J. Hargis, E. L. Hart, P. Haridas, K. Hasegawa, T. Hayashino, D. Q. Huang, ⁽¹⁾ R. I. Hulsizer, S. Isaacson, M. Jobs, G. E. Kalmus, D. P. Kelsey, J. Kent, T. Kitagaki, J. Lannutti, A. Levy, P. W. Lucas, M. MacDermott, W. A. Mann, T. Maruyama, R. Merenyi, R. Milburn, C. Milstene, K. C. Moffett, J. J. Murray, A. Napier, S. Noguchi, F. Ochiai, S. O'Neale, A. P. T. Palounek, I. A. Pless, M. Rabin, P. Rankin, W. J. Robertson, A. H. Rogers, E. Ronat, H. Rudnicka, T. Sato, J. Schneps, S. J. Sewell, J. Shank, A. M. Shapiro, C. K. Sinclair, R. Sugahara, A. Suzuki, K. Takahashi, K. Tamai, S. Tanaka, S. Tether, H. B. Wald, W. D. Walker, M. Widgoff, C. G. Wilkins, S. Wolbers, C. A. Woods, Y. Wu, A. Yamaguchi, R. K. Yamamoto, S. Yamashita, G. Yekutieli, Y. Yoshimura, G. P. Yost, and H. Yuta

Birmingham University, Birmingham B15 2TT, England, and Brown University, Providence, Rhode Island 02912, and Duke University, Durham, North Carolina 27706, and Florida State University, Tallahassee, Florida 32306, and Imperial College, London SW7 2BZ, England, and National Laboratory for High Energy Physics (KEK), Ohno-machi, Tsukuba-gun, Ibaraki 305, Japan, and Oak Ridge National Laboratory, Oak Ridge, Tennessee 37830, and Rutherford Appleton Laboratory, Didcot, Oxon OX11 0QX, England, and Stanford Linear Accelerator Center, Stanford University, Stanford, California 94305, and Technion-Israel Institute of Technology, Haifa 32000, Israel, and Tohoku University, Sendai 980, Japan, and Tufts University, Medford, Massachusetts 02155, and University of California, Berkeley, California 94720, and University of Tel Aviv, Tel Aviv, Israel, and University of Tennessee, Knoxville, Tennessee 37916, and Weizmann Institute, Rehovot, Israel

(Stanford Linear Accelerator Center Hybrid Facility Photon Collaboration)

(Received 2 May 1983)

Forty-seven charm events have been observed in an exposure of the SLAC Hybrid Facility bubble chamber to a 20-GeV backward-scattered laser beam. Thirty-seven events survive all the necessary cuts imposed. Based on this number the total charm cross section is calculated to be $63 \pm \frac{11}{2}$ nb.

PACS numbers: 13.60.Le, 13.60.Rj

In this Letter we present results on the charm photoproduction cross section in an experiment using the SLAC Hybrid Facility. Results on lifetimes of charmed particles based on part of the data were published earlier.¹

The SLAC 1-m hydrogen bubble chamber was exposed to a 20-GeV photon beam produced by Compton scattering of laser light by the 30-GeV electron beam. It was collimated to 3 mm in diameter. The photon beam energy spectrum is shown in Fig. 1. It peaks at 20 GeV with a full width at half maximum of 2 GeV. Most of the data were taken at photon intensities of 20–30 γ /pulse. In order to detect decays of charmed particles, a fourth camera with high-resolution optics having a resolution of 55 μ m over a depth of field ± 6 mm was used. The cameras were triggered either on the passage of a charged particle through three multiwire proportional chambers and pointing back to the fiducial volume of

the bubble chamber or on a sufficient energy deposition in an array of lead-glass blocks. Particle identification was provided by ionization measurements in the bubble chamber and light detection in two large-aperture Cherenkov counters. More details of the experimental setup and trigger are given in Ref. 1.

The results presented here are based on 270 000 hadronic interactions found in a restricted fiducial volume. All hadronic events were closely examined for the decays of short-lived particles within 1 cm of the production vertex. When such a decay was found, the following cuts were applied to ensure that the decays which survived were genuine charm decays: (a) Decays with less than two charged products were rejected. (b) Two-prong decays consistent with either photon conversions or strange-particle hypotheses were rejected. To eliminate K^0 decays, the two-body (assumed to be $\pi\pi$) invariant

William Bugg, Professor of Physics at the University of Tennessee, Knoxville, and a participant in this experiment, explained how such an experiment is done. This experiment, which recorded charm events and measured the lifespan of the charmed particles, was one of a series of experiments costing perhaps \$10 million. After it was funded, about 50 man/years were spent making the needed equipment and the necessary improvements in the Stanford Linear Accelerator. Then approximately 50 physicists worked perhaps 50 man/years collecting the data for the experiment. When the data were in, the experimenters divided into five geographic groups to analyze the data, a process which involved looking at 2½ million pictures, making measurements on 300,000 interesting events, and running the results through computers in order to isolate and measure 47 charm events. The “West Coast group” that analyzed about a third of the data included 40 physicists and technicians who spent about 60 man/years on their analysis.

