

## Chapter One

# Should Knowledge Come First?

What is the place of knowledge within epistemology? This is a methodological question of first importance. Timothy Williamson argues that knowledge should come first. Methodologically, this means that we shouldn't expect an informative analysis of knowledge in terms of belief, truth, plus some further set of necessary, non-circular conditions. Nor should we accept an impoverished conception of evidence or epistemic normativity that would be acceptable to a skeptical interlocutor. Rather, we should begin doing epistemology by focusing on epistemic access itself, and knowledge itself is the most natural candidate for this access. We can then understand other notions such as evidence and justification in terms of knowledge, which is a factive, world-involving state. In opposition to Williamson's view, Trent Dougherty and Patrick Rysiew argue for an *experience-first* approach to epistemology. Experience is where it all begins – indeed, where it *must* begin. Understanding experience to include perceptual awareness, introspective awareness, and rational insights, Dougherty and Rysiew argue that it is our basic evidence – that in virtue of which all else is made evident – that must come first in any intellectual project. On this approach, it is experience that ultimately justifies belief, guides rational thinkers, signifies truth, and settles disputes.

## Knowledge First

*Timothy Williamson*

Epistemology matters.<sup>1</sup> It is not just fascinating in itself; its concerns arise in every serious form of human inquiry. How much does this evidence support that claim? Is this alleged evidence just another claim, itself in need of evidence? Is the category of evidence even relevant here? Such disputes are often resolved by the internal standards of the inquiry. But sometimes they go deeper, raising issues adequately addressed only at a level of abstract reflection characteristic of philosophy. It is all the more striking

that in the second half of the twentieth century epistemology acquired a reputation (of which many practitioners still seem unaware) as sterile and inward-looking. Some of that was just the usual complaint about analytic philosophy, that it is boring and inaccessible to the untrained. But much of it came from other analytic philosophers, well acquainted with analytic epistemology, who still found it small-minded and old-fashioned in relation to its proper task.

One aspect was the post-Gettier industry of trying to analyze knowledge in terms of belief, truth, and non-circular further conditions. The increasingly gerrymandered definitions were obvious signs of a degenerating research program. Most of them, if correct, seemed to make knowledge too grue-like to be worth analyzing. But in any case they succumbed one after another to counterexamples. Moreover, no prior reason to expect knowledge to have such an analysis withstood scrutiny. Evidence accumulated that few if any words of natural language are understood by means of complex definitions. Nor does the nature of knowledge provide any clear evidence that it has such an analysis. That knowing entails believing truly does not show that for some non-circular condition C, knowing is equivalent to believing truly and meeting C. In consequence, the project of analyzing knowledge has lost its importance to analytic epistemology.

Another common charge is that epistemology is obsessed with the problem of skepticism, wasting its time on an imaginary opponent. The charge is partly unfair. Not only is analytic epistemology not mainly concerned to answer skepticism, the issue usually arises when apparently legitimate ways of criticizing ill-founded beliefs, applied more systematically, lead to skeptical conclusions. Epistemologists did not willfully introduce the skeptical tendency from outside; it was already in us. That does not mean that the skeptic is right, just that we should consider the issue.

However, skepticism plays a further role in defining the framework of much analytic epistemology. Many practitioners take the key epistemological notion to be not knowledge but *justification*, in a specifically epistemic sense (a qualification henceforth understood). They typically explain the difference between knowledge and justification by contrasting an everyday situation with a skeptical scenario in which everything appears the same to one but one is a brain in a vat. In the good case, one knows that one has hands. In the bad case, since one lacks hands, one does not know that one has them; one merely appears to oneself to know that one has hands. By contrast, such epistemologists claim, in the two cases one is justified to exactly the same degree in believing that one has hands. They intend this notion of justification for general epistemological purposes, not only for handling skepticism. In particular, they treat it as the appropriate normative standard for criticizing beliefs.

That view casts appearances in a leading role: justification supervenes on them. But what is epistemologically so special about appearances? One answer is that if cases appear the same to one, then one cannot discriminate between them, and cannot fairly be criticized in one case with respect to a feature it does not share with the other, since one cannot discern it. The underlying principle is that justification is exactly the same in cases indiscriminable to the subject. But that principle falls to an objection from the non-transitivity of indiscriminability. Consider a long sorites series of cases  $\alpha_0, \dots, \alpha_n$ , where the subject  $\alpha_i$  and  $\alpha_{i+1}$  are so similar in appearance that they are indiscriminable ( $i=0, \dots, n-1$ ), but  $\alpha_0$  and  $\alpha_n$  are so different in appearance that they are easily discriminable. By the principle, justification is exactly the same in  $\alpha_i$  and  $\alpha_{i+1}$  ( $i=0, \dots, n-1$ ). Therefore, by the transitivity of exact sameness in a given respect,

justification is exactly the same in  $\alpha_0$  and  $\alpha_n$ . But that is absurd, for justification differs between easily discriminable cases for some proposition. To avoid such objections, one must stick to the weaker principle that when appearances are exactly the same, so is justification. But why predict so tight a link between justification and appearances, if it is not mediated by indiscriminability?

Another idea is that in some sense one is always fully *acquainted* with present appearances to one, even if one cannot discriminate slight differences between them. This may help to explain the epistemological privileging of appearances (if one is never fully acquainted with anything else). The upshot is some sort of phenomenal conception of evidence. But by now, alarm bells should be ringing. The idea of full acquaintance with appearances has no basis in contemporary psychology or cognitive science. Nor has it much phenomenological plausibility. When I ask myself what I am acquainted with, the physical objects in front of me are far more natural candidates than their appearances. If I try to introspect or otherwise catch how things appear to me, I experience confusion, characteristic of embarking on an ill-defined task. Rather, full acquaintance with appearances is a wild postulate of a specific type of epistemological theory, one that requires something to be fully and unproblematically given to the subject to serve as the basis for justification. This is a barely updated Cartesianism: however vulnerable I am to doubt, ignorance, and error, something in me is clear and distinct.

To a depressing extent, epistemology has served as the refuge of an otherwise discredited philosophy of mind, supporting and supported by the definition of epistemic normativity in terms of a skeptical challenge. For justification was explained by the contrast between the ordinary case and the skeptical scenario. There may even be a distorting selection effect, by which those inclined to think along such lines are disproportionately drawn to, and rewarded in, epistemology rather than other branches of philosophy. It is Cartesianism that makes epistemology the starting point.

Suspicious of full acquaintance with sensory appearances, we might strip them out of the picture. What that leaves of the inner is a formal structure of beliefs, for which the norm of justification above reduces to mere internal coherence. Subjective Bayesianism is the best developed such view. Despite its mathematical virtues, it fails to make most distinctions of epistemological significance. It treats alike you and someone with the same credences but a radically different perceptual experience of the world.

The starting point of Cartesian epistemology is the comparison between the good case and the corresponding bad case. From a contemporary perspective, what they most obviously share is an internal microphysical state *S*. Consequently, they also share any mental state that supervenes on *S*. But *S* has no privileged *epistemological* status. Our internal microphysical states are typically unknown to us, and can only become known through arduous scientific investigation. Similarly, any mental state that supervenes on *S* has as such no epistemological privilege. It may be a state of depression that we cannot introspect ourselves to be in. What matters is whether we are aware of being in the given state, whether we *know* that we are in it.

Epistemologies that explicitly make knowledge a secondary phenomenon may nevertheless implicitly put it first, because they select the mental states to which they officially assign a privileged epistemological status for their supposed amenability to being known. They typically take that amenability for granted, rather than subjecting their crucial choice of starting point to open reflection. The less amenable the selected mental states turn out to being known, the more arbitrary becomes their promotion to

a privileged epistemological status. Once we give up Cartesian fantasies about the mind, we can recognize that no special sort of fact is automatically amenable to being known, although many sorts of fact often are known. Rather than seeking a domain to which we have privileged epistemic access, we should concentrate on epistemic access itself. By far the clearest explication of “epistemic access” is simply knowledge. Thus attempts to start epistemology with something much more internal than knowledge nicely illustrate the naturalness of starting with knowledge itself.

On one knowledge-first view, our total evidence consists of facts we know, irrespective of whether they are facts about our mental states. We are in no position to use facts we don't know as evidence. When we acquire new evidence in perception, we do not first acquire unknown evidence and then somehow base knowledge on it later. Rather, acquiring new evidence *is* acquiring new knowledge. That knowledge need not itself be based on further evidence, nor is it evidence for itself in some non-trivial way. But it is evidence for or against potential answers to questions to which we do not yet know the answer.

Equating evidence with knowledge helps reconnect epistemology with other fields. For one of the ways in which it marginalized itself was by depicting evidence as utterly unsuited to its role in science. The evidence for a well-confirmed scientific theory is typically a matter of public record. At least to a first approximation, it consists of facts intelligibly related to the theory and available to be known by anyone of suitable intelligence and training who takes the trouble to find out. It does not consist of facts about the present mental states of scientists or anyone else, facts that are no matter of public record and whose evidential relation to the theory itself has never been properly explained. Although the fact that a physicist believes a physical theory may raise its probability, the link requires an auxiliary sociological hypothesis and is hardly an evidential relation of the sort with which physics typically deals.

