

WILLIAMSON ON EVIDENCE AND KNOWLEDGE

JIM JOYCE

The University of Michigan

Timothy Williamson's *Knowledge and its Limits* is a treasure trove of philosophical riches. By reversing the usual explanatory order and putting "knowledge first" in epistemology (p. v) Williamson reconceives the relationships between knowledge, belief, justification and evidence. While there is much to be said about Williamson's ambitious project, I will focus on his imaginative account of *evidence* (Ch. 9). Most epistemologists explain knowledge in terms of evidence: knowledge requires justification, and justification is a matter of having sufficient evidence for one's beliefs. This approach assumes that evidence can be understood independently of knowledge. Williamson denies this, and analyses evidence in terms of knowledge. He defends the " $E = K$ thesis" that "knowledge, and only knowledge, constitutes evidence. . . . S 's evidence [is] S 's knowledge, for every individual or community S in any possible situation" (p. 185).

While $E = K$ does capture some aspects of our concept of evidence, I will argue that it fails to secure others. Williamson ends up characterising an austere 'high church' evidence that leaves out many core features of our cognitive and social practices: among them, the fact that evidential status is generally a matter of degree, the fact that we often speak of a person's evidence when presenting *her own* subjective reasons for believing what she does, and the fact that what counts as evidence varies with context.

Let me begin by saying that I agree with Williamson on a number of points of methodology and substance. First, he correctly recognises the central role for a theory of evidence in epistemology (pp. 188–190). Second, he sees that a theory of evidence should be judged on the basis of its ability to explain core features of our evidential practices. Third, he is right that only something *propositional* can do the job that evidence does. (For those tempted by the idea of non-propositional evidence I can recommend no better tonic than Section 9.6 of *Knowledge and its Limits*.) Finally, Williamson is one of the few authors who explicitly recognises that a proposition e 's ability to serve as evidence for an hypothesis h depends both on e 's relationship to h and on e 's own epistemic status. This divides the theory of evidence into what I call a *theory of evidential relevance*, which seeks to determine whether and how much " e speaks in favour of h ", and a *theory of evidential status*, which specifies the sort of "creditable standing" that e must have in order to count as evidence for anything (p. 186). While many epistemologists focus on only part of the job, Williamson addresses both issues.

Williamson's theory of evidence can be stated in two theses:

Status. A proposition e has the status of evidence for a person S just in case S knows e .

Relevance. e is (positively) evidentially relevant to h for S just in case the objective evidential probability of h conditional on e and all of S 's other evidence exceeds the objective evidential probability of h conditional on S 's other evidence alone.

The first principle identifies a privileged class of S 's beliefs that constitute her evidence. The second determines which further beliefs are (objectively) justified in light of this evidence. If e_{total} is the conjunction of all propositions S knows, then the degree to which S is justified in believing p is just p 's objective probability conditional on e_{total} . If S is epistemically rational she will seek to "proportion her beliefs to her evidence" by adjusting her level of confidence in each proposition to its probability conditional on e_{total} . While I have misgivings about *Relevance*,¹ my focus here will be on *Status*.

For Williamson, evidence is whatever it is that justifies beliefs in a truth-directed sense (p. 185). Accordingly, S 's belief that e has evidential standing only if S can legitimately offer e as a consideration in favour of believing another proposition h . Notice, however, that the mere fact that S infers e from h , or takes e to be a good reason for h never suffices to make e a part of S 's evidence. Unless e meets the standards for evidence it cannot figure into genuine justifications (although it might figure into what S mistakes for justifications). According to $E = K$ the relevant standard is that S know e .

While I grant that all knowledge is evidence, requiring all evidence to be knowledge sets the bar too high. As I will argue, $E \subseteq K$ entails that evidential status is a categorical, non-contextual property. This, I shall claim, is at odds with our common conception of evidence, according to which evidential status is *gradational* and *contextual*.

