

Knowledge as Evidence

TIMOTHY WILLIAMSON

It is argued that a subject's evidence consists of all and only the propositions that the subject knows.

I

Tradition has it that the main problems of philosophy include the nature of knowledge. But, in recent decades, questions of knowledge seem to have been marginalized by questions of justification. Thus, according to Crispin Wright,

... knowledge is not really the proper central concern of epistemologico-sceptical enquiry. [...] We can live with the concession that we do not, strictly, *know* some of the things we believed ourselves to know, provided we can retain the thought that we are fully justified in accepting them. (1991, p. 88; Wright's italics)

Similarly, John Earman argues that accounts of knowledge are irrelevant to the philosophy of science, because in it

... the main concern is rarely whether or not a scientist "knows" that some theory is true but rather whether or not she is justified in believing it. (1993, p. 37)¹

Once Gettier showed in 1963 that justified true belief is insufficient for knowledge, and therefore that knowledge is unnecessary for justified true belief, it became natural to ask: if you can have justified true beliefs, why bother with knowledge?²

There is a lacuna in the case for the unimportance of knowledge. Grant, for the sake of argument, that knowledge is important only if it is some-

¹ Earman is discussing externalist accounts of knowledge, but the quoted comment would clearly apply to internalist accounts too. Earman's further point, that "because science is a community enterprise the only forms of justification that are scientifically relevant are those which are stateable and open to public scrutiny", may be most relevant to externalist accounts. See also Craig (1990, p. 272). For the contrary view that scepticism about knowledge entails scepticism about rationality and justification see Unger (1975, pp. 197–249).

² Kaplan (1985) argues along similar lines.

how essential to the justification of belief.³ Although it has been shown that *what is justified* need not be knowledge, even when it is true, it has not been shown that *what justifies* need not be knowledge. Only one end of the justification relation has been separated from knowledge. Suppose that knowledge, and only knowledge, justifies belief. That is, in any possible situation in which S believes a proposition p , that belief is justified, if at all, by propositions q_1, \dots, q_n (usually other than p) which S knows. On that supposition, if justified belief is central to epistemologico-sceptical enquiry and the philosophy of science, then so too is knowledge. Now assume further that what justifies belief is *evidence* (I briefly discuss this assumption in §VIII). Then the supposition just made is equivalent to the principle that knowledge, and only knowledge, constitutes evidence. This paper defends that principle; it equates S's evidence with S's knowledge, for every individual or community S in any possible situation (throughout, "knowledge" means propositional knowledge).⁴ Call this equation $E = K$.⁵

The proposed account uses the concept of knowledge in partial elucidation of the concepts of evidence and justification. It may therefore seem to get things back to front. For although knowledge is more than justified true belief, many philosophers still expect to use concepts such as *evidence* and *justification* in a more complex explanation of the concept of knowledge; it would then be circular to use the latter to explain the former. Others prefer to use concepts of a different kind, such as *causation* or *reliability*, to explain the concept of knowledge; but even they are likely to regard the concept of knowledge as so much in need of explanation itself that its pre-theoretic use would lack explanatory value.

In response, the first point to note is that the equation $E = K$ could be true without being knowable a priori. It is a universal generalization over all metaphysically possible situations, and therefore necessarily true, if true at all; but there is no presumption that a necessary truth is knowable a priori.⁶ $E = K$ equates the extensions of the concepts of knowledge and evidence in any possible situation; that is enough to make it an informa-

³ Wright speaks of justified acceptance rather than belief; although Earman speaks of belief, some philosophers of science contrast it with acceptance and regard the latter as a more appropriate attitude towards scientific theories. The distinction is not important here.

⁴ The communal case is needed. Science depends on public evidence, which is neither the union nor the intersection of the evidence of each scientist. "It is known in S" and "We know" (which is not synonymous with "Some/many/most/all of us know") express the corresponding kind of knowledge.

⁵ The principle is stated, and applied to the problem of vagueness, in Williamson (1994, pp. 245–7).

⁶ The text assumes the S4 principle that a necessary truth is necessarily necessary. Nothing in the rest of the paper hangs on this.

tive thesis. It does not equate the concepts themselves; nor is it to be read as offering a conceptual analysis of either *evidence* or *knowledge*, or as making one concept prior to the other in any sense. Of course, in offering *arguments* of a broadly a priori kind for $E = K$, as I do, I commit myself to at least its a priori plausibility; in the best case for those arguments, they would provide a priori knowledge of $E = K$. However, even if the concepts are equivalent a priori, it does not follow that one is prior to the other. In particular, we still lack good reason to think that the concept of knowledge must or can be analysed in terms of more basic concepts, whether these are supposed to be the concepts of justification, evidence or anything else. The equation “knowledge = true belief + X ” need not have a non-circular solution—one which specifies X without using the concept of knowledge. It might be like the equation “red = coloured + X ”, in which one cannot specify X without using the concept of red. Not all our concepts can be broken down into more basic concepts, on pain of infinite regress. After Gettier, massive industry went into the production of such analyses of knowledge; their history is one of failure. It constitutes good inductive evidence that no such analysis provides even necessary and sufficient conditions for knowledge. The bad record of attempts to provide non-circular necessary and sufficient conditions for other philosophically interesting concepts, such as *meaning* and *causation*, suggests that such conditions are unavailable for most such concepts. Why should the concept of knowledge be an exception?⁷

More positively, one may speculate that standard accounts of justification have failed to deal convincingly with the traditional problem of the regress of justifications—what justifies the justifiers?—*because* they have forbidden themselves to use the concept of knowledge. $E = K$ suggests a very modest kind of foundationalism, on which all one’s knowledge serves as the foundation for all one’s justified beliefs. Perhaps one can understand how something could found belief only by thinking of it as knowledge.

II

The present enterprise is limited. Most discussions of evidence concern the relation in which something is evidence *for* something. When is e evidence for the hypothesis h ? When is evidence for h sufficient to justify belief in h ? This paper does not primarily address such questions, impor-

⁷ Williamson (1995) develops an account of knowledge in line with these remarks.

tant though they are. It concerns the nature of the first relatum e of the evidential relation rather than its relation to the second relatum h .⁸ Analogously, informative accounts of the relata of the causal relation leave open many other aspects of the relation. That evidence is evidence *for* something constrains what evidence itself can be, but one can use that constraint without a systematic theory of the evidential relation. Of course, evidence has other parameters too: the subject for whom it is evidence, the time at which it is evidence, and perhaps others. Henceforth, “theories of evidence” means theories of the dependence of evidence on these parameters, not theories of what it is evidence for. What justifies this methodological abstraction is the idea that one has a *body of evidence*; one has evidence for h insofar as that body supports h . The idea is elaborated in what follows.