If our evidence is what we know, the evidence differs between the good and bad cases, contrary to what skeptics and many other epistemologists assume. For in the good case but not the bad, the subject's evidence includes the fact that she has hands. Although one can stipulate an alternative sense for the word “evidence” in which the evidence is the same, the challenge is to give epistemological significance to the new sense. The preceding reflections suggest that any idea that the two cases are evidentially equal is no basic insight but a product of misconceived epistemological theorizing. In particular, the fact that for all one knows in the bad case one is in the good case does not entail that one has the same evidence; it just means that for all one knows in the bad case one has the same evidence as in the good case. One is not always in a position to know whether one's evidence includes a given proposition. Although we might prefer a notion of evidence that does not work like that, we have no right to expect one.

A knowledge-first approach discourages trying to explain knowledge in terms of belief. We may even try the reverse, explaining belief in terms of knowledge. Here is a simple picture. Beliefs are the products of cognitive faculties whose function is to produce knowledge. When and only when all goes well, beliefs constitute knowledge. Even if something goes wrong, the belief may still be true, just as someone's scheme for getting rich may fail while they become rich by an unintended chain of events. Thus knowing is the successful state, believing the more general state neutral between success and failure. Knowing corresponds to doing, believing to trying. Just as trying is naturally understood in relation to doing, so believing is naturally understood in relation to knowing.

If justification is the fundamental epistemic norm of belief, and a belief ought to constitute knowledge, then justification should be understood in terms of knowledge too. Indeed, a belief is fully justified if and only if it constitutes knowledge. Although your belief that you have hands is fully justified, the corresponding brain in a vat's belief is not. But the brain in a vat has a good excuse for believing that it has hands, because, for all it knows, its belief that it has hands is justified, since, for all it knows, it knows that it has hands. Confusion between justifications and excuses undermines much talk of epistemic justification.

Some beliefs fall shorter of justification than others. In that respect we can grade beliefs by their probability on the subject's evidence, that is, on the subject's knowledge. A theory of evidential probability can be developed along such lines. It fills a gap between purely subjective probabilities, Bayesian credences ("degrees of belief"), and purely objective chances in indeterministic physics. When we ask how probable a theory is on our evidence, we want something less dependent on our doxastic state than a credence but more dependent on our epistemic state than a chance.

The most salient feature of knowing as the focus of epistemology is that it is a world-involving state. For it is factive: knowing that P, unlike believing that P, entails that P. Thus the state of the external environment constitutively constrains one's epistemic state. More specific factive epistemic states include perceiving that P and remembering that P. But even believing involves the world in another way. For the external environment constitutively constrains the contents of most intentional states; belief and knowledge are no exceptions. Croesus knew and believed that he was rich in gold. Despite being in the same internal microphysical states, Twin-Croesus on Twin-Earth neither knew nor believed that he [Croesus] was rich in gold. He had no knowledge or beliefs about Croesus, since he never had any suitable contact with him, however indirect. Nor did he know or believe that he [Twin-Croesus] was rich in gold. He wasn't. He was rich in another material, superficially like gold. Although many attempts have been made to define some sort of "narrow content" for Croesus and Twin-Croesus to share, all rely on deeply problematic assumptions. In any case, the intentional states that normally matter to us are broad states like believing that Croesus was rich in gold and wanting to be rich in gold oneself. Factive states involve the world twice over, in both their contents and their attitudes to those contents.

The ways in which an intentional state involves the world are not impurities. They are its point. The function of intentional states is to enable us to engage intelligently with the world. Take another example. Only those suitably related to Heloise can be in the state of loving her. The idea that the real core of loving Heloise is a mental state one can be in even if there is no Heloise looks more like a symptom of pathological self-absorption than a serious philosophy of mind. Without something loved or hated, there is at most an illusion of love or hate. Similarly, the connection to water is not accidental to desiring water or believing that there is water over there. The same goes for factiveness. A connection to the external environment is not accidental to the mental nature of seeing that it is raining, nor is a connection to past events accidental to the mental nature of remembering that it was raining. Misperceiving must be understood as a deviation from perceiving, and misremembering as a deviation from remembering. A neutral state that covers both perceiving and misperceiving is not somehow more basic than perceiving, nor is a neutral state that covers both remembering and misremembering more basic than remembering, for what unifies the various cases

of each neutral state is their relation to the successful state. Likewise for the generic factive state of knowing: a neutral state that covers both knowing and seeming to oneself to know is not somehow more basic than knowing, for what unifies the various cases of seeming to oneself to know is their relation to knowing. That whenever one is in such a broad state one is in some specific internal microphysical state too is no reason to postulate a corresponding narrow mental state.

An internal starting point for epistemology is false to the nature of mental life. In bracketing the differences between the good case and the skeptical scenario, the internalist approach to justification does not isolate a purely mental dimension; it merely ignores those aspects of the subject's awareness present in the good case but absent in the bad one. Much epistemology has been in denial about the depth of externalist developments in the philosophy of mind over recent decades, as though broad mental states could be analyzed into narrow mental states and their causal relations to the external world. But the postulated underlying layer of narrow mental states is a myth, whose plausibility derives from a comfortingly familiar but obsolescent philosophy of mind. Knowledge-first epistemology is a further step in the development of externalism.

A closely related contrast occurs in the philosophy of language. Truth-conditional, referential semantics is an externalist program. On such a theory, in a context of utterance the atomic expressions of a language refer to worldly items, from which the truth conditions of sentences are compositionally determined. Just like the contents of someone's intentional states, the truth-conditional semantics even of their idiolect is far from supervening on their internal microphysical states; it also depends on what external objects they are in causal contact with. Reference and truth, like knowledge, are matters of success, not of something neutral between success and failure. Insofar as there is a competing internalist program, it is conceptual role or inferential semantics, on which the meaning of an expression is something like its place in a web of inferential relations. Although inferential semantics can be given an externalist twist, in practice many of its proponents are motivated by internalist sympathies. On internalist inferential semantics, the inferential relations of an expression do not depend on what, if anything, it refers to, although there may be dependence in the opposite direction if reference is determined by a combination of inferential and external causal relations. Inferentialism faces grave problems of principle, for instance in separating patterns of inference that are to count as essential to the meaning of an expression from those that will count as accidental (a form of the analytic/synthetic distinction). Moreover, the internalist version has particular difficulty in establishing an adequate relation between meaning and reference. Even more striking is the disparity in practice between the success of referentialist truth-conditional semantics as a flourishing research program, pursued by both philosophers of language and linguists, which has provided invaluable insight into semantic phenomena in natural languages, concerning both overall structure and the behavior of specific expressions, and the lack of progress of inferentialism, which remains in a largely programmatic state. By that pragmatic criterion, referentialism beats inferentialism hands down. Internalism has proved to be an obstacle to new insights in the philosophy of language.

The fruitfulness of referential semantics is an encouraging precedent for knowledge-first epistemology, since both take as basic the central forms of success distinctive of their field: truth and reference in semantics, knowledge in epistemology. In the long run, knowledge-first epistemology too should be judged by its fruitfulness

as a research program, compared to its competitors. I have shown elsewhere how a knowledge-first methodology casts light on such matters as the nature of indiscriminability and the norm of assertion. Here, for reasons of space, I will focus on another issue: applications of epistemic logic in epistemology.

In recent years it has become clear that formal models of epistemic logic enable us to analyze some epistemic phenomena in a more disciplined, systematic, and rigorous way than we can achieve through qualitative description in ordinary prose, even eked out with epistemological jargon. We can sometimes learn far more about the target phenomena by mathematically exploring the consequences of a model than by the shaky reasoning and appeals to the “obvious” characteristic of so much traditional epistemologizing. Of course, the models usually involve simplifications and idealizations, just like mathematical models in natural science. Humans are no more logically omniscient than planets are point masses. As always, formal methods will not give good results when applied with bad judgment. It takes experience and skill to know which simplifications and idealizations are appropriate for a given problem. But epistemologists are gradually acquiring the relevant experience and skills.

What is the connection between epistemic logic and knowledge-first epistemology? In standard epistemic logic, the basic epistemic operator is  $K$ , read “The agent knows that ...,” and interpreted in each model by means of an accessibility relation of epistemic possibility between worlds. A world  $x$  is epistemically accessible from a world  $w$  just if for all one knows in  $w$  one is in  $x$ , that is, everything one knows in  $w$  is true in  $x$ . A formula  $KP$  is true in a world  $w$  just if the formula  $P$  is true in every world epistemically accessible from  $w$ . Since epistemic accessibility is required to be reflexive, every instance of the factiveness schema  $KP \rightarrow P$  is true in every world in every model. Often the framework is multi-agent:  $K$  has different subscripts for different agents, and each agent has their own accessibility relation. The truth condition for  $KP$  is not intended as an analysis of knowledge in independent terms, for epistemic accessibility is itself explained in terms of knowledge. Rather, the semantics simply decodes the information about knowledge conveniently encoded in the accessibility relation. This framework accords with a knowledge-first methodology.