1. Categorical versus Gradational Evidence

In arguing for $E \subseteq K$, Williamson imagines that S has just observed 1000 red balls drawn in consecutive random sampling with replacement from an urn that contains red or black balls in unknown proportion. A 1001st draw has been made and the ball is in fact red, but its colour has yet to be revealed. Intuitively, S has excellent evidence for believing that the last ball is red, but she does not *know* this until she sees it. Williamson maintains that the proposition Red_{1000} (= the 1000th draw is red) has the status of evidence, while Red_{1001} does not. To show this, he considers two false hypotheses

- h : The first 1000 balls drawn were red, but the last was black.
- h^* : The first ball drawn was black, but the last 1000 were red.

1. I am very sceptical of the notion of objective evidential probability, and I think Williamson, like many others, runs together two distinct ways of thinking about evidential relevance. For discussion of the second point, see my *The Foundations of Causal Decision Theory* (Cambridge University Press, 1999), pp. 203–212.

He writes,

It is natural to say h is consistent with S 's evidence and h^* is not. In particular, it is consistent with S 's evidence that draw 1001 was black; it is not consistent with her evidence that draw 1 was black. Thus, S 's evidence does not include the proposition that draw 1001 was red. Why not? After all, by hypothesis S has a justified true belief that it was red. The obvious answer is that S does not know that draw 1001 was red; the unsatisfied necessary condition for evidence is knowledge. (p. 201, with incidental changes in notation)

Williamson is claiming that, whatever S 's evidence consists in, h is consistent with it but h^* is not. This makes Red_1 part of her evidence but excludes Red_{1001} . Since S has justified true beliefs in both propositions, Red_1 's status as knowledge must be what makes the difference.

This last point is debatable. There are other relevant differences between Red_1 and Red_{1001} that might explain why only the first is evidence. For example, because S 's belief that Red_1 is justified on the basis of direct observation and her belief that Red_{1001} is justified by inference from observations, the two will be *resilient* under different learning experiences.² Imagine that balls will be drawn from the urn another 100 times, and consider the proposition e^* that 95 of these balls are black. Learning e^* will not change S 's level of confidence in Red_1 or alter its status as knowledge, but it will greatly reduce her confidence in Red_{1001} . This difference seems relevant to the evidential status of the two beliefs. One might suggest, therefore, that the line between evidence and non-evidence can be drawn on the basis of differences in the resilience of beliefs in the face of information about future draws.

This objection is probably not decisive. Williamson can plausibly respond that these differences stem from the fact that Red_1 is known while Red_{1001} is not. Still, even if resilience is not what explains the difference in evidential status, Williamson still needs to convince us that *only* knowledge will.

At any rate, I have a more serious worry about Williamson's argument. He begins from the premise that, whatever S 's evidence consists in, h is logically consistent with it but h^* is not. Since h entails Red_1 & $\neg Red_{1001}$ while h^* entails $\neg Red_1$ & Red_{1001} , he concludes that Red_1 is a part of S 's evidence but Red_{1001} is not. The problem is that Williamson's initial premise is only plausible on a *categorical*, as opposed to a *gradational*, conception of evidence. On a categorical conception, the question of whether a belief has evidential standing has an unqualified yes/no answer. A person's evidence can be specified as a set of believed propositions, each member of which has the same, univocal status. On a gradational view, one speaks not of evidential status *tout court*, but of the *degree* to which a believed proposition counts as evidence. Evidential status

2. A belief that h is resilient with respect to e^* to the extent that learning e^* does not greatly alter the believer's level of confidence in h . See B. Skyrms, *Pragmatics and Empiricism* (Yale University Press, 1984), pp. 54–56.

falls along a spectrum that ranges from the best sort of evidence, through intermediate grades, to beliefs that are not evidence at all. Even though Williamson construes many epistemological concepts in gradational terms (belief, evidential relevance, and justification), he adopts a strictly categorical view of evidential status, and uses it to sharply distinguish h from h^* . Whatever the evidence is, he argues, it definitively rules out h^* by being inconsistent with it, and definitively permits h by being consistent with it.