Why does it matter what counts as evidence? Consider the idea that one should proportion one’s belief in a proposition to one’s evidence for it. How much evidence one has for the proposition depends on what one’s evidence is. More precisely, a theory of evidence is needed to give bite to what Carnap calls *the requirement of total evidence*:

In the application of inductive logic to a given knowledge situation, the total evidence available must be taken as a basis for determining the degree of confirmation. (1950, p. 211; compare Hempel 1965, pp. 63–7)⁹

If too much or too little is counted as evidence, then inductive principles will be misapplied. Given the requirement of total evidence, disputes between different theories of evidence are not merely verbal; they involve disagreements as to which inductive conclusions are warranted. Formulations of the total evidence requirement in terms of knowledge encourage $E = K$, which identifies the total evidence available with the total knowledge available. For example, Peirce writes:

I cannot make a valid probable inference without taking into account whatever knowledge I have (or, at least, whatever occurs to my mind) that bears on the question. (1932, p. 461)

⁸ The papers in the representative collection Achinstein (1983), for example, are largely concerned with questions about the evidential relation at the expense of questions about its first relatum.

⁹ The total evidence available must *not* be taken as a basis for determining the degree to which e increases the confirmation of h , for if e is part of the total evidence available then the confirmation of h prior to the acquisition of the total evidence presently available is also relevant. This is in effect the problem of old evidence; see Glymour (1980, pp. 85–93), Earman (1992, pp. 119–35), Howson and Urbach (1993, pp. 403–8), Maher (1996) and Williamson (1997). But that does not undermine the claim that the total evidence now available must be taken as a basis for determining the degree to which h is now confirmed in the non-comparative sense.

Carnap himself describes the evidence as (observational) knowledge. Given $E = K$, the original idea becomes something like this: one should proportion one's belief in a proposition to the support which it receives from one's knowledge.

Theories of evidence also play a role when theses of the underdetermination of theory by data are assessed, for the data in question are the actual or potential evidence. If too much or too little is counted as evidence, then the standard for underdetermination is set uninterestingly high or uninterestingly low. $E = K$ implies that underdetermination theses of the relevant kind must count all knowable facts as data. Although this condition does not automatically make any argument for underdetermination circular, it is not easily met. Consider, for example, the underdetermination thesis that the theoretical facts do not supervene on the evidential facts: two possible worlds can differ in the former without differing in the latter.¹⁰ One cannot establish this claim just by showing that the theoretical facts do not supervene on the observable facts (in some sense of "observable"); one must also show that they do not supervene on all the knowable facts. The gap would be filled by an argument that the knowable facts supervene on the observable facts, for then whatever failed to supervene on the observable facts would fail to supervene on the knowable facts too, by the transitivity of supervenience. But any such argument risks begging the question against the view that at least some theoretical facts are knowable. The issue cannot be pursued here; the present point is just how theories of evidence interact with underdetermination theses.

III

What cannot be expected of a theory of evidence is a recipe for deciding in practice whether our evidence includes a given item. In general, a philosophical theory of a concept is not required to give a recipe for deciding in practice whether any given item falls under that concept. Nevertheless, it might be objected that the concept of evidence is special in this respect. For if it were problematic whether one's evidence included something, one would need evidence to decide whether one's evidence included it; an infinite regress looms. Therefore, the objection runs, in some basic sense of "evidence" it must be unproblematic whether one's evidence includes any given item, and an adequate theory of evidence must explain how it manages to be so unproblematic.

¹⁰ Compare EI₃ in the useful classification of kinds of empirical indistinguishability in Earman (1993, p. 21).

$E = K$ fails this constraint. There is no infallible recipe for deciding in practice whether we know a proposition p , and therefore (given $E = K$) for deciding whether our evidence includes p . Sometimes we reasonably believe ourselves to know p , when in fact we do not know p , because p is false. Reputable authorities assert that Henry V died in 1422; I have no grounds for doubting them, and reasonably believe myself to know by testimony that Henry V died in 1422. But it is not inconceivable that he died in 1423, some elaborate conspiracy being responsible for present evidence to the contrary. According to $E = K$, if I do know that Henry V died in 1422, then my total evidence includes the proposition that Henry V died in 1422; but if Henry V died in 1423, then my belief that my total evidence includes that proposition is mistaken—my total evidence includes only the proposition that reputable authorities assert that Henry V died in 1422. $E = K$ is an *externalist* theory of evidence, in at least the sense that it implies that one's evidence does not supervene on one's internal physical states.¹¹

How does $E = K$ avoid the threatened regress of evidence? The regress comes if evidence-based belief in a proposition p must always be preceded by evidence-based belief in a proposition about the evidence for p . We can distinguish two senses of "evidence-based". Call S 's belief in p *explicitly* evidence-based if it is influenced by prior beliefs about the evidence for p . Explicitly evidence-based beliefs may be more common in science than in everyday life. Call S 's belief in p *implicitly* evidence-based if it is appropriately causally sensitive to the evidence for p . A belief can be both explicitly and implicitly evidence-based. Now explicitly evidence-based belief in p is not always preceded by *explicitly* evidence-based belief in a proposition about the evidence for p ; this is consistent with $E = K$ and most other theories of evidence. An explicitly evidence-based belief is influenced by a prior state of belief in a proposition about the evidence for p , and something has gone wrong if the latter belief is not at least implicitly evidence-based; but it need not be explicitly evidence-based. Thus there is no regress of explicitly evidence-based belief. There would be a different regress if implicitly evidence-based belief in p were always preceded by implicitly evidence-based belief in a proposition about the evidence for p . But the causal sensitivity of the belief in p to the evidence for p need not be mediated by further *beliefs* about the evidence for p ; there need be no such beliefs.

How can a belief in p be implicitly evidence-based, if we are liable to misidentify the evidence for p ? If the real evidence differs from the apparent evidence, won't the belief be causally sensitive to the latter rather than

¹¹ One's evidence does supervene on one's *mental* states, on the account of knowledge in Williamson (1995).

the former? But causal sensitivity need not be perfect to be genuine. There can be a non-accidental rough proportionality between the strength of the belief and the strength of the evidence, even if distortions sometimes occur.

Similar questions arise about explicitly evidence-based belief. How can one follow the rule “Proportion your belief in p to your evidence for p ” when one doesn’t know exactly what one’s evidence is? Given $E = K$, the rule becomes “Proportion your belief in p to the support which p receives from your knowledge”: but isn’t one at best following the rule “Proportion your belief in p to the support which p receives from *what you believe to be* your knowledge”? Consider an analogy. We can follow the rule “Proportion your voice to the size of the room”. This is not because we are infallible about the size of the room. We sometimes make mistakes; but it does not follow that we are really following the rule “Proportion your voice to *what you believe to be* the size of the room”. After all, it is often quite hard to know what beliefs one has about the size of a room; we are fallible in our beliefs about such beliefs. In general, if the fallibility of our beliefs about X posed a problem, it would not be solved by the move to our beliefs about our beliefs about X , because they are fallible too. But fallibility does not pose a problem here. To make a mistake in following a rule is not to follow a different rule. The rule is a standard of correctness for action, not a description of action. To apply the rule “Proportion your voice to the size of the room”, one needs beliefs about the size of the room, but they need not be true—although when they are false, one’s application is faulty. Similarly, to apply the rule “Proportion your belief in p to the support which p receives from your knowledge”, one needs beliefs about the support which p receives from one’s knowledge, and therefore about one’s knowledge, but those beliefs need not be true—although when they are false, one’s application is faulty.