However, the accord so far is rather superficial. We can equally well introduce a basic doxastic operator  $B$ , read “The agent believes that ...,” and interpreted in each model by means of an accessibility relation of doxastic possibility between worlds. A world  $x$  is doxastically accessible from a world  $w$  just if everything one believes in  $w$  is true in  $x$ . A formula  $BP$  is true in a world  $w$  just if  $P$  is true in every world doxastically accessible from  $w$ . Since doxastic accessibility is not required to be reflexive, not every instance of the schema  $BP \rightarrow P$  is true in every world in every model. This framework accords with a belief-first methodology.

A clue that the accord with a knowledge-first methodology goes deeper is that in by far the most widely applied class of models of epistemic logic accessibility is an equivalence relation (reflexive, symmetric, and transitive), which automatically validates factiveness, and thereby favors interpreting the operator in terms of knowledge rather than belief. Such models are routinely used in most applications of epistemic logic in computer science and economics, for example to model common knowledge. The users are not motivated by any philosophical prejudice in favor of knowledge-first epistemology. They prefer models in which accessibility is an equivalence relation for their simplicity, tractability, and naturalness.

Once we go into finer-grained epistemology, we can no longer require epistemic accessibility to be an equivalence relation, for doing so validates the positive and negative “introspection” schemas  $KP \rightarrow KKP$  (if one knows that P, one knows that one knows that P) and  $\neg KP \rightarrow K\neg KP$  (if one does not know that P, one knows that one does not know that P), which fail for important and interesting reasons. Negative introspection fails because in the skeptical scenario one does not know that one has hands, but one is in no position to know that one does not know that one has hands. Positive introspection fails for reasons closely connected with the non-transitivity of indiscriminability discussed above, since it is equivalent to the transitivity of epistemic accessibility. They are comparatively harmless in most applications in computer science and economics because there the emphasis is on understanding essentially multi-agent epistemic phenomena, such as problems about the transmission of information. For those purposes, failures of introspection within a single agent constitute noise, best eliminated from the model to enable us to discern more clearly the phenomena distinctive of the multi-agent case. In epistemology, by contrast, our purpose is usually to understand phenomena that already occur in the single-agent case, so idealizing away failures of introspection is not harmless. This exemplifies the sort of consideration relevant to selecting an appropriate model. In any case, the fruitfulness of what is in effect a knowledge-first methodology when epistemic logic is applied to other disciplines is at least some indication that knowledge-first epistemology is on the right track.

A more direct point is that the knowledge-first equation of one’s total evidence with one’s total knowledge facilitates a natural way of introducing evidential probabilities into models of epistemic logic. At a coarse-grained level of individuation adequate for present purposes, a proposition in a model is simply a set of worlds, those in which it is true. One proposition entails another just if the former is a subset of the latter. For each world  $w$ , the set  $R(w)$  of worlds accessible from  $w$  is the strongest thing known in  $w$ , in the sense that the propositions known in  $w$  are just those  $R(w)$  entails. Since  $R(w)$  is what one knows in  $w$ , on the equation of evidence with knowledge the worlds consistent with one’s evidence in  $w$  are exactly those in  $R(w)$ . Now suppose that we have a prior probability distribution on the space of worlds. In the simplest case, there is a finite number  $n$  of worlds, each of prior probability  $1/n$ . Then we define the probability of a hypothesis  $h$  on the evidence in a world  $w$  as the conditional probability of  $h$  on  $R(w)$ . That is, we eliminate the worlds inconsistent with the evidence and scale up the probabilities of those that remain to total 1. Since accessibility is reflexive,  $R(w)$  always contains  $w$ , so we never conditionalize on the empty set. Such models can be used to explore the evidential probabilities of states of knowledge, ignorance, and evidential probability in far more systematic and controlled ways than were previously available. For example, it can be argued that sometimes our evidence eliminates a hypothesis  $h$ , even though it is almost certain on our evidence that our evidence does not eliminate  $h$ .

Such models with belief or excusable belief in place of knowledge are implausible. They make evidence non-factive, so a true proposition is sometimes inconsistent with one’s evidence. But that surely misdescribes a case of a true proposition inconsistent with what one mistakenly takes to be one’s evidence.

The obvious place to look for an alternative formal account of evidential probability is the more familiar Bayesian tradition. But it fails to integrate probabilities with an epistemology of evidence. The standard way of updating probabilities is by conditionalizing them on what is called “new evidence,” but rarely with any account

of what constitutes evidence. Although ultra-subjective Bayesians may count whatever the agent updates on as evidence, the claim that neo-Nazis have evidence inconsistent with the occurrence of the Holocaust is disreputable. A more flexible form of updating is Jeffrey conditionalization, on which there need be no proposition that is conditionalized on, but for the same reason it is even harder to integrate with anything that deserves to be called epistemology. Whether a particular change of probabilities is justified is not a purely formal matter; the usual Bayesian tradition provides only formal criteria.

Some rivals to knowledge-first epistemology seem to make evidence non-propositional (for example, if it comprises perceptual appearances). However, it is hard to explain how non-propositional evidence contributes to updating probabilities. We might therefore try to weld such an account of evidence onto the Bayesian framework by equating the new evidence (whatever it is) with the set of worlds in which the agent acquires just that new evidence, and requiring updating by conditionalization on that new evidence proposition. That is in effect to treat exact sameness of new evidence as the relevant accessibility relation. But, as already noted, exact sameness in a given respect is an equivalence relation. Thus such models are isomorphic to the simplest models of epistemic logic above, which accord with a knowledge-first methodology. Even though the accessibility relation is no longer characterized in terms of knowledge, we are back with something suspiciously like a knowledge-first framework. However, for the same reasons as before, we still have the positive and negative introspection principles that if one's evidence entails that P, then one's evidence entails that one's evidence entails that P, and that if one's evidence does not entail that P, then one's evidence entails that one's evidence does not entail that P. Those principles exclude a sort of higher-order uncertainty about the boundaries of one's evidence. But if evidence provides an epistemologically serious way of discriminating propositions, that positive introspection principle will be vulnerable to a version of the argument from the non-transitivity of indiscriminability. Thus the envisaged reduction of apparently non-propositional evidence to propositions yields the analogue of the least sophisticated, most elementary and objection-prone version of the knowledge-first account. A proper knowledge-first epistemology integrates epistemic logic and evidential probability in a way that makes more epistemological sense than any available alternative.

The rigorous development of knowledge-first epistemology began quite recently. We might therefore expect its longer-established rivals to be far ahead, judged by their capacity to generate epistemologically illuminating formal models. Instead, they are already in danger of being left behind. That is one of many signs that they belong to a stagnating research program. It is time to give knowledge-first epistemology its chance to do better.

## Note

- 1 The ideas sketched here are explained and argued for in more detail in my *Identity and Discrimination* (Oxford: Wiley-Blackwell, 1990; revised and updated, 2013), pp. 4–47; *Knowledge and its Limits* (Oxford: Oxford University Press, 2000); “On Being Justified in One’s Head,” in M. Timmons, J. Greco, and A. Mele, eds., *Rationality and the Good* (Oxford: Oxford University Press, 2007); *The Philosophy of Philosophy* (Oxford: Blackwell, 2007), pp. 208–277, and “Very Improbable Knowing,” *Erkenntnis* (forthcoming). Those works cite earlier contributors to the development of knowledge-first epistemology.

# What Is Knowledge-first Epistemology?

*Trent Dougherty and Patrick Rysiew*

What is knowledge-first epistemology? It is not yet clear to us how to answer that question. One key assertion seems to be that knowledge is “unanalyzable” – that is, not neatly factorizable into component parts. As evidence for this claim, Williamson cites the facts that (i) the Gettier problem hasn’t been solved in four decades, and (ii) attempts to solve it have led to clumsy analyses of knowledge. He also mentions there is no reason to think it is factorizable in the first place. He notes that its entailing belief and entailing justification does not entail that those things are constituents of knowledge. We agree, but note that there being such provides a pretty natural explanation for the relevant entailments. It is unclear what Williamson’s alternative explanation of the latter are, but, before we comment briefly on (i) and (ii), his explanation seems to go like this. Knowledge is a kind of success, and mere true beliefs are failures to achieve this success, even when justified. So we can think of (mere) belief, (mere) true belief, and (mere) justified belief as “botched knowledge” (Williamson, 2000, p. 47). But of course Williamson doesn’t think they are really a kind of knowledge. “Former president” doesn’t pick out a special kind of president. It is not as though there is the genus *Presidents*, one species of which is *Former*. The same goes for knowledge and “botched knowledge.” It is not as though there is the genus *knowledge*, one species of which is *botched*. Clearly, this is not at all what Williamson has in mind. But then it is hard to see how thinking of the relation between knowing and certain other states and goods on the model of the relation between doing and trying explains the data in question – namely, that knowledge entails belief and justification. Traditional epistemologists (and they might be knowledge-first in some sense; see Conee, “Truth Connection,” chapter 10 of Conee and Feldman 2004) have a simple explanation for this. For entailment is modeled in formal semantics as set inclusion. *Cat* entails *mammal* because the cats are a subset of the mammals. On the traditional view, *knowledge* entails *belief* because knowledge is a kind of belief, the kind that meets the conditions required to be knowledge. *Knowledge* entails *justified belief* for the same reason: it is a kind of justified belief, the kind that meets further conditions for being knowledge. Truth is one such further condition. Being based on one’s evidence in such a way to avoid a deviant causal chain is another. The latter kind of condition is hard to spell out, since there are so many ways a causal chain can go wrong between takeoff and landing, even if one gets to the right destination in the end.