Gradationalists deny this. On any reasonable analysis, both h and h^* are extremely unlikely in light of S 's evidence. Hence, *both* conflict with S 's evidence to a high degree, but neither is definitively 'ruled out' by it. Red_1 and Red_{1001} are more alike in epistemic status than Williamson's categorical analysis allows in that each contradicts a hypothesis that is highly unlikely in light of the evidence. More generally, the gradationalist idea is that propositions are almost never categorically ruled out by a person's evidence: since evidential status is a matter of degree, so too is conflict with evidence. Williamson's argument fails because it incorrectly assumes that the evidence definitively permits h and precludes h^* .

Those wedded to categorical thinking may think they smell a fallacy. When I say that h and h^* are unlikely in light of the evidence doesn't that mean that their probabilities are low when conditioned on all S 's evidence, so that $\mathbf{P}(h/e_{\text{total}})$ and $\mathbf{P}(h^*/e_{\text{total}})$ are close to zero? If so, then there is a clear categorical distinction between h and h^* : since h is consistent with e_{total} but h^* is not, we have $\mathbf{P}(h/e_{\text{total}}) > 0 = \mathbf{P}(h^*/e_{\text{total}})$. Even gradationalists must recognise the difference between a positive probability and no probability at all!

This reply misses the point of the gradationalist proposal, and so begs the question against it. Anyone who thinks evidential standing is a matter of degree will deny that S 's total evidence can be captured in a conjunction like e_{total} . To characterise S 's total evidence it is necessary both to specify a set of propositions she believes, and to assign a degree of evidential status to each of its elements. Since these degrees of status will not all coincide, no one proposition can capture all S 's evidence. Consequently, h 's probability in light of the evidence cannot be expressed as $\mathbf{P}(h/e_{\text{total}})$, and likewise for h^* . In general, when a hypothesis q is logically incompatible with some proposition e , the degree to which e rules q out will be an increasing function of e 's evidential status (and a decreasing function of its probability given e). For q to be definitively ruled out it must be inconsistent with some proposition e of *maximal* evidential status.

This might seem to play into Williamson's hand. Nothing, it might seem, can be better evidence than knowledge. Items of knowledge, like Red_1 , have maximal evidential status, while propositions like Red_{1001} have some lesser status. Thus, since h^* contradicts a proposition of maximal status it really is ruled out.

The problem with this line of reasoning is that most items of knowledge, including Red_1 , do *not* have maximal status. While I have no theory of evidential status to offer, I do claim that, minimally, e 's status as evidence improves with the subject's level of confidence in its truth and its resilience in the face of learning experiences that add true beliefs to her evidence. The *best* sort of evidence must be believed with certainty and be completely resilient under

the learning of truths. Since known propositions are confidently believed and highly resilient to the truth, they are always evidence of high quality. They are not, however, always evidence of the highest quality, at least not on a view that, like Williamson's, rightly recognises that knowledge does not require subjective certainty and that it can be overturned by future evidence.

That said, gradationalists can recognise a difference between Red_1 and Red_{1000} , and can even explain why the former seems like evidence while the latter does not. Red_1 is clearly *better* evidence than Red_{1001} since it is likely to be more resilient to the truth and more confidently believed. Indeed, the important element of truth in $E = K$ is that knowledge is always a better calibre of evidence than non-knowledge. A gradationalist can use this fact to explain why it is typically inappropriate for S to *offer up* Red_{1001} as part of her evidence for any conclusion.

We must distinguish the epistemological question of what beliefs *can* be used as evidence *per se* from the pragmatic question of what beliefs it is permissible to *cite* as evidence in a given conversational context. The rule is that one should not offer up a proposition as evidence when one can cite propositions of higher evidential status that support the conclusion equally well. Given that Red_1, \dots, Red_{1000} will justify any conclusion as well as Red_{1001} does,³ S should never cite the latter when she is in a position to cite the former. That said, Red_{1001} does have some level of evidential status since there are circumstances in which she would be able to cite it. If, say, S forgets Red_1, \dots, Red_{975} , and so does not recall her justification for Red_{1001} (but still remembers that she had one), then it would be permissible for her to cite Red_{1001} as a reason for believing Red_{1002} .