None of this would be much consolation if our beliefs about our knowledge were hopelessly unreliable. Sceptics say that those beliefs have no rational basis, but they say the same about most of our other beliefs too; I assume that we need follow them in neither case. Although we have no infallible recipe for deciding whether we know p , in practice we are often in a position to know whether we know p .¹²

Other theories of evidence try to make it something that we can infallibly identify. They are unlikely to succeed. Arguably, there is no non-trivial condition such that we are always in a position to know whether it

¹² Nor are the ways in which we decide whether we know p simply the ways in which we decide whether we believe that we know p . If I want to check whether I now really know that Henry V died in 1422, I may return to my sources; but it would be irrelevant to return to them if I merely wanted to check whether I now really believe that I know that he died in 1422.

obtains (Williamson 1996). In particular, we are not always in a position to know whether our evidence includes p . A fortiori, it is possible to judge that something is evidence and be mistaken. Someone can incorrectly take his evidence to include the proposition that Henry V died in 1423. Attempts are sometimes made to solve the problem by interiorizing evidence: it becomes one's present experience, or one's present degrees of belief, or the like. These attempts are quaint relics of Cartesian epistemology. Knowledge of the present contents of one's own mind is neither unproblematic nor prior to knowledge of other things. It is not obvious to me how many shades of blue I am presently experiencing, or to what degree I believe that there was once life on Mars. Moreover, if one's evidence were restricted to the contents of one's own mind, how could it play the role that it actually does in science? The evidence for the proposition that the sun is larger than the earth is not just my present experiences or degrees of belief. If the evidence is widened to include other people's experiences or degrees of belief, or my past ones, then my identification of it becomes even more obviously fallible. In any case, that does not seem to be the right widening; it is more plausible that the evidence for a scientific theory is the sort of thing which is published in scientific journals. If that is what evidence is, our identification of it is certainly fallible.

IV

Here is a schematic argument for $E = K$:

All evidence is propositional.
 All propositional evidence is knowledge.
 All knowledge is evidence.

All and only knowledge is evidence.

The argument is obviously valid, but its premises are contentious. Its aim is simply to divide the contentiousness of the conclusion into manageable portions; §§V, VI and VII respectively defend the three premises. Since "knowledge" here means propositional knowledge, each premise follows from the conclusion; thus the conclusion is equivalent to the conjunction of the premises.

What is it for evidence to be propositional? The idea is that we should be able to refer to evidence by using "that" clauses. My evidence for the conclusion that the house was empty is *that* it was silent, *that* no lights were on in the evening, *that* the telephone went unanswered, For simplicity, "that" clauses (including "that") are treated as referring to abstract

objects: propositions.¹³ Which proposition such a clause refers to may depend on context. Propositions are the objects of propositional attitudes, such as knowledge and belief; they can be true or false. One's evidence is propositional just in case it is a set of propositions.

V

Why should all evidence be propositional? It would not be on a broad interpretation of "evidence". In the courts, a bloodied knife is evidence. It is natural to say that my evidence that I am getting a cold includes various sensations. Some philosophers apply the term "evidence" to nonpropositional perceptual states; Quine restricts it to the stimulation of sensory receptors (1969, p. 75). How can "All evidence is propositional" do more than stipulate a technical use for the word "evidence"?

Indiscriminate description of the ordinary use of a term and arbitrary stipulation of a new use are not the only options. We can single out theoretical functions central to the ordinary concept of evidence, and ask what serves them. That strategy will be pursued here. The argument substantiates the familiar claim that only propositions can be reasons for belief.¹⁴ It also suggests a further conclusion: one grasps (can think) the propositions which are one's evidence.

Consider inference to the best explanation (Lipton 1991). We often choose between hypotheses by asking which of them best explains our evidence—which of them, if true, would explain the evidence better than any other one would, if true. Fossil evidence enables us to answer questions about terrestrial life in this way. Even if inference to the best explanation is not legitimate in all theoretical contexts, what matters for present purposes is that, where evidence does enable us to answer a question, a central way for it to do so is by inference to its best explanation. Thus evidence is the kind of thing which hypotheses explain. But the kind of thing which hypotheses explain is propositional. Therefore evidence is propositional.

The kind of thing which hypotheses explain is propositional. Inference to the best explanation concerns why-explanations, which can be put in the form "--- because ...", which is ungrammatical unless declarative sentences, complements for "that", replace the blanks. We cannot simply

¹³ On alternative views of "that" clauses, they sometimes refer to facts rather than propositions, or never refer at all. The arguments of this paper can be adapted to such frameworks.

¹⁴ See e.g. Unger (1975, pp. 204–6) and Davidson (1986); for opposing views, Moser (1989, pp. 47–125) and Millar (1991).

explain Albania, for “Albania because ...” is ill-formed. We can sometimes make sense of the injunction “Explain Albania!”, but only when the context allows us to interpret it as an injunction to explain why Albania exists, or has some distinctive feature. What follows “why” is a declarative sentence, expressing the proposition to be explained—*that* Albania exists, or *that* it has the distinctive feature.¹⁵ Likewise for events: “Explain World War I!” enjoins one to explain why it occurred, or had some distinctive feature. Again, the sensation in my throat is evidence for the conclusion that I am getting a cold in the sense that the hypothesis that I am getting a cold would best explain why I have that sensation in my throat. The evidence to be explained is *that* I have that sensation in my throat (not just that I have *a* sensation in my throat). Even in the courts, the bloodied knife provides evidence because the prosecution and defence offer competing hypotheses as to why it was bloodied or how it came into the accused’s possession; the evidential proposition is *that* it was bloodied or *that* it came into the accused’s possession. The knife is a source of indefinitely many such propositions.

One can use an hypothesis to explain why --- only if one grasps the proposition that ---. Thus only propositions which one grasps can function as evidence in one’s inferences to the best explanation. By this standard, only propositions one grasps count as part of one’s evidence.