There is another analogy available, akin to the one Williamson himself favors, which also incorporates teleological thinking into our conception of knowledge. Perhaps belief (for example) is akin to *intention* (rather than to Williamson’s *trying*) and knowing to action. Just as intentions are intentions to perform some action, believings “aim at” knowledge; and just as there are failed intentions, there are failed attempts at knowing (merely true beliefs, say). *This* allows that there is some good sense in which it might be proper to think of belief in terms of its relation to knowledge; so too, it might explain why knowing entails believing. But it also undercuts the motivation for an approach that puts “knowledge first” in some more interesting sense.

For it is very natural to think of actions as individuated in part by the intentions involved – to think, that is, that actions (vs. mere movements, etc.) “include” intentions. Likewise, it is natural to think of knowings as including believings – indeed, to think that what’s known (when it is) is identified in part by the belief(s) involved. But all this merely rehearses the traditional idea that belief enters into our understanding of knowledge, and that knowing is a species of belief.

## **(Un)analyzability**

Williamson thinks considerations pertaining to the Gettier literature show that the justified true belief (JTB) approach to knowledge is moribund at best. And he seems to think that this spells trouble for the value of justification relative to knowledge. But this is far from clear. On the contrary, in the olden JTB days, one might have thought that justification was important only because it was part of an analysis of knowledge. But the unanalyzability of knowledge – if unanalyzable it is – could in fact be seen as a liberation of justification to assume importance in its own right. Kvanvig (1992; 2003, p. 192) and Greco (2010, p. 9ff.), in defending an externalist notion of knowledge, theorize that the intuitions of epistemic justification internalists might be about not knowledge, but understanding, where understanding stands between knowledge and wisdom in value. Suppose knowledge is unanalyzable. Either understanding entails knowledge or not. If it doesn’t (Kvanvig, 2003), then knowledge doesn’t help us understand understanding, whereas justification might. If it does (Grimm 2006), then it appears to be a special subset of knowledge where (at a minimum) certain internalist goods are added (seeing connections, etc.), which might make it a particularly prized kind of knowledge. It then remains an open question whether these goods apart from knowledge are more valuable than knowledge apart from these goods.

Other examples in the same vein are available. There is, for instance, Sosa’s well-known distinction between “animal” and “reflective” knowledge (1991) – though it is perhaps not clear in what sense these are different *kinds* of knowledge, as opposed to a single kind with some “extra goods” sometimes being added, it being an open question what, if anything, accounts for the value of the latter (Kornblith, 2004). (The same goes for Lehrer’s distinction between “discursive” and “primitive” knowledge (2000).) And there is Foley’s argument (2004) that epistemologists have for too long thought that egocentrically rational and reliable belief – roughly, internalistically and externalistically justified belief, respectively – must converge, as opposed to each calling for its own theory and having its own distinctive value.

Much of Williamson’s discussion of “traditional epistemology,” it seems to us, equates it with a particular strain of internalistic theorizing; and much of Tim’s dissatisfaction, we think, is with the presumption that a certain form of internalistic justification *must* be a component of knowledge. This strikes us as a good worry to have. But the examples just mentioned serve to illustrate that abandoning that presumption does not commit one to any specifically “knowledge first” ideas.

So too, all the relevant parties can reject the idea that knowledge admits of any neat analysis. “Maybe,” as Plantinga says, “there isn’t any neat formula, any short and snappy list of conditions (at once informative and precise) that are severally necessary and jointly sufficient for warrant; if so, we won’t make much progress by grimly pursuing

them” (1993, p. 20). Perhaps “the program of analysis,” as Williamson calls it (Williamson, 2000, p. 31), is a hangover from the heyday of logical atomism (Williamson, 2000); perhaps it is rooted in a faulty theory of concepts<sup>1</sup>; perhaps the best we can reasonably hope for is “reflective understanding” (Williamson, 2000, p. 33). Still, it’s an open question whether, within such an understanding, knowledge will have to be taken as unanalyzable, and justification (for example) understood as an entirely derivative notion.

## **Justification and Excuses**

Suppose you unwittingly receive a perfect forgery of an authorization granting you permission to explore a protected piece of land. Note in hand, you proceed past the many “No Trespassing” signs. After about half an hour, you encounter a patrol officer who inspects your alleged permit, detects the forgery, and escorts you off the premises but does not prosecute you. Here are two competing descriptions of the event. First comes the one we take to be the natural interpretation: You were justified in crossing the “No Trespassing” signs *because* you had the misleading signification of the note. Not only are you not to blame, your behavior is not subject to any legitimate criticism. The bare fact of being at odds with a law is irrelevant from a normative point of view, and knowledge is a normative notion. Next comes Williamson’s interpretation. He seems to imply that your behavior was simply unjustified. Your ignorance excuses you from any punishment for this unjustified behavior, but the behavior was unjustified from start to finish. Saying that one was justified in believing they were justified does not change this result.

Suppose the latter interpretation is right. Suppose, that is, that your behavior was unjustified *full stop*, and that you are merely blameless. Taking the epistemic case: suppose, as Williamson holds, that in “the bad case” one is not justified but merely blameless. Doesn’t this go along with thinking that justification, understood apart from knowledge, just doesn’t have much real work to do? Not at all. Kent Bach (1985), for example, has argued that much theorizing about justification conflates issues of justified belief with issues of justified/blameless believers. The result, Bach thinks, isn’t that justification is of no great theoretical interest, but that theorists are freed up to pursue externalist theories thereof, including ones whereby those in the bad case don’t have any justified beliefs. Similarly, Rysiew (2011) considers as a live option the view that those in the bad case lack any real (as opposed to apparent) evidence – “contrary,” as Williamson says, “to what sceptics and many other epistemologists assume” (p. 4). He does so, however, without endorsing “E=K” (and, in fact, while allowing for the sort of view of evidence we outlined in our opening statement). Once again, then, there are various extant theories that preserve the result Williamson ultimately recommends, but without going the knowledge-first route.

## **The Good and Bad Cases**

Williamson insists that what is common to perceiving and misperceiving is not more basic than perceiving. Fortunately, we do not need any basicity claim. All we need is that there *is* something in common. What the veridical and illusory cases have in

common, clearly enough, is what it is like to be in those states. In addition to shared qualitative character, they have the same representational contents. Both experiences assert, as it were, the same thing. And, of course, this phenomenology is *in a way* more direct to us than the external world, since it is in virtue of our experience that we are aware of the world. (Note that this only makes experience the medium, and not the object, of awareness.) Williamson has strong words indeed for the view that narrow content is the “real core” of experience. We are quite glad, then, that “experience first” epistemology is committed to no such thesis. Again, all we need is that there is something *shared* between them. And this shared thing – what it’s like to have the world look that way – is where we need to begin in finding out how the world is. Whether or not the terminology or theoretical machinery of “narrow content” is the best way of getting at it, *what it’s like* to see a red mug is exactly *what it’s like* to see a perfect hologram of a red mug; no cognitive science can reveal *that* to be false. In the good case, *as a matter of fact* what it’s like is causally hooked up to the world in the right kind of way, such that the experiential signs are not misleading and so we have not only justified belief but knowledge. In the bad case, we lack knowledge, but there is some credit in heeding misleading evidence.

Of course, once again, maybe the latter such credit should not be identified with justification of the sort required for knowing. Even so, it should be clear that resisting putting knowledge first does not require – and in our case, does not involve – throwing ourselves behind “the veil of ideas,” seeing subjects as being acquainted only with “appearances,” and suchlike. We take our view to be perfectly compatible with direct realism about the external world and not to be a version of sense-datum theory (Chisholm adopts the adverbial view to defend direct realism, but is an “experience first” epistemologist; Chisholm, 1989, esp. pp. 66ff.). See also Huemer, 2001, chapter IV, esp. sect. 5). So too, it should be clear that ours is not a view that is born out of an obsession with the problem of skepticism. That experience comes first in the indicated sense is simply a fact of our everyday epistemic lives.

## Regarding Indiscriminability

Williamson offers a sorites argument against the principle that “justification is exactly the same in cases indiscriminable to the subject.” Indiscriminability, he hypothesizes, is the link between appearances and justification. (We are a bit worried about the shift in that discussion from talk of *beliefs* being justified to *persons* being justified, for it suggests a possible running together of the idea of epistemic justification and epistemic responsibility.) What he may have in mind is this: It seems that our unfortunate envatted counterparts are just as justified as we are in believing in an external world, for that is *all we have to go on*. We have already indicated how our own central claims do not require that result. But suppose it is right. What is the problem? Williamson claims that this picture falls prey to a sorites argument. One problem with his argument is that it assumes experiences are fully determinate rather than vague. Yet fairly early on in the development of contemporary empiricism – in response to the Problem of the Speckled Hen – both Ayer (1940) and Chisholm (1942)<sup>2</sup> noted that experiences or the characters of sensing are in fact not fully determinate (see Tye, 2009, for an excellent discussion of a representationalist perspective on this problem and Dougherty, 2011,

for more context). We do not have  $RED_{\lambda=.72744}$  experiences and  $RED_{\lambda=.72745}$  experiences. Rather, we have “reddish” or even “somewhat reddish” experiences. All basic evidence is necessarily vague. But even if higher-order anti-luminosity prevents this, nothing important follows, for, as Jeffrey showed, we can get along perfectly well with uncertain evidence.