Gradationalism also handles a more theoretical argument that Williamson offers for $E \subseteq K$. He writes,

If evidence required only justified true belief, or some other good cognitive status short of knowledge, then a critical mass of evidence could 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. (p. 201)

To see his concern, suppose that 998 of the first 1000 draws were red. Let X be the proposition that states this evidence. Since the subject knows X she believes that there is a small probability of getting a black ball on any future draw. Now, if Red_{1001} counts as evidence when it is justified on the basis of X , then (Red_{1001} & Red_{1002}) is evidence too since it can be justified on the basis of

3. This is because Red_{1001} is not what Williamson refers to as *independent* evidence (p. 207). Williamson admits that even an item of knowledge may be evidence, but not independent evidence, in which case it would be inappropriate to *cite* it as one's evidence. The same goes, I am saying, for propositions of less exalted evidential status.

$X \& Red_{1001}$. Likewise, Red_{1003} is evidence because it can be justified on the basis of $X \& Red_{1001} \& Red_{1002}$. Continuing in this way, it follows that $X \& Red_{1001} \& Red_{1002} \dots \& Red_n$ will be evidence for any natural number n , which is absurd given that this conjunction becomes increasingly unlikely as n grows.

Williamson prevents such chain reactions by restricting a person's evidence to what she knows. Whereas true beliefs based on inductively strong inferences from justified true beliefs are justified, true beliefs based on inductively strong inferences from knowledge need *not* be knowledge. Thus, in any chain of justifications there will be a point at which knowledge ends. According to $E \subseteq K$, this is also where evidence ends and something less than evidence begins.

This is overkill. 'Chain reactions' only threaten categorical conceptions of evidence. Gradationalists will simply say that evidential status diminishes as one goes further down the chain of inductive inferences, eventually to vanish completely. The evidential status (and subjective probability) of $X \& Red_{1001} \& Red_{1002} \dots \& Red_n$, for example, will slowly decline as n grows until at some (perhaps vague) point it becomes so unlikely that it ceases to be evidence altogether.

2. Must Evidential Beliefs be True?

Williamson's view of evidence as that which justifies belief in the truth-directed sense is ambiguous between two senses of 'justifies'. When inquiring into a person's justification for some belief we might be interested in *her own* reasons for holding it, or we might want to know which of her reasons would pass muster with an objective third party. Williamson is particularly concerned with the latter question. He tends to focus on those aspects of our evidential practices that involve presenting evidence for public consumption or evaluating evidence from an objective, third-person perspective. This is part of the reason why he insists that evidential beliefs must be true. In a telling passage he writes,

If e is evidence for h then e must be true. . . . If e is not true then at most a counterfactual holds: if e had been true, then e would have been evidence for h . 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* 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. (pp. 201–202)

Notice how Williamson speaks of 'the evidence' *per se*. He seems to mean those parts of the proceedings that a juror (or a well-informed third party) can *safely* use to justify a verdict. He is not, however, speaking of any juror's *own* reasons for believing the accused innocent. The witness's testimony, taken by itself, provides no reason at all for the juror to believe this. Barring *recherché* cases, a juror will only regard the witness's assertion as evidentially relevant to the guilt of the accused if she holds the further (false) belief that the witness is reliable. Thus, if we want to present the reasons that a juror might have for

believing the accused innocent then we have no option but to specify some false beliefs that she holds.