Similar points apply to explicitly probabilistic reasoning. The best way of comparing the conditional probabilities of two hypotheses h and h^* on evidence e , $P(h|e)$ and $P(h^*|e)$, is often by calculating the inverse probabilities of e on h and h^* , $P(e|h)$ and $P(e|h^*)$. Take a simple case; the comparison of hypotheses about more important matters on the basis of more complex evidence can have the same kind of underlying structure. A bag contains ten red or black balls; we wish to estimate how many of them are red; we can gain evidence only by sampling one ball at a time, noting its colour and replacing it. A good way to compare the probabilities of hypotheses about the number of red balls is by calculating the probabilities of the actual outcome e of the sampling (say, red fifteen times and black five times) on those hypotheses. One way of using those probabilities is to regard h as more probable than h^* ($P(h|e) > P(h^*|e)$) just in case h makes e more probable than h^* does ($P(e|h) > P(e|h^*)$). Bayesians take this method to involve assigning the same prior probability to h and h^* ($P(h) = P(h^*)$); they treat as equally legitimate assignments of unequal prior probabilities to the hypotheses—perhaps reflecting differences in explanatory virtues such as simplicity and ele-

¹⁵ It makes no significant difference if what is to be explained is why e rather than f , e.g. why Albania rather than Bosnia was peaceful in 1995 (Lipton 1991, pp. 75–98). The evidence in question would be the propositions that Albania was peaceful in 1995 and that Bosnia was not.

gance.¹⁶ To allow for such cases, their general rule weights the probability of e on h by the prior probability of h ; thus $P(h|e) > P(h^*|e)$ just in case $P(h)P(e|h) > P(h^*)P(e|h^*)$.¹⁷ For present purposes, it does not matter whether Bayesians are right to introduce prior probabilities here. The point is that such probabilistic comparisons of hypotheses on the evidence depend on the probabilities of the evidence on the hypotheses. But what has a probability is a proposition; the probability is the probability *that ...*. At least, that is so when “probability” has to do with the evidential status of beliefs, as now; if we speak in this connection of the probability of an event, we mean the probability *that it occurred*.¹⁸ We might initially suppose that, in $P(x|y)$, only x need be a proposition, but the relation between $P(x|y)$ and $P(y|x)$ means that y must be a proposition too; what gives probability must also receive it. Moreover, these probabilities, as measures of degrees of belief warranted by evidence, are idle unless the subject grasps x and y .

More straightforward uses of evidence also require it to be propositional. In particular, our evidence sometimes rules out some hypotheses by being *inconsistent* with them. For example, the hypothesis that only males have varicose veins is inconsistent with much medical evidence. But only propositions can be inconsistent in the relevant sense. If evidence e is inconsistent with an hypothesis h in that sense, it must be possible to *deduce* $\sim h$ from e ; the premises of a deduction are propositions. Moreover, the subject who deduces $\sim h$ from e must grasp e .

Only propositions that we grasp serve the central evidential functions of inference to the best explanation, probabilistic confirmation and the ruling out of hypotheses. Could non-propositional items count as evidence by serving other central functions of evidence? For example, they might serve as the inputs to a non-inferential process whose outputs were beliefs. But suppose that we are choosing between hypotheses according to which best explains our evidence, or is most probable on our evidence, or is not ruled out by our evidence. The argument so far shows that non-propositional evidence would be irrelevant to our choice. Moreover, in choosing between hypotheses in those ways, we can use only proposi-

¹⁶ Perhaps the probabilistic reasoning can be assimilated to inference to the best explanation, or vice versa; if so, no matter.

¹⁷ Assume that $P(e)$, $P(h)$ and $P(h^*)$ are all non-zero.

¹⁸ Objective probabilities, in the sense of chances determined in the natural world independently of our beliefs, are irrelevant here. We sometimes speak too of the probability of one property, concept or predicate conditional on another, e.g. the probability of lung cancer conditional on smoking; but the probabilities relevant to the argument are the probabilities of hypotheses, which, unlike properties, concepts and predicates, have truth-values. That someone has lung cancer is evidence that he smoked; the unattributed property of lung cancer is not by itself evidence of anything.

tions we grasp. In those respects, any evidence other than propositions we grasp would be impotent. Evidence may well have central functions additional to those considered above; the point is that genuine evidence would make a difference to the serving of the functions considered above, whatever else it made a difference to. Certainly, defences of non-propositional evidence have not been based on an appreciation of its impotence in those respects. Since only propositions we grasp make a difference of the requisite kind, only propositions we grasp are our evidence.

My positive argument for that conclusion is now complete. Nevertheless, perceptual experience is often regarded as a kind of non-propositional evidence. Do the considerations above somehow fail to do it justice? In the remainder of this section, I will rebut objections to the view that our perceptual evidence consists of propositions that we grasp. I will not argue against contrary views; in effect, that has just been done. What follows is only an additional check.

Experiences provide evidence; they do not consist of propositions. So much is obvious. But to provide something is not to consist of it. The question is whether experiences provide evidence just by conferring the status of evidence on propositions. On that view, consistent with $E = K$, the evidence for an hypothesis h can include propositions e_1, \dots, e_n which count as evidence for me only because I am undergoing a perceptual experience ε (as a limiting case, h might be e_1). The threatening alternative is that ε can itself be evidence for h , without the mediation of any such e_1, \dots, e_n . Both views permit ε to have a non-propositional, non-conceptual content, but only the latter permits that content to function directly as evidence. What could e_1, \dots, e_n be?

Consider an example. I am trying to identify a mountain by its shape. I can see that it is pointed; that it is pointed may be part of my evidence for believing that it is not Ben Nevis. However, the proposition that it is pointed does not begin to exhaust my present perceptual evidence. No description of the mountain in words seems to capture the richness of my visual experience of its irregular shape. But it does not follow that my evidence is non-propositional. If I want to convey my evidence, I might point and say "It is that shape".¹⁹ Of course, the mere linguistic meaning of the sentence type "It is that shape" does not convey my evidence, for

¹⁹ See McDowell (1994, pp. 56–9) for a similar proposal. It is not being used here to deny that perceptual experience has non-conceptual content (see Peacocke 1992, pp. 84). Christensen (1992, p. 545) discusses such a proposal in a Bayesian context. It is no part of the present proposal that the demonstrative beliefs have a special certainty. Christensen asks whether they can "connect with other beliefs in the way that would be necessary for them to fulfill their intended evidential role". The connections will not be purely syntactic, but fifty years of confirmation theory have shown that confirmation is not a purely syntactic matter.

it is independent of the reference of “that shape” in a particular context of utterance. Only by using the sentence in an appropriate context do I express the proposition at issue. My token of “that shape” still expresses a constituent of that proposition, even if you cannot grasp that constituent without having a complex visual experience with a structure quite different from the constituent structure of the proposition. The proposition that the mountain is that shape is contingent (it could have been another shape) and known a posteriori (I do not know a priori that I am not including the tip of another mountain behind in the profile). But in ordinary circumstances I can know that the mountain is that shape, and a fortiori grasp the proposition that it is, when “that shape” does not refer to an absolutely specific shape. Of course, I cannot see *exactly* what shape the mountain is; I can only see roughly what profile it presents to me, and cannot see round the back.²⁰ Nor is the knowledge that the mountain is that shape restricted to the present context; you can have it too, and we can retain it in memory. Properties other than shape are similar in those respects.