## Regarding Formal Methods

We agree that formal models can sometimes be illuminating of epistemic phenomena. And it is hard to disagree that “shaky reasoning” isn’t the way to go. However, we suggest that traditional epistemologists sometimes exercise the same kind of careful, skilled inquiry as do formal epistemologists, even when they are not working with numbers or special symbols. Williamson mentions several of the many risks of using formal models, and it seems to us that the cost-to-benefit ratio is about the same in formal and non-formal modes of inquiry, when pursued conscientiously by capable parties.

It is unclear how any kind of Bayesianism is an *alternative* to (almost) any epistemic logic, for it is unclear whether they are designed to do the same thing. For example, if we let the weak epistemic modality be “it is permissible for the agent to believe that,” there are forms of Bayesianism perfectly compatible with the theorems such an epistemic logic would include. Likely this is so also if the strong operator reads “It is certain that.” Also, we don’t know in advance whether the simplifications and idealizations of the respective theories would “hook up” in a direct enough way to constitute a rivalry. And if there were rivalry, we would perceive no threat to experience-first epistemology, because it is not at all clear that the success of epistemic logic in being illuminating furnishes any reason to adopt the knowledge-first approach.

Williamson says Bayesianism fails to integrate probabilities with an epistemology of evidence. Like all philosophers, Bayesians start out with a set of problems to solve. Most Bayesians are concerned with ways of characterizing coherence properties, which are generally agreed to be good-making features of one’s noetic structure. We doubt Williamson disagrees. Other Bayesians are interested in formal learning theory and are not concerned with where the evidence comes from, but rather with what one ought to do with it when one gets it. So a theory of evidence is simply outside the scope of standard Bayesian pursuits. But of course many Bayesians have a favored view of evidence (Swinburne, 2001, for example). Many probabilists don’t commit to any formal learning theory and simply see probability logic as exactly parallel to first-order logic. Its job is to tell you what choices you face given your current commitments (Howson and Urbach, 1996). In Jeffrey’s case, the assigning of probabilities to basic propositions is a *techne* – it requires an art of judgment that is generally acquired by practicing in the appropriate community (Jeffrey, 1992). Williamson says, “When we ask how probable a theory is on our evidence, we want something less dependent on our doxastic state than a credence but more dependent on our epistemic state than a chance” (p. 5). This is just what we get on an experiential theory of evidence wedded to a Chisholm-like theory of evidence that states objective, material epistemic principles. We end up with a form of epistemic probability that strikes the appropriate balance between subjectivity and objectivity.

Williamson says that traditional epistemology, since it has been going on much longer than knowledge-first epistemology, should be ahead when “judged by their capacity to generate epistemologically illuminating formal models” (p. 9). It is not at all clear why it should have been expected to generate *any* such models. Nor is it clear that what he calls “traditional” epistemology *has* been around longer. Indeed, it is sometimes criticized as having started essentially with Descartes, whereas some kind of “knowledge first” view seems to go back to Aristotle, and to have been expounded by Scholastic philosophers.<sup>3</sup> Mathematical philosophy germinated in the nineteenth century, budded in the 1920s, and began to blossom in the mid to late 1970s. So we doubt we should expect any kind of theory to bear heavy formal fruit yet. (And it is worth noting that it has been only a little more than four decades since Gettier’s paper, and that some people are satisfied that they have solved the problem, and in a very non-grue-like fashion (Feldman, 2003, p. 125), and interesting work continues to be done on the topic (Bernecker, forthcoming). Furthermore, it has now been a fourth of that time since publication of *Knowledge and Its Limits*, and there remain many, many details to work out. It would be unreasonable to expect of any view, however, that a project worth doing could be done in the lifetime of one philosopher.) However, Bayesianisms and more general probabilisms are alive and well, indeed flourishing across the globe today, and most of them are compatible with most theories of evidence, including the commonsense empiricist theory that our fundamental evidence – at the base or frame of our noetic structure – consists in the experiences by which the world is revealed to us but which sometimes lead us astray.

Finally, Williamson notes that some rivals to knowledge-first epistemology take evidence to be non-propositional. For example, my warm feeling can be evidence that it is hot in here. But then he says that it is hard to explain how non-propositional evidence contributes to updating probabilities. This issue has been discussed quite a bit elsewhere,<sup>4</sup> including in our opening remarks, so there is no need to add to them here.

## Notes

- 1 Cf. Kornblith (2007), who notes trouble for the traditional philosophical project of conceptual analysis: “Knowledge, for example, may be analyzed, on certain views, as justified, true belief meeting some additional, and difficult to specify, condition. It is taken for granted that the form of a proper analysis is just some such set of individually necessary and jointly sufficient conditions. The idea that our concepts are mentally represented in this form is what psychologists refer to as the Classical View of concepts. Since the early to mid-1970s, it has become increasingly clear that the Classical View is not correct.” The moral Kornblith draws from this, however, isn’t Tim’s. Rather, it is that the standard philosophical method, which includes liberal appeal to intuitions, is not a reliable method of understanding our concepts – and, he thinks, our concepts are not plausibly viewed as the target of philosophical understanding anyway.
- 2 In fact, Chisholm explicitly considers the non-transitivity of indistinguishability (1942, p. 371). He notes that Russell was dealing with this problem as early as 1921. He claims that this allows basic evidence to be certain, though we are not committed to that.
- 3 “Whereas Descartes seeks to place philosophy and science on firm foundations by subjecting all knowledge claims to a searing methodological doubt, Aristotle begins with the conviction

that our perceptual and cognitive faculties are basically dependable, that they for the most part put us into direct contact with the features and divisions of our world, and that we need not dally with sceptical postures before engaging in substantive philosophy” (<http://plato.stanford.edu/entries/aristotle/>). Yet Aristotle held that knowledge of the external world was by means of “sensible species.” Reid has an interesting discussion of Aristotle’s epistemology (*IP* IV 2, *W* 372a-b).

4 See note 9 from our opening statement “Experience First.”

## References

- Ayer, A.J. (1940) *The Foundations of Empirical Knowledge*. London: Macmillan.
- Bach, K. (1985) A rationale for reliabilism. *The Monist* 68 (2): 246–263.
- Bernecker, S. (forthcoming) Keeping track of the Gettier problem. *Pacific Philosophical Quarterly*.
- Chisholm, R. (1942) The problem of the speckled hen. *Mind* 51 (204): 368–373.
- Chisholm, R. (1989) *Theory of Knowledge*, 3rd edn. Upper Saddle River, NJ: Prentice Hall.
- Conee, E. and Feldman, R. 2004. *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.
- Dougherty, T. (2011) Evidence. In *Oxford Bibliographies Online: Philosophy*. <http://oxfordbibliographiesonline.com/view/document/obo-9780195396577/obo-9780195396577-0132.xml> (last accessed February 27, 2013).
- Feldman, R. (2003) *Epistemology*. Englewood Cliffs, NJ: Prentice Hall.
- Foley, R. (2004) A trial separation between the theory of knowledge and the theory of justified belief. In *Ernest Sosa and His Critics*, ed. J. Greco. Oxford: Blackwell.
- Greco, J. 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge University Press.
- Grimm, S. (2006) Is understanding a species of knowledge? *British Journal for the Philosophy of Science* 57 (3).
- Howson, C. and Urbach, P. (1996) *Scientific Reasoning: The Bayesian Approach*. Open Court.
- Huemer, M. (2001) *Skepticism and the Veil of Perception*. Lanham, MD: Rowman & Littlefield.
- Jeffrey, R. (1992) *Probability and the Art of Judgment*. Cambridge: Cambridge University Press.
- Kornblith, H. (2004) Sosa on human and animal knowledge. In *Ernest Sosa and His Critics*, ed. J. Greco, 2004. Oxford: Blackwell.
- Kornblith, H. (2007) Naturalism and intuitions. *Grazer Philosophische Studien* 72: 27–49.
- Kvanvig, J. (1992) *The Intellectual Virtues and the Life of the Mind: On the Place of the Virtues in Contemporary Epistemology*. Lanham, MD: Rowman & Littlefield.
- Kvanvig, J. (2003) *The Value of Knowledge and the Pursuit of Understanding*. New York: Cambridge University Press.
- Lehrer, K. (2000) Discursive knowledge. *Philosophy and Phenomenological Research* 60 (3): 637–653.
- Plantinga, A. (1993) *Warrant and Proper Function*. Oxford: Oxford University Press.
- Reid, T. (1895/1983) *Thomas Reid: Philosophical Works*, 8th edn. Edited by William Hamilton. New York: Georg Olms Verlag.
- Rysiew, P. (2011) Making it evident: evidence and evidentness, justification and belief. In *Evidentialism and Its Discontents*, ed. T. Dougherty. Oxford: Oxford University Press.
- Sosa, E. (1991) *Knowledge in Perspective: Selected Essays in Epistemology*. New York: Cambridge University Press.
- Swinburne, R. (2001) *Epistemic Justification*. Oxford: Oxford University Press.
- Tye, M. (2009) A new look at the speckled hen. *Analysis* 69 (2): 258–263.
- Williamson, T. (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.