I grant that there is one sense of 'evidence' in which a lying witness provides jurors with no evidence at all. When I evaluate your reasons for p , *with an eye to the question of whether you, or I, or anyone else should believe p on the basis of those reasons*, I dismiss any premises in your justification that I know to be false. When I evaluate the fitness of your reasons 'for public consumption' I may reasonably impose a standard of truth because by endorsing your reasons for p I imply that others can legitimately believe p on the basis of those reasons. This is not, however, the only, or even the main, way to evaluate reasons for belief. Often when we speak of a person's evidence our aim is to *make sense* of (or criticise) her beliefs by presenting her subjective reasons for believing what she does. We try to give a *rationalising explanation* (or critique) of some things she believes, usually by citing other things she believes. Here the standard of truth is inappropriate. Given that the witness has convinced the juror, the fact that he was lying does nothing to impugn the rationality of her verdict. It is quite true that her verdict would have been more *reliably* justified had she not believed the false testimony. Still, she did believe it, and used this as the basis for her verdict. This makes the belief *misleading* evidence, not *non-evidence*.

It is open to Williamson to suggest that we rationalise the juror's verdict by showing how it makes sense in light of what she *mistakenly took her evidence to be*, but we should resist this suggestion. First, it requires us to speak in odd ways. When the juror learns, after the trial, that the witness lied, her excuse for acquitting should be that the witness presented false evidence, not that his false statement misled her as to what her evidence was! Second, this suggestion mandates an *error theory* of 'rationalising' uses of evidence talk. A person speaks falsely when he aims to rationalise the juror's decision by saying (as we ordinarily do) that she had evidence that the defendant was asleep. Yet, by asserting this falsehood he achieves his aim by successfully communicating that the juror believed the defendant was asleep, took this to be a good reason for acquitting, and used this as the reason for her verdict. It is simpler, and better methodology, to interpret people as speaking truly whenever we can. To do this, in the present case, we need only recognise that 'evidence' sometimes applies to those of a subject's beliefs that she *regards* as especially reliable and uses to justify other beliefs.

3. Proportioning Beliefs to Evidence

Williamson also thinks that evidential beliefs must be true because subjects would otherwise have no reason to proportion their beliefs to their evidence. He is clearest on this point in his criticism of the ' $E = B$ thesis', the idea that a person's evidence is just the set of propositions she believes. $E = B$, Williamson writes, permits one to "manufacture evidence for one's favourite theories by manipulating oneself into a state of certainty about appropriate propositions. . . . [This] does not capture the spirit of the injunction to proportion one's belief to one's evidence" (p. 222). The claim is that if one's

evidence is not composed of truths, then one can satisfy the ‘proportioning requirement’ by fitting one’s evidence to one’s beliefs rather than by fitting one’s beliefs to one’s evidence.

I am not convinced that $E = B$ condones such manipulation. An example will make the point. Suppose I desperately want to believe in god, but that I am a Cliffordian who does not believe anything on the basis of insufficient evidence. I now have no evidence for thinking that god exists. In particular, while I recognise that certain commonly reported miracles would be excellent evidence for god’s existence, I am sure that no such miracle has occurred. I also know, however, that if I hear a sermon by the persuasive Monsignor Menteur I will come out believing that there have been many such miracles. Suppose, despite my doubts about Menteur’s veracity, I attend the sermon so as to bring myself to hold beliefs that, given $E = B$, will provide me with strong evidence for theism. Williamson is suggesting that $E = B$ condones such an act, but I do not see how it does. The conjunction of $E = B$ with the requirement to proportion my beliefs to my evidence does require me to believe in god *once I come to believe in miracles*. This, however, is not where my mistake lies: if I believe in miracles then I *should* believe in god since (*ex hypothesi*) the latter is the best explanation for the former. My error was in choosing to attend the sermon when I was sure that doing so would lead me to believe falsehoods. Before the sermon my evidence (i.e., what I believe) is that there have been no miracles and that Menteur is a liar. Despite this, I knowingly take a step that will lead me to hold opinions that directly conflict with my evidence. *This* is where I violate the proportioning requirement; I do not violate it after hearing the sermon. So, while $E = B$ does enjoin us from criticising a person for drawing conclusions on the basis of evidence arrived at via manipulation, it *requires* us to criticise the manipulation itself.