In unfavourable circumstances, one fails to gain perceptual knowledge, perhaps because things are not the way they appear to be. One does not know that things are that way, and $E = K$ excludes the proposition that they are as evidence. Nevertheless, one still has perceptual evidence, even if the propositions which it supports are false. If that evidence consists of propositions, what are they? The obvious answer is: the proposition that things appear to be that way (the mountain appears to be that shape). Of course, unless one has reason to suspect that circumstances are unfavourable, one may not consider the cautious proposition that things appear to be that way; one may consider only how they really are. But it does not follow that one does not know that things appear to be that way, for one knows many propositions without considering them. When one is walking, one normally knows that one is walking, without considering the proposition. Knowing is a state, not an activity. In that sense, one can know without consideration that things appear to be some way. When I believe falsely that circumstances are favourable, I believe falsely that I am gaining perceptual knowledge about the environment, and therefore that my evidence includes those propositions believed to be known. But our fallibility in identifying our evidence is nothing new, and my actual evidence may justify my false beliefs about my evidence.

²⁰ “That shape” must be unspecific enough to give my knowledge that the mountain is that shape an adequate *margin for error* in the sense of Williamson (1994, pp. 226–30). The unspecificity makes the present proposal closer to that of McDowell (1994, pp. 170–1) than to that of Peacocke (1992, pp. 83–4).

In order to grasp the proposition that things appear to be some way, one needs the concept of appearance.²¹ Although one's grasp of it may be inarticulate, one must have some inkling of the distinction between appearance and reality. For instance, one should be willing in appropriate circumstances to give up the belief that things were that way while retaining the belief that they appeared to be that way; in the absence of such dispositions, it would be implausible to attribute the qualified belief that things appear to be that way rather than the unqualified belief that they are that way. Perhaps some young children and animals have beliefs and perceptual experiences without even implicitly grasping the concept of appearance. Suppose that such a simple creature is given a drug which causes the hallucinatory appearance that there is food ahead; as a result, it comes to believe falsely that there is food ahead. Does it have any evidence for that belief? According to $E = K$, its evidence cannot be that things appear some way, for it cannot grasp that proposition. Perhaps it knows that the situation is *like* one in which there is food ahead, where the concept of likeness covers both likeness in appearance and other kinds of likeness indifferently, so that grasp of the concept of likeness does not require grasp of the concept of appearance. If the creature does not even know that the situation is like one in which there is food ahead, then it is plausible to deny that it has perceptual evidence that there is food ahead. It does not recognize the features of its perceptual experience which, recognized, would provide it with evidence. We can use the proposition that there appears to be food ahead as evidence, but the simple creature cannot. Although the hallucinatory appearance causes a belief, the causal relation is not an evidential one.

Very simple creatures have no concepts, grasp no propositions, and have no beliefs or knowledge (it is sometimes even argued that any creature which lacks the distinction between appearance and reality is in this predicament). Such creatures have no evidence, for they have no degrees of belief, and degrees of belief are what evidence justifies.

S can use as evidence only propositions which S grasps. Since S can use S's evidence as evidence, propositions which S grasps are S's evidence. What has not yet been argued is that those propositions count as evidence by being known.

²¹ A concept here is a constituent of a proposition; this assumes that the semantically significant constituents of a sentence express constituents of the proposition expressed by the whole sentence, but it does not assume any particular theory about the nature of those constituents.

VI

Why should all propositional evidence be knowledge? The thesis is that if S's evidence includes a proposition *e*, then S knows *e*. If I do not know that the mountain is that shape, then that it is that shape is not part of my evidence. As in the previous section, the argument is from the function of evidence.²² Indeed, the thesis draws support from the role of evidence cited there, in inference to the best explanation, probabilistic reasoning and the exclusion of hypotheses. When we prefer *h* to *h** because *h* explains *e* better than *h** does, we are standardly assuming *e* to be known; if we do not know *e*, why should *h*'s capacity to explain *e* confirm *h* for us? It is likewise hard to see why the probability of *h* on *e* should regulate our degree of belief in *h* unless we know *e*. Again, an incompatibility between *h* and *e* does not rule out *h* unless *e* is known. But it is prudent to consider the matter more carefully.

Suppose that balls are drawn from a bag, with replacement. In order to avoid issues about the present truth-values of statements about the future, assume that someone else has already made the draws; I watch them on film. The following situation can arise. I have seen draws 1 to N; each was red (produced a red ball). I have not yet seen draw N+1. I reason probabilistically, and form a justified belief that draw N+1 was red too. My belief is in fact true. But I do not know that draw N+1 was red. Consider two false hypotheses:

h: Draws 1 to N were red; draw N+1 was black.

*h**: Draw 1 was black; draws 2 to N+1 were red.

It is natural to say that *h* is consistent with my evidence and that *h** is not. In particular, it is consistent with my evidence that draw N+1 was black; it is not consistent with my evidence that draw 1 was black. Thus my evidence does not include the proposition that draw N+1 was red. Why not? After all, I have a justified true belief that it was red. The obvious answer is that I do not *know* that draw N+1 was red; the unsatisfied necessary condition for evidence is knowledge. An alternative answer is that I have not *observed* that draw N+1 was red. That is equally good for the purposes of this section (although not for those of the next), for observing the truth of *e* includes *e* in my evidence only by letting me know *e*. If I observe the truth of *e* and then forget all about it, my evidence no longer includes *e*. It is hard to see how evidence could discriminate between hypotheses in the way we want it to unless it had knowledge as a necessary condition.

²² For linguistic arguments that if S's reason or justification is that --- then S knows that --- see Unger (1975, pp. 206–14).

If evidence required only justified true belief, or some other good cognitive status short of knowledge, then a critical mass of evidence would set off a kind of chain reaction. Our known evidence justifies belief in various true hypotheses; they would count as evidence too, so this larger evidence set would justify belief in still more true hypotheses, which would in turn count as further evidence The result would be very different from our present conception of evidence.

That propositional evidence is knowledge entails that propositional evidence is true. We may treat false propositions as evidence, but it does not follow that they are evidence. No true proposition is inconsistent with my evidence, although I may think that it is. That the ground is wet is evidence that it rained last night only if the ground *is* wet (but not: only if it rained last night). If the ground is not wet, then only a counterfactual holds: if it had been wet, that would have been evidence that it rained last night.²³ If the convincing but lying witness says that the accused was asleep at the time of the murder, then it is part of the evidence for the innocence of the accused that the witness *said* that he was asleep then. It is not part of the evidence for his innocence that he was asleep, for it is consistent with the evidence that he was not. The rival view, that a false proposition can become evidence through a sufficient appearance of truth, gains most of its appeal from the assumption, disposed of in §III, that we must have an infallible way of identifying our evidence.

Once it is granted that all propositional evidence is true—and therefore, by the previous section, that all evidence consists of true propositions—adjusting our beliefs to the evidence has an obvious point; it is a way of adjusting them to the truth. True evidence can still support false conclusions, but not without limit. The maxim “Proportion your belief to your evidence” requires more than the mere internal coherence of one’s belief system; it does so because evidence must be true. Even if an internally coherent belief system cannot be wholly false, a given belief system with a given degree of internal coherence can be better or worse proportioned to the evidence, depending on what the evidence is. But the evidence is not a wholly external standard either, if it is known.