# Experience First

*Trent Dougherty and Patrick Rysiew*

In what way is experience first, epistemically speaking? It is first, at least, in the order of immediacy: in short, experience is where we begin (which is not itself to say that it is what we know best). This has important implications that we will discuss shortly, but first we are happy to acknowledge that knowledge might indeed be first in some ways as well. For example, we appear to acquire the concept very early, before that of belief (much less justification) (Perner, 1993; Bartsch and Wellman, 1995; Williamson, 2000, p. 33, n.7). From this, though, nothing follows about the correct relative priority of the concepts in epistemological theory. (Still less does it tell us about the metaphysics of knowledge itself.) One might also think knowledge first in one or another teleological way. For instance, it may be that belief “aims at” knowledge, or that knowledge is “the norm of belief.” But, if true, this doesn’t entail any more substantive sense in which belief must be understood in terms of knowledge, not the other way around. Alternatively, one might think the main purpose of our cognitive system as a whole – or, at least, the portion(s) thereof in which epistemologists are traditionally interested – might be to acquire knowledge. This thesis faces opposition from either end. From a broadly naturalistic perspective, it might seem that mere true belief, or indeed just “getting by,” is the purpose of our cognitive architecture.<sup>1</sup> At the same time, there are familiar but understudied epistemic goods that, according to some (Kvanvig, 1992, 2003; Zagzebski, 1996; Greco, 2010), are clearly more valuable than knowledge: understanding and wisdom. Indeed, one of the big shifts in epistemology toward the end of the twentieth century was a return to interest in epistemic virtues. So knowledge may turn out to be just a middling epistemic desideratum. More radically, some Bayesians find no need for the notion at all.

The view just rejected as insufficient was a particular teleological interpretation of the “knowledge first” slogan, one which makes knowledge the first on the list of epistemological ends. Most of our lives, however, are occupied with securing means. An end having been set, the question is, What do I do now? This is where experience is first: in the quest for true belief, justification, knowledge, understanding, and wisdom we have no other starting point than experience. Our experiences (broadly construed to include what it’s like to have intuitions and rational insights, etc.) are our basic evidence, in the light of which all else that is evident is made evident. Experiences play well the roles that characterize evidence. We will consider four such roles (see Kelly, 2011, for more on the functional approach) and show how experiences are well suited to the task.

Let us be clear from the outset: it might be that no single thing can play all the roles typically ascribed to evidence (Kelly, 2011). Nevertheless, a good argument can be made that experiences can satisfy all adequately.<sup>2</sup> Furthermore, we do not here argue that experiences are the *only* thing that can play these evidential roles. The thesis, rather, is that experiences *can and do* play these roles, and that they do so, moreover, in a way that can reasonably be described as “basic” or “ultimate,” in the sense that they are evidential regress stoppers. Any chain of cited evidence must end with the

way the world appears to us to be. So on this view, experience is first in that it inhabits the ground floor of the intellectual edifice. The four roles we will consider as evidence are: (1) what justifies belief; (2) what rational thinkers judge by; (3) a guide to truth (evidence as sign, symptom, or mark); and (4) neutral arbiter (objectivity, publicity, intersubjectivity).

1. *Experience is what ultimately justifies belief*, for it is ultimately to your experience that your beliefs must be called to account. Its very ultimacy sometimes conceals this fact, for we rarely need to dig that deep. Our “derived evidence,” evidence that is based on immediate experience,<sup>3</sup> is rarely called into question in ordinary life. So, for example, you say you saw a bear and are challenged to provide evidence. You reply “Well, I saw a very large mammal with thick black fur foraging for berries.” That will usually suffice, for it is rare that we doubt that someone is able to accurately describe their experiences. But however socially awkward it might be, asking for further evidence for that claim is perfectly coherent and, in special cases, appropriate. When pushed to this deeper level, one has nothing else to appeal to besides one’s experiences. And note this as well. Once one *has* appealed to one’s experiences – construed here as the way things appear to them to be – there is no question of further evidence. Once you have said, “I had an experience with these properties (blackness, certain geometric patterns)” – or, if you prefer, “there seemed to be something large-ish and black-ish (etc.)” – calls for further evidence defending the claim that you had such an experience are wrong-headed (this feature makes them attractive to foundationalists as regress-stoppers). We have reached epistemic rock bottom. And of course, sometimes citing one’s experiences as evidence is the most natural thing in the world. I assert that the temperature has fallen. You ask me why I think this. I say that I feel cold. I do not have in mind the *fact that* I feel cold, but, rather, *my feeling* cold.
2. *Rational thinkers judge by their experience*. Given that our ultimate evidence consists in experiences, it is platitudinous that rational thinkers judge by their experiences. To continue the example introduced just above: given that you feel cold but didn’t before, it will, other things being equal, be rational for you to judge that the temperature has dropped. Or, consider another example: in your large backyard, you see a bird and can’t tell whether it is a female cardinal or a juvenile male. The female will be slightly more gray with slightly more orange beak and a slightly rounder body. You strain your eyes to get a better view and attend more closely to the bird and(!) to the experiences you are having (and note that the experience of the selfsame bird will change as you squint or put on your glasses, etc.). To come to a judgment concerning whether it is a female cardinal or a juvenile male, you will, if you are rational, judge by your experience. Note that the claim that one is attending to and being guided by features of one’s experience does *not* imply that experiences are the primary objects of knowledge (see Crane, 2011, for more on this) or that one is not *at the same time*<sup>4</sup> aware of the objects in the world – assuming that one is not hallucinating – which give rise to those experiences. Also note that being aware of a feature is not obviously itself a form of knowledge.

3. *Experience is a guide to truth: Evidence as sign, symptom, or mark.* The idea that, say, smoke is a sign of fire, is common enough. So, seeing smoke gives one evidence that there is fire. It is tempting to say that this is because smoke is a reliable indicator of fire. But what if it turned out that your experience was very atypical and that, in fact, most of the time smoke was not correlated with fire? You might say, "Well, it was a reliable indicator *in my experience*." Thus it seems that it is not *mere* objective correlation which makes something evidence. With this we are close to the notion of experience as evidence, but not all the way there. There is a notion that we might call "scientific evidence," where we are adopting an idealized third-person perspective; here, we say that something is evidence when we are already aware that Fs are positively correlated with Gs (see §1.1 of Conee and Feldman, 2008). But the notion epistemologists are interested in is such that all evidence is *someone's* evidence or evidence *for someone*. If your background experiences are different from mine, the same observation can be evidence that *p* for you and that  $\sim p$  for me.

In its basic sense, evidence is what *makes evident* some proposition. It is that in the light of which a proposition seems true. But this light is cast, fundamentally, by experience. So experience is our ultimate evidence. For example, we have all had, when considering some theorem of logic, that "aha" experience, the moment at and in virtue of which the theorem is made evident. It is most in evidence in self-evident propositions. It could be that self-evidence is factive, but it seems that as in all other areas, evidence can be misleading. The obviousness of, say, the naive axiom of comprehension seemed just as clear as the obviousness of some true axioms.<sup>5</sup> Experience reveals the world to us in perception. The features of a certain experience *make evident* to me that there is an elm before me, for example. Chisholm (who combined experience-first epistemology with direct realism) puts it this way: "In the case of being appeared to, there is something, one's being appeared to in a certain way, that one interprets as being a *sign* of some external fact."<sup>6</sup> Experiences are not the *objects* of knowledge, but they are the *medium* for knowledge.<sup>7</sup> For every state of knowledge, there is some experience that makes the fact in question evident. This is a central way in which experience is prior to knowledge. One's *feeling* cold is different from one's *being* cold, in the sense that one's core body temperature has dropped. But one's experience of coldness typically makes evident that one's core body temperature has dropped. In this way experience is the sign of what the world is like.

4. *Experience is a neutral arbiter among disputants.* How could private experiences play such a seemingly public role? Aren't they too "subjective"? It's not hard, actually. The way in which experience plays an intersubjective role is familiar to all. You are sailing with friends off the Gulf Coast, and in your periphery you think you see a dolphin jump. You turn to your friends and say "Did you see that?! I think I just saw a dolphin jump!" The friend standing closest says, "I saw something in that direction too, but it seemed too small to be a dolphin." Another says, "I thought I saw a white cap on the top of that shape." Another, "Me too." Your experience is overruled by their experiences. When the collective experiences are taken together, the evidence suggests that, though dolphin sightings are not infrequent here, what you saw was a stray wave rather than a dolphin. This is intersubjective in a way sufficient to satisfy all the constraints of scientific, legal, and medical inquiry.<sup>8</sup>

## Conclusion

As we mentioned above, it could well be that there is such a thing as derived evidence and that beliefs or propositions can play the roles too. Basic knowledge might be near the evidential foundations, but we reject both that all knowledge is evidence and that only knowledge is evidence. In fact, maybe we should all, like Thomas Reid, be *pluralists* about evidence. Reid says that “[w]e give the name of evidence to whatever is the ground of belief” (*IP* II 20, *W* 328a); and, he thinks, there are different types or sources of evidence: there is the evidence of sense, of memory, of consciousness, of axioms, of reasoning, and so on (*IP* IV 20, *W* 328a). Various kinds of experiences (perceptual, memorial, introspective), arguments, testimony, the judgment of recognized authorities, the marks or signs by which we distinguish between kinds of things, a person’s past actions, various “signs” of another’s mind and/or character (gestures, facial expression, etc.), observed connections in the world – these are all things which Reid seems to count as evidence.<sup>9</sup> Perhaps all these usages can be reduced to a single experiential notion, such as in phenomenal conservatism (see Huemer, 2001), or perhaps not. At any rate, there is no reason to think all and only knowledge is evidence, and there is ample reason to consider experiences a legitimate, and even the most basic form of evidence.