We should also not be too quick to grant that proportioning one’s beliefs to one’s evidence only has a point when “it is way of adjusting them to the truth” (p. 202). In defending this claim Williamson writes:

Once it is granted that all propositional evidence is true . . . adjusting one’s beliefs to the evidence has an obvious point. It is a way of adjusting them to the truth. Although true evidence can still support false conclusions, it will tend to support truths. The maxim ‘Proportion your beliefs to your evidence’ requires more than mere internal coherence of one’s belief system; it does so because evidence must be true. (p. 202)

Everyone will agree that there is a compelling rationale for proportioning beliefs to the evidence when it is true, but we should resist the thought that this is the *only* legitimate way to vindicate the proportioning requirement.

We need to proportion our beliefs to the evidence because propositions do not wear truth-values on their sleeves. We rely on evidence because (a) it is often easier for us to determine which of our beliefs have the status of evidence than it is to directly determine the truth-value of an arbitrary belief, and (b) there is reason to think that beliefs with this status are more likely true than garden-variety beliefs. In an ideal world, it would be easy to detect

evidential status, and there would be an ironclad guarantee that evidential beliefs are true. Descartes thought the world was ideal in this way. He thought a sufficiently careful agent could tell whether a belief was 'clear and distinct' and that a non-deceiving god would vouchsafe the truth of such beliefs. Likewise, some empiricists have held that the perceptual character of a belief is manifest to its holder, and that perceptual beliefs are surely true. But these are pipe dreams; there is no hope of finding a class of infallible, easily identifiable beliefs that justify any substantial number of other beliefs. We thus seem forced to choose between (a) and (b).

At one end of the spectrum sits $E = B$, which has us believe whatever seems true. This makes evidential status easy to assess, but evidential beliefs have no special claim to truth. Views like Williamson's sit at the other end of the scale. By advocating $E = K$, and recognising that knowledge is not an epistemically transparent state (pp. 93–113), he makes the connection between evidence and truth very tight, but makes it hard for a subject to know a belief's evidential status (since they would have to know what they know).

The error in both positions lies in thinking that there is a single choice between (a) and (b) that applies across all situations. Evidential status is *contextual*. Justification clearly depends on context, but it is tempting to think that this only has to do with the *amount* of evidence required to justify a belief. This is wrong: evidential *status* is also contextually constrained. In different contexts we have different practical and epistemic aims, which makes it appropriate to employ different standards of evidence. In particular, we are often willing to exchange some reliability to attain accessibility. It is, I think, *always* required that evidential beliefs be more reliable, as a class, than garden-variety beliefs (which is why $E = B$ is wrong), but it is not usually required that they all be true or known. We often forgo some reliability to make it easier for subjects to tell which of their beliefs constitute their evidence. This is why (a modest) empiricism is so plausible as a theory of evidence. We are not infallible about which of our beliefs are directly caused by observation, but we are *fairly* good at detecting this. Likewise, while beliefs caused by observations are not entirely reliable they are *more* reliable than most beliefs. Thus, it often makes sense to invest all empirical beliefs with evidential status whether or not they are *all* true. For most purposes it suffices that a sufficiently high proportion of them are true, with the 'sufficiently' varying from context to context depending on how important it is for us to be able to discern evidential status. Note, however, that even when we invest all empirical beliefs with evidential status, we still have a compelling rationale for proportioning our other opinions to them. After all, we have fairly reliable access to our empirical beliefs, and these beliefs are more likely to be true than others. We do not need anything nearly as strong as $E = K$ to make this argument.

4. A Moral

I hope to have conveyed at least some of the richness of our multifaceted concept of evidence. Evidential status typically is a matter of degree. Sometimes

we attribute this status to beliefs in order to present a person's subjective reasons without endorsing them. Other times we use such attributions to offer up reasons for public consumption. What we count as evidence varies with context. Given all this richness, it should be clear that no single identity like $E = K$ will ever capture all there is to say about evidence.