Another consequence of the claim that propositional evidence is knowledge is that propositional evidence is believed—at least, if knowledge entails belief, which is granted here. The case of perception may seem to

²³ Stampe (1987, p. 337) takes a similar view of reasons. Millar (1991, p. 65) says that we ordinarily think of evidence as consisting in facts; that is also how we ordinarily think of the objects of knowledge. If facts are distinguished from true propositions, then the arguments of this paper can be adjusted accordingly, but the individuation of facts must be reconciled with the individuation of knowledge and evidence (presumably, the fact that Hesperus is bright is the fact that Phosphorus is bright).

suggest that propositional evidence is not always believed. In conformity with the previous section, a piece of perceptual evidence is a proposition e that things are *that way*. According to $E = K$, my evidence includes e because I know that things are that way. But, it is suggested, that does not go back far enough; my evidence includes e because it is perceptually apparent to me that things are that way, whether or not I believe that they are that way. Even if I do believe e , my evidence included e before I came to believe it; I came to believe it because it was perceptually apparent. If “It is perceptually apparent that ---” entails “---”, this view allows that evidential propositions are always true; what it denies is that they are always believed, and therefore that they are always known.

If my evidence includes a proposition e , then I grasp e , as the previous section argued. Thus, if I fail to believe e , my problem is not conceptual incapacity. Perhaps I have simply not had time to form the belief; perhaps I suspect, for good or bad reasons, that I am the victim of an illusion. For my evidence to include e , must I at least be *in a position* to know e ? If so, then the proposed view does not differ radically from $E = K$. Given $E = K$, the evidence in my actual possession consists of the propositions which I know, but there is also the evidence in my potential possession, which consists of the propositions which I am in a position to know. The alternative to $E = K$ takes my evidence to be the evidence in my potential possession, not just the evidence in my actual possession. To bring out the difference between the two views, suppose that I am in a position to know any one of the propositions p_1, \dots, p_n without being in a position to know all of them; there is a limit to how many things I can attend to at once. Suppose that in fact I know p_1 and do not know p_2, \dots, p_n . According to $E = K$, my evidence includes only p_1 ; according to the alternative, it includes p_1, \dots, p_n . Let q be a proposition which is highly probable given p_1, \dots, p_n together, but highly improbable given any proper subset of them; the rest of my evidence is irrelevant to q . According to $E = K$, q is highly improbable on my evidence; according to the alternative, q is highly probable on my evidence. $E = K$ gives the more plausible verdict, because the high probability of q depends on an evidence set to which as a whole I have no access.

The contrast with $E = K$ is more radical if the alternative allows my evidence to include e even when I am not in a position to know e . For example, it is perceptually apparent to me that things are some way; I am not hallucinating; but since I know that I have taken a drug which has a 50% chance of causing me to hallucinate, I am not in a position to know that things are that way. According to the radical alternative, my evidence nevertheless includes the proposition that things are that way, because it is perceptually apparent to me that they are; thus my evidence is inconsistent

with the hypothesis that I am hallucinating and things are not way, even though, for all I am in a position to know, that hypothesis is true. According to $E = K$, my evidence includes at best the proposition that things appear to be that way. Surely, if I proportion my belief to my evidence, I shall not dismiss the hypothesis that I am hallucinating and things are not that way; $E = K$ gives the better verdict. Perceptual cases do not show that we sometimes fail to believe our evidence.

A truth does not become evidence merely by being believed, nor even by being justifiably believed, as the example of the proposition that draw N+1 was red showed above. Nothing short of knowledge will do. But is even knowledge enough?

VII

Why should all knowledge be evidence? In general, we want as much evidence as possible. Any restriction on what counts as evidence must be well-motivated by the function of evidence. Given what has already been argued, one's evidence includes only propositions one knows. Is any further restriction needed? If one is assessing an hypothesis, and knows something which bears on its truth, shouldn't that thing be part of one's total evidence? If one knows a proposition, one is in a position to use it as evidence in inference to the best explanation, probabilistic reasoning and the exclusion of hypotheses; the presumption is that it counts as part of one's evidence. The burden of proof is on the claim that not all knowledge is evidence. This section examines attempts to discharge that burden, and finds them wanting.

If our evidence includes a proposition e , it might be objected, then e is evidence for itself; but since little of what we know is self-evident, our evidence includes little of what we know. The objection exploits an ambiguity in "self-evident". If e is evidence for h whenever e entails h and the evidence includes e , then trivially e is evidence for itself whenever the evidence includes e .²⁴ The objection provides no reason to doubt that what we know is self-evident in that sense. But "self-evident" also suggests a stronger sense, in which self-evident propositions are *essentially* evidence. Most of what we know is obviously not self-evident in this stronger sense, but that is quite consistent with the claim that all knowledge is evidence, for that claim does not entail that all knowledge is self-evident in the stronger sense. Indeed, given obvious facts about knowledge, $E = K$

²⁴ The condition supposed to be sufficient for " p is evidence for q " is not also supposed to be necessary.

implies that none of our evidence is self-evident in the stronger sense, for no proposition is essentially known.

The objector may not be satisfied. For let e be known. Then, if all knowledge is evidence, it should be correct to cite e in answer to the question “What is the evidence for e ?”. But when e is a non-obvious truth (e.g. that Henry V died in 1422 or that blood circulates), such an answer would be insulting. Certainly, our knowledge of such truths depends on evidence (e.g. about documents or experiments) related to them in less trivial ways; the questioner is entitled to ask for this evidence. But it does not follow that e is not evidence. For what is wrong with citing e in answer to the question “What is the evidence for e ?” may be conversational inappropriateness rather than untruth. It would be conversationally inappropriate to cite Mary in answer to the question “Who lives in the same house as Mary?”; nevertheless, it is true that Mary lives in the same house as Mary (Grice 1989). We assume that the questioner does not need to be told the obvious. For conversational reasons, someone who asks for my evidence for p presumably wants evidence which (in the circumstances) I could have without knowing p ; call this *independent* evidence for p .

Some may insist that all evidence for p is independent evidence for p . That is consistent with $E = K$, because $E = K$ does not say when evidence is evidence *for* p . “Independent” might come from “for”, not from “evidence”. What $E = K$ does imply is that if I know p , then my total evidence entails p .

$E = K$ implies that whether one proposition is evidence may depend on whether other propositions are evidence for it. My evidence includes the proposition that Henry V died in 1422, but would not have done so if there had not been documentary evidence for that date. This is hardly an objection to $E = K$, for we have no good reason to expect evidence to consist of epistemically self-sufficient nuggets of information. Given what was argued in the previous section, *my* evidence does not include the propositions which historians know about the documents, because I do not know those propositions; but if I did learn them, that would not extrude the proposition that Henry V died in 1422 from my evidence. Would it be redundant? A proposition e is a redundant part of one’s evidence only if the rest of one’s evidence entails e , for otherwise one’s evidence including e rules out (i.e. is inconsistent with) some hypotheses (e.g. $\sim e$) which one’s evidence excluding e does not rule out. The proposition that contemporary documents say that Henry V died in 1422 does not make the proposition that Henry V died in 1422 redundant in this sense; nor does the proposition that I seem to remember that Henry V died in 1422. Moreover, redundant evidence is still evidence. Indeed, it

is logically possible for each proposition in one's evidence to be redundant.