Of course, we have left many questions open – for instance, just *why* experience confers justification. That it does can be combined with any number of epistemological theories, internalist and externalist. Even externalist Alvin Plantinga recognizes that:

My perceptual beliefs are not ordinarily formed on the basis of *propositions about* my experience; nonetheless they are formed on the basis of my experience. You look out of the window: you are appeared to in a certain characteristic way; you find yourself with the belief that what you see is an expanse of green grass. You have *evidence* for this belief: the evidence of your senses. Your evidence is just this way of being appeared to; and you form the belief in question *on the basis of* this phenomenal imager, in this way of being appeared to. (1993, p. 98)<sup>10</sup>

The limitation of evidence to propositions (of any kind),<sup>11</sup> however, seems to us to reflect an over-intellectualization of inference and epistemic support.<sup>12</sup> To some extent, the focus on propositions is perfectly natural. As noted above, our everyday epistemic practices seldom oblige us to descend to the level of what we’ve talked about here as “ultimate evidence” (non-propositional experiences). And, when we do, the fact that we’re thinking and talking about such matters itself renders them into propositional form. Still, commitment to principles like “what gives probability must also receive it” (Williamson, 2000, p. 196) seems to lay evidence upon a Procrustean bed (not to mention the circularity or regress worries they might raise). We will say a bit more about this in our rejoinder.

## Notes

- 1 For more deflationary views, see for example Churchland (1987) and Stich (1990).
- 2 For a treatment of evidence that unifies the roles without assuming that any one thing can fulfill them all, see Rysiew (2011).
- 3 There is some question whether this is in fact a kind of evidence. See Littlejohn (2011).

- 4 Aquinas did not think this possible. See Adler (1985, pp. 14ff.).
- 5 Saying such things is, of course, fully compatible with a broadly reliabilist, or otherwise externalist (e.g., proper functionalist, etiological functionalist, etc.), approach to evidence and/or other epistemic goods. See Rysiew (2011).
- 6 Chisholm (1989, p. 67).
- 7 In the language of the medieval Aristotelian epistemologists, ideas are the *qua* of knowledge, not the *quod*. See, for example, Aquinas's "Treatise on the Powers of Man" in the *Summa Theologica*.
- 8 Of course, when we talk about evidence the way we do in a court of law – where physical objects in the room are referred to as "evidence" – we are, as always, speaking in a way that reflects what is mutually obvious: here, the fact that no one has any doubt that everyone is having the same kind of experience. That we in this way take for granted ultimate evidence speaks to its very ultimacy.
- 9 The discussion is not always explicit and systematic. For example, in comparing "the evidence of sense" with that of reasoning and consciousness, Reid doesn't come right out and say just what the evidence of sense is. Some of the relevant passages include: *IP* II 20, *W* 328aff.; VII 3, *W* 481bff.; VI 5, *W* 441aff.
- 10 Based on the way Plantinga has defined "direct awareness" on p. 53, it follows that we are not directly aware of the grass. But on p. 189, he says that "in another and perfectly good sense of "directly aware of," I *am* directly aware of [the object]."
- 11 Tim defends this move in §9.5 of Williamson (2000). Conee and Feldman (2008) respond, then Dougherty (2011) pressed Williamson's case further, but reconverted after reading Conee and Feldman's (2011) reply.
- 12 It would be an instance of the "argumentational" view of evidence discussed by Rysiew (2011).

## References

- Adler, M.J. (1985) *Ten Philosophical Mistakes*. New York: Touchstone.
- Bartsch, K. and Wellman, H.M. (1995) *Children Talk about the Mind*. New York & Oxford: Oxford University Press.
- Chisholm, R. (1989) *Theory of Knowledge*, 3rd edn. Englewood Cliffs, NJ: Prentice Hall.
- Churchland, P. (1987) Epistemology in the age of neuroscience. *Journal of Philosophy* 84: 544–553.
- Conee, E. and Feldman, R. (2008) Evidence. In *Epistemology: New Essays*, ed. Q. Smith. Oxford: Oxford University Press.
- Conee, E. and Feldman, R. (2011) Reply to Dougherty. In *Evidentialism and Its Discontents*, ed. T. Dougherty. Oxford: Oxford University Press.
- Crane, T. (2011) The problem of perception. In *The Stanford Encyclopedia of Philosophy*, ed. E.N. Zalta. <http://plato.stanford.edu/entries/perception-problem/> (last accessed February 27, 2013).
- Dougherty, T. (2011) In defense of a propositional theory of evidence. In *Evidentialism and Its Discontents*, ed. T. Dougherty. Oxford: Oxford University Press.
- Dougherty, T., ed. (2011) *Evidentialism and Its Discontents*. Oxford: Oxford University Press.
- Greco, J. (2010) *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.
- Huemer, M. (2001) *Skepticism and the Veil of Perception*. Lanham, MD: Rowman & Littlefield.
- Kelly, T. (2011) Evidence. In *The Stanford Encyclopedia of Philosophy*, ed. E.N. Zalta. <http://plato.stanford.edu/entries/evidence/> (last accessed February 27, 2013).
- Kvanvig, J. (1992) *The Intellectual Virtues and the Life of the Mind: On the Place of the Virtues in Contemporary Epistemology*. Lanham, MD: Rowman & Littlefield.
- Kvanvig, J. (2003) *The Value of Knowledge and the Pursuit of Understanding*. New York: Cambridge University Press.

- Littlejohn, C. (2011) Evidence and knowledge. *Erkenntnis* 74: 241–262.
- Perner, J. (1993) *Understanding the Representational Mind*. Cambridge, MA: MIT Press.
- Planginta, A. (1993) *Warrant and Proper Function*. Oxford: Oxford University Press
- Reid, T. (1895/1983) *Thomas Reid: Philosophical Works*, 8th edn. Edited by W. Hamilton. New York: Georg Olms Verlag.
- Rysiew, P. (2011) Making it evident: evidence and evidentness, justification and belief. In *Evidentialism and Its Discontents*, ed. T. Dougherty. Oxford: Oxford University Press.
- Stich, S. (1990) *The Fragmentation of Reason*. Cambridge, MA: MIT Press.
- Williamson, T. (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.
- Zagzebski, L. (1996) *Virtues of the Mind*. Cambridge: Cambridge University Press.

## Knowledge Still First

*Timothy Williamson*

Trent Dougherty and Patrick Rysiew (D&R) advocate “experience-first epistemology.” They regard experience as our only starting point for inquiry. That may sound like common sense: surely we must start with what we have learnt over the course of our lives. But D&R don’t mean the slogan that way. Climbing Toubkal is one of my valued experiences, but it isn’t an experience at all in their sense. For D&R, the “experience” was *what it was like to climb Toubkal* or *how climbing Toubkal appeared to me*, ways of experiencing I share with someone merely in an appropriate reality simulator, unlike climbing Toubkal itself. Thus their appeals to everyday comments such as “It was a reliable indicator in my experience” are quite misleading. D&R mean that *appearances* are our only starting point for inquiry, which is not common sense but theory much in need of argument. One might have hoped to start with obvious facts about the external world, but they are not appearances: the fact that there are other people is not the fact that there appear to be other people.

D&R argue that all chains of justification for belief ultimately lead back to appearances: if you cite something other than appearances, you can be challenged for its justification, but once you cite appearances, any such further challenge is “wrong-headed.” Well, if the challengers are D&R, they can challenge whatever is not an appearance and stop when they are given appearances, but that shows something about them, not about the epistemic role of appearances. Several of their illustrative appeals to experience, such as “There seemed to be something large-ish and black-ish (etc.),” raise an immediate difficulty. Appeals to *past* experience can easily be challenged by other evidence: the police may discover that immediately after the sighting I wrote in my diary, “There did not seem to be anything large-ish or blackish,” and that my memory is notoriously unreliable. Perhaps D&R’s repeated use of the past tense is a (revealingly natural) slip; on reflection they may prefer to restrict one’s ultimate justifiers to one’s *present* experiences. The need for such a restriction dramatizes their distance from the commonsense conception of judging on the basis of one’s experience.

D&R’s inchoate account of intersubjective evidence is revealing too. They treat it as a matter of pooling appearances to different subjects, not mentioning that on their view each subject starts only with the appearances of the others’ reports of their

appearances, on a par with other appearances of the external world, so that nobody is in a position to do the pooling. Their glib comment that “It’s not hard, actually” suggests inexperience of the complex problems facing any attempt to develop a serious theory of shared scientific evidence within their framework (problems related to the failure of Carnap’s program in the *Aufbau*).