Obviously, no one person could have done this experiment—in fact, Bugg reports that no one university or national laboratory could have done it—and many of the authors of an article like this will not even know how a given number in the article was arrived at.¹⁰ Furthermore, even if one person could know enough and live long enough to do such an experiment, there would be absolutely no point in his attempting to do so, for his results would have become obsolete long before he completed the experiment. Although Bugg expresses confidence that the team’s measurement of the lifespan of charmed particles is a good one, he estimates that within three years some other group will have come up with another technique that will give considerably better results. He consequently expects that within five years the paper will no longer be of general interest.

Finally, Bugg notes that the article’s 99 authors represent different specializations with particle physics, but all are experimentalists, so none would be able to undertake the theoretical revisions which might be required as a result of this experiment and which provide a large part of the rationale for doing it. On the other hand, most

¹⁰ Of course, only a few people actually write the article, but it does not follow that these people are masterminds for the whole procedure or that they completely understand the experiment and the analysis of the data. According to Bugg, although a few persons—“the persons most actively involved in working on the data and who therefore understand most about it”—wrote up the experiment (this article is 3½ journal pages long), they really only prepared a draft for revisions and corrections by the other authors. The team then met to argue substantive points about the techniques for analyzing the data and how the article should be presented to best enable other physicists to understand it.

theoreticians would not be competent to conduct the experiment—and neither the experimentalists nor the theoreticians are competent to design, build, and maintain the equipment without which the experiment could not be run at all.

Obviously, this is an extreme example, though not all that extreme in the realm of particle physics.¹¹ However, we can see how dependence on other experts pervades any complex field of research when we recognize that most footnotes that cite references are appeals to authority. And when these footnotes are used to establish premises for the study, they involve the author in layman-expert relationships even within his own pursuit of knowledge. Moreover, the horror that sweeps through the scientific community when a fraudulent researcher is uncovered is instructive, for what is at stake is not only public confidence. Rather, each researcher is forced to acknowledge the extent to which his own work rests on the work of others—work which he has not and could not (if only for reasons of time and expense) verify for himself.

Thus in very many cases *within* the pursuit of knowledge, there is clearly a complex network of appeals to the authority of various experts, and the resulting knowledge could not have been achieved by any one person. We then have something like the following:

- A* knows that *m*.
- B* knows that *n*.
- C* knows (1) that *A* knows that *m*, and (2) that if *m*, then *o*.
- D* knows (1) that *B* knows that *n*, (2) that *C* knows that *o*, and
(3) that if *n* and *o*, then *p*.
- E* knows that *D* knows that *p*.

Suppose that this is the only way to know that *p* and, moreover, that no one who “knows” that *p* knows that *m*, *n*, and *o* except by knowing that others know them. Does *D* or *E* know that *p*? Does anyone know that *p*? Is that *p* known?

Unless we maintain that most of our scientific research and scholarship could *never*, because of the cooperative methodology of the enterprise, result in knowledge, I submit that we must say that *p* is known in cases like this. But if *D* or *E* knows that *p*, we must also say that someone can know “vicariously”—i.e., without possessing the evidence for the truth of what he knows, perhaps without even fully understanding what he knows. And this conclusion would

¹¹ Of the 42 articles on elementary particles and fields published by *Physical Review Letters* in the three months from April 25 to July 18, 1983, 11 listed more than 10 authors, 9 listed more than 20 authors, and 5 more than 40 authors. In the same period only 5 articles were by single authors.

require dramatic changes in our analysis of what knowledge must be.

If the conclusion is unpalatable, another is possible. Perhaps that p is known, not by any one person, but by the *community* composed of A , B , C , D , and E . Perhaps D and E are not entitled to say, "I know that p ," but only, "We know that p ." This community is not reducible to a class of individuals, for no one individual and no one individually knows that p . If we take this tack, we could retain the idea that the knower must understand and have evidence for the truth of what he knows, but in doing so we deny that the knower is always an individual or even a class of individuals. This alternative may well point to part of what Peirce may have had in mind when he claimed that the *community* of inquirers is the primary knower and that individual knowledge is derivative.

The latter conclusion may be the more epistemologically palatable; for it enables us to save the old and important idea that *knowing* a proposition requires understanding the proposition and possessing the relevant evidence for its truth. But it will not be very comfortable for those who have a taste for desert landscapes, intellectual autonomy, or epistemic individualism; for it undermines the methodological individualism that is implicit in most epistemology. I believe that it is also deeply disturbing because it reveals the extent to which even our rationality rests on trust and because it threatens some of our most cherished values—individual autonomy and responsibility, equality and democracy. But that is a story for another occasion.

Thus if the arguments of this paper are accepted, some very basic changes in our epistemologies are required. We must recast our conception of what it means for beliefs and persons to be rational. We must also either agree that one can know without possessing the supporting evidence or accept the idea that there is knowledge that is known by the community, not by any individual knower.

JOHN HARDWIG

East Tennessee State University