If all my knowledge is evidence, and I know p , then my evidence entails that no evil demon has tricked me into falsely believing p . But my independent evidence for p , which I could have without knowing p , need not entail that no evil demon has tricked me into falsely believing p . This may help to explain our reluctance to accept that we know that no evil demon has tricked us into falsely believing p . When I know p , and deduce q from p , my independent evidence for p sometimes fails to support q . Thus q looks unsupported. But why assume that all evidence is independent evidence? Even when prior (therefore independent) evidence for p is necessary for knowing p , it is not sufficient (granted true belief in p); knowledge of p has epistemic powers beyond those of its evidential precondition. Without independent evidence for p , I may be unable to convince a sceptical interlocutor of p ; but I can still know p .²⁵

Sometimes the underlying worry is a different one. For most known propositions p , the question "What is the evidence for p ?" tends to elicit ways in which p , although well-supported, is not beyond all doubt; we can conceive our present belief in p to be mistaken. We are uneasy with the idea of uncertain evidence. Is this just the old Cartesian prejudice that only unshakable foundations will do? The worry cannot be so easily dismissed. It takes a particularly sharp form in a Bayesian context. The standard way of accommodating new evidence e is by conditionalizing on it. The new unconditional probability of a proposition is its old probability conditional on e (where the old probability of e was non-zero), thus:

$$P_{\text{new}}(h) = P_{\text{old}}(h|e) = P_{\text{old}}(h \& e) / P_{\text{old}}(e).$$

In particular, the new probability of e itself is 1. Now if the old probability of h was 1, so is its new probability; for if $P_{\text{old}}(h) = 1$ then $P_{\text{old}}(h \& e) = P_{\text{old}}(e)$. Thus once a proposition is conditionalized on as evidence, it acquires probability 1, and retains it no matter what further evidence is conditionalized on (provided that the prior probability of all this evidence is non-zero). But most of our knowledge has no such status. Further evidence could undermine it.

I put exactly one red ball and one black ball into an empty bag, and will make draws with replacement. Let h be the proposition that I put a black ball into the bag, and e the proposition that the first ten thousand draws are all red. I know h by a standard combination of perception and memory, because I saw that the ball was black as I put it into the bag a moment ago.

²⁵ If the argument is correct, appeals to the context-dependence of "know" in response to scepticism may be unnecessary. However, the argument does not imply that "know" is context-independent. The matter obviously deserves far more extended discussion.

Nevertheless, if after ten thousand draws I learn e , I shall have ceased to know h , because the evidence that I shall then have will make it too likely that I was somehow confused about the colours of the balls. Of course, what I know now is true, and so will never be discovered to be false, but it does not follow that there will never be misleading future evidence against it. My present knowledge is consistent with e ; on simple assumptions, e has a probability of $1/2^{10,000}$ on my present evidence. If I subsequently learn e , the probability of h on that future evidence will be less than 1. But if conditionalization on evidence of non-zero probability will give h a probability less than 1, then the present probability of h is less than 1, so h is not part of my present evidence. The problem is general: if misleading future evidence of positive probability can undermine my knowledge that I put a black ball into the bag, it can undermine most of my present knowledge. It looks as though h should count as part of my present evidence, and therefore receive probability 1, only if h is bound to be a rational belief for me in the future, come what may. Few propositions will pass that test. Indeed, not even e passes the test, for later evidence may make it rational for me to believe that I had misremembered the outcome of the first ten thousand draws; several eyewitnesses may insist that I was misremembering; but since uncertainty about e does not make h certain, it does not rehabilitate h as evidence. By this line of argument, either we know very little or very little of our knowledge is evidence.

What empirical propositions qualify as evidence by the proposed test? One might suppose that the best candidates would be propositions about the present (traditionally, propositions about the subject's present mental states). But the test requires evidence to remain certain as time passes, so if a proposition about the present is to be evidence, it must remain certain long after the time it is about has passed. But even if it is absolutely certain for me today that I seem to see a blue patch, it will not be absolutely certain for me tomorrow that I seemed to see a blue patch today; my memories will not be beyond question.²⁶ It is hard to see what empirical propositions would qualify as evidence by the proposed test. Thus the very possibility of learning from experience is threatened.²⁷

²⁶ It would only exacerbate the problem to individuate propositions so that the present-tensed sentence "I seem to see a blue patch" expressed the same proposition at different times, for even if such a proposition is certain and so true now, it will be false and so uncertain in the future.

²⁷ Jeffrey conditionalization is a way of modifying probabilities as a result of experience without assigning evidential propositions probability 1 (Jeffrey 1983, pp. 164–83). It does not address the question "What is S's evidence?", and gives no way of making sense of the idea that S's degrees of belief are coherent but out of line with S's evidence; indeed, Jeffrey says that he is concerned with the causes of belief, not the reasons for it (pp. 184–5). Thus Jeffrey conditionalization does not solve the present problem. Jeffrey motivates his proposal by an example

The model assumes that probabilities change only by conditionalization on new evidence. This is to assume that evidence can be added but not subtracted over time. The assumption is obviously false in practice, because we sometimes forget. But even if the model is applied to elephants, idealized subjects who never forget, the assumption that evidence cannot be lost is implausible. On any reasonable theory of evidence, an empirical proposition which now counts as evidence can subsequently lose its status as evidence without any forgetting, if future evidence casts sufficient doubt on it; on $E = K$, this process is the undermining of knowledge.

The underlying problem is that the model does not keep separate track of evidence and probabilities, even though it requires an exogenous notion of new evidence. Within a broadly Bayesian framework that does keep separate track of evidence and probability, the evidential probability of an hypothesis h for S at a time t can be identified with $P(h|e_t)$, the probability of h conditional on e_t , the conjunction of S 's evidence at t . P is a fixed initial probability distribution; its justification is far from unproblematic, but Bayesian accounts of evidential probabilities require such a probability distribution in any case.²⁸ If evidence is lost between t and a later time t^* , then e_{t^*} does not entail e_t , so $P(h|e_{t^*})$, the later evidential probability of the hypothesis, can be less than 1 even though $P(h|e_t)$, its earlier evidential probability, was 1. In particular, $P(e_t|e_{t^*})$, the later evidential probability of the earlier evidence, can be less than 1. One still conditionalizes, but each time one conditionalizes the initial probability distribution on the totality of old and new evidence. By contrast, on the standard Bayesian model, one conditionalizes the probability distribution resulting from previous conditionalizations on the old evidence on the new evidence. The alternative proposed here is consistent with any reasonable theory of evidence. In particular, it is consistent with $E = K$. Trivially, on this combination of views, the evidential probability of anything known is 1. In that sense, it is evidentially certain, but only because certainty is measured relative to knowledge. There is no guarantee that what is certain today will be certain tomorrow.²⁹

in which no proposition expressible in the English language encapsulates S 's perceptual evidence (p. 165); his argument is vulnerable to the considerations about perceptual demonstratives in §V. See Christensen (1992) and Williamson (1997) for further discussion.

²⁸ In particular, they require it to handle the problem of old evidence; see Williamson (1997).

²⁹ The proposed combination of views is elaborated in Williamson (1997).

It is, of course, equally consistent with the modified analysis of evidential probability that only knowledge of the present, or observational knowledge, counts as evidence.³⁰ But once it is recognized that evidence is not obliged to meet unusual standards of certainty, such restrictions on evidence look *ad hoc*. Although knowledge of the present or observational knowledge may be easier to obtain than some other kinds of knowledge, that is no reason against counting other kinds of knowledge as evidence, when we obtain them. If we believe that we know *p*, then we can use *p* in the ways in which we use evidence; if our belief is true, then we are right to use *p* in those ways. It does not matter what kind of proposition *p* is; as Austin said, “any kind of statement could state evidence for any other kind, if the circumstances were appropriate” (1962, p. 116). All knowledge is evidence.

VIII

The present case for $E = K$ is now complete. If evidence is what justifies belief, then knowledge is what justifies belief. But is all justified belief justified by evidence? Why can't experience itself, or practical utility, justify belief? Why can't belief sometimes *be* justified without being justified *by* anything at all? These large questions can receive only the most cursory treatment here.

The pragmatic justification of belief need not be by evidence; it is far from clear that belief could be epistemically justified except by evidence. Without any evidence at all, someone believes that her child somehow survived an air crash, and will one day return to her. The belief is the only thing which keeps her going; without it, she would commit suicide. Perhaps it is on balance a good thing that she has the belief, and in that sense the belief is justified. But this is not the sense of “justified” in which justified belief appeared to have marginalized knowledge within epistemology. The latter sort of justification aims at truth in a way—admittedly hard to define—in which the former sort does not. It is far from obvious that any belief is justified in the truth-directed sense without being justified by evidence. It appears otherwise when evidence is conceived narrowly, for

³⁰ For Maher, “*E* is evidence iff *E* is known directly by experience” (1996, p. 158; he relativizes “directly” to a set of propositions, cf. pp. 160–1). On his view, if I know *e*, that a substance *S* dissolved when placed in water, and deduce *h*, that *S* is soluble, thereby coming to know *h*, then *e* but not *h* is evidence for me (p. 158). On the present view, both *e* and *h* are evidence for me, but *h* is not independent evidence for *e*, since in the circumstances I cannot know *h* without knowing *e*. If someone now tells me that *S* is salt, *h* may become independent evidence for *e*.

then the evidence looks too scanty to justify all the beliefs which are in fact justified. But if anything we know can be evidence to anchor a chain of justification, as $E = K$ implies, then evidence may suffice for all truth-directed justification. If we are aiming at the truth, we should proportion our belief to the evidence.

$E = K$ supports the plausible equation of truth-directed justification with justification by evidence, and therefore with justification by knowledge. On this view, if truth-directed justification is central to epistemology, so too is knowledge.³¹

Department of Philosophy
University of Edinburgh
Edinburgh
UK

TIMOTHY WILLIAMSON

REFERENCES

- Achinstein, Peter (ed.) 1983: *The Concept of Evidence*. Oxford: Oxford University Press.
- Austin, J. L. 1962: *Sense and Sensibilia*, G. J. Warnock (ed.) Oxford: Oxford University Press.
- Carnap, Rudolf 1950: *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Christensen, David 1992: "Confirmational Holism and Bayesian Epistemology". *Philosophy of Science*, 59, pp. 540–57.
- Craig, Edward 1990: "Three New Leaves To Turn Over". *Proceedings of the British Academy*, 76, pp. 265–81.
- Davidson, Donald 1986: "A Coherence Theory of Truth and Knowledge", in E. LePore (ed.), *Truth and Interpretation*, Oxford: Blackwell, pp. 307–19.
- Earman, John 1992: *Bayes or Bust?* Cambridge, Mass.: MIT Press.
- Earman, John 1993: "Underdetermination, Realism and Reason". *Midwest Studies in Philosophy*, 18, pp. 19–38.
- Gettier, Edmund 1963: "Is Justified True Belief Knowledge?". *Analysis*, 23, pp. 121–3.
- Glymour, Clark 1980: *Theory and Evidence*. Princeton: Princeton University Press.
- Grice, Paul 1989: *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.

³¹ Thanks for helpful comments to the anonymous referees and Michael Ayers, Alexander Bird, John Campbell, Neil Cooper, Keith DeRose, Olav Gjelsvik, Keith Lehrer, Michael Martin, Hugh Mellor, Peter Milne, Nicholas Nathan, Stathis Psillos, John Skorupski and others in audiences at Edinburgh, Oslo and Oxford.

- Hempel, Carl 1965: *Aspects of Scientific Explanation and other Essays in the Philosophy of Science*. New York: The Free Press.
- Howson, Colin and Urbach, Peter 1993: *Scientific Reasoning: The Bayesian Approach*. 2nd ed. Chicago: Open Court.
- Jeffrey, Richard 1983: *The Logic of Decision*. 2nd ed. Chicago: University of Chicago Press.
- Kaplan, Mark 1985: "It's Not What You Know That Counts". *Journal of Philosophy*, 82, pp. 350–63.
- Lipton, Peter 1991: *Inference to the Best Explanation*. London: Routledge.
- Maher, Patrick 1996: "Subjective and Objective Confirmation". *Philosophy of Science*, 63, pp. 149–74.
- McDowell, John 1994: *Mind and World*. Cambridge, MA: Harvard University Press.
- Millar, Alan 1991: *Reasons and Experience*. Oxford: Clarendon Press.
- Moser, Paul. 1989. *Knowledge and Evidence*. Cambridge: Cambridge University Press.
- Peacocke, Christopher 1992: *A Study of Concepts*. Cambridge, MA: MIT Press.
- Peirce, C. S. 1932: *Collected Papers of Charles Sanders Peirce, vol. 2: Elements of Logic*, C. Hartshorne and P. Weiss (eds.). Cambridge, Mass.: Harvard University Press.
- Quine, W. V. 1969: "Epistemology Naturalized", in *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- Stampe, Dennis 1987: "The Authority of Desire". *The Philosophical Review*, 96, pp. 335–81.
- Unger, Peter 1975: *Ignorance: A Case for Scepticism*. Oxford: Oxford University Press.
- Williamson, Timothy 1994: *Vagueness*. London: Routledge.
- 1995: "Is Knowing a State of Mind?". *Mind*, 104, pp. 533–65.
- 1996: "Cognitive Homelessness", *Journal of Philosophy*, 93, pp. 554–73.
- 1997: "Conditionalizing on Knowledge". Typescript.
- Wright, Crispin 1991: "Scepticism and Dreaming: Imploding the Demon". *Mind*, 100, pp. 87–116.