Even when we restrict discussion to present appearances to the subject, D&R’s theory is quite implausible. For example, suppose I see two complicated figures and report that the figure on the left appears to me in exactly the same way shapewise as the figure on the right. When someone points out a difference in shape between the two figures, I notice it immediately. A plausible hypothesis is that all along the figure on the left appeared to me in a slightly different way shapewise from the figure on the right, but I failed to notice the difference; even if the two figures appeared to me to be exactly the same in shape, it does not follow that the way they appeared shapewise was exactly the same. Since such errors about present appearances to one are possible, I may be appropriately challenged to justify my claim that the two figures appear to me in exactly the same way shapewise, and in other circumstances may aptly reply that my vision is normal and the two figures were cast from the same mold. The idea that justification ends with appearances is neither a datum nor a discovery, but wishful thinking in the service of an outdated epistemological theory.

Somehow attributing ultimate justification to the appearances themselves, rather than facts about them, does not help. To vary the example, suppose that although two figures really do appear to me in exactly the same way, my past record provides massive evidence that I have failed to notice a slight difference between two ways of appearing. Am I justified in claiming that the ways of appearing are identical? D&R’s account is too undeveloped to engage with such questions.

D&R later invoke Jeffrey conditionalization to handle uncertain evidence, but the appeal is notoriously empty without constraints on which Jeffrey conditionalizations are epistemically permissible in given circumstances, which they do not attempt to provide. Instead, as D&R interpret Jeffrey, “the assigning of probabilities to basic propositions is a *techne* – it requires an art of judgment that is generally acquired by practicing in the appropriate community.” Maybe so, but that does not answer the question, for it says nothing about the content of the norm good judges are better at satisfying than bad ones. Changing horses, D&R claim that “on an experiential theory of evidence wedded to a Chisholm-like theory of evidence” we “end up with a form of epistemic probability that strikes the appropriate balance between subjectivity and objectivity,” with no hint of the hard formal and interpretative work needed to turn such programmatic remarks into a serious option.

D&R write: “it seems to us that the cost-to-benefit ratio is about the same in formal and non-formal modes of inquiry, when pursued conscientiously by capable parties.” But we do not have to choose between exclusively formal and exclusively non-formal modes of inquiry. A better mode of inquiry than either combines formal and informal considerations, using informal reasoning to select appropriate formal models and formal models to discipline informal reasoning. The key is to integrate the two sides. Although various forms of Bayesianism are formally consistent with experience-first epistemology, D&R do not come to grips with the challenge of integrating them into a combined approach.

In “Knowledge First” I discussed the underlying motivation for experience-first epistemology. One source is the principle that justification is exactly the same in cases

indiscriminable to the subject. I used the non-transitivity of indiscriminability to refute that principle. D&R complain that my argument falsely “assumes experiences are fully determinate rather than vague.” Their complaint is mistaken. My argument makes no such assumption, as the reader can check. It works perfectly well if experiences are vague. Thus the refutation of the principle stands.

Although much more could be said, the upshot remains that D&R have failed to rehabilitate experience-first epistemology. I turn to their objections to knowledge-first epistemology. For all their puzzlement about its nature, it is an easily recognizable way of doing epistemology. Naturally there are borderline cases, just as there are for the rival ways of doing epistemology, such as D&R’s. For reasons of space, I focus on just four issues.

1. D&R wonder how knowledge-first epistemology can explain the presumed entailment from knowledge to belief. They make heavy weather of my suggestion that knowing stands to believing roughly as doing stands to trying. The picture is that when one (intentionally) does something, one tries to do it and succeeds, although for Gricean reasons it is conversationally misleading for someone to say just that one tried to do it. Thus doing entails trying. The converse fails; unsuccessful trying is mere trying, botched doing, trying without doing. Similarly, when one knows something, one believes it; one as it were aspires to know it and succeeds, although for Gricean reasons it is conversationally misleading for someone to say just that one believes it. Thus knowing entails believing. The converse fails; believing that falls short of its aspiration to knowledge is mere believing, botched knowing, believing without knowing. Of course, such analogies never go all the way, otherwise they would be identities, but D&R’s comments raise no difficulty for the picture just sketched.

Equally puzzling is D&R’s alternative explanation on behalf of traditional epistemologists. The key is supposed to be that “entailment is modeled in formal semantics as set inclusion.” But epistemologists of any persuasion can use that commonplace to articulate the entailment from knowing to believing: the set of cases in which one knows that P is a subset of the set of cases in which one believes that P. D&R have identified no substantive issue here.

2. D&R suggest that *understanding* may trump knowledge as an epistemic value. The presupposition that understanding does not involve knowledge is widespread (for example, in discussions of what philosophy aims at), but hardly withstands scrutiny. If you do not know how a jet engine works, you do not understand how it works. If you do not know why the plane crashed, you do not understand why it crashed. Fortunately, D&R allow understanding to involve knowledge. In that case, they suggest, understanding is “a special subset of knowledge where (at a minimum) certain internalist goods are added (seeing connections, etc.),” and these goods may have a value independent of knowledge. They supply no warrant for the gloss “internalist” here. Understanding why the plane crashed may involve seeing the connection between the pilot’s tiredness and his poor decision making. What’s internalist about that? What matters is whether the investigators come to know that the pilot’s tiredness was connected in the relevant way with his poor decision making, and then use that knowledge to acquire other relevant knowledge of the crash. There is no comfort here for attempts to displace knowledge as the central epistemic value.

An essential prerequisite for useful discussion of the relation between knowledge and understanding is systematic explicitness about *what* is to be known or understood. The better that discipline is maintained, the less understanding looks like a rival value to knowledge. Rather, some things are more worth knowing than others. Knowing why or how something happened is often much more valuable than just knowing that it happened. We should not imagine that knowing why or how something happened is non-propositional knowledge, just because it goes beyond knowing that it happened. Knowing why it happened is knowing that it happened for a certain reason; knowing how it happened is knowing that it happened in a certain way. In knowing the answers to some questions, we know far more than we do in knowing the answers to other questions. A good challenge for any epistemologist is to develop a non-trivial theory about that dimension of comparison. But such an investigation need have no tendency to downgrade the central value of knowledge.

3. “The rigorous development of knowledge-first epistemology began quite recently,” I noted, compared to the corresponding development of experience-first epistemology; I gave as one sign that the latter is a stagnating research program its relative lack of progress in generating epistemologically illuminating formal models. In response, D&R make various historical observations. As they indicate, knowledge-first epistemology itself is not new; it has deep roots in the Aristotelian tradition (the introduction to *Knowledge and its Limits* starts with an epigraph from Aristotle’s *Metaphysics*). My point concerned its recent integration with formal epistemology, in the shape of both epistemic logic and objective Bayesianism. Although formal epistemology in an experience-first spirit goes back much longer, attempts to integrate the formal and informal sides remain epistemologically naive or empty, in particular with respect to the general theory of evidence. As already explained, D&R do not come to terms with this challenge.
4. D&R repeatedly emphasize the variety of rivals to knowledge-first epistemology, including approaches that accept some knowledge-first theses while rejecting others. As in natural science, you don’t get far in criticizing a research program just by listing its rivals, without showing that one of them is doing better, or at least equally well. In the long run, we can adequately judge a serious program only by letting its proponents develop its potential. Experience-first epistemology has had a long run, with poor results. Time for alternatives, of which knowledge-first epistemology has made most recent progress.

## Still Nowhere Else to Start

*Trent Dougherty and Patrick Rysiew*

Unlike climbing Toubkal, something’s looking tall or seeming upon reflection to be true aren’t typically things to write home about. Often, such experiences aren’t noticed (as such) at all. For well-functioning humans in normal environments, however, they

too are ways of being engaged with the world. That some facts about that world are obvious (evident), and that we sometimes start there, is perfectly compatible with our view, as is the fact that such beliefs as we have about our experiences can be mistaken.

Williamson mentions that (in our framework) pooling evidence requires the use of testimony as if that's a problem. But "a serious theory of shared scientific evidence" will surely be an endeavor in social epistemology involving trust and an acknowledgment of the epistemic significance of (seeming) comprehension of others' presentations-as-true.

Williamson tends to associate experience-first epistemology with some pretty radical forms of internalism (e.g., Descartes's, Carnap's). But, as it was for Reid (hardly a Cartesian!), it is for us the "*belief-evoking experiences* characteristic of [our] faculties" (Wolterstorff, 2004, p. 80, emphasis added), and of evidence generally, that is epistemically central. "The evidence of sense," for instance, "is neither the proposition itself nor some other proposition but the sensory experience one is having" (Wolterstorff, 2004, p. 83; see Plantinga, 1993, p. 98 for a contemporary externalist avowal of experience as evidence).

Regarding the phenomenal sorites argument, as we've stressed, the "experience first" approach doesn't stand or fall with the idea that argument targets. Even so, we ask that the reader consult the footnotes and decide for themselves whether the argument succeeds, and to consider as well whether it doesn't require the dubious thesis that there can be indiscernible non-identical qualia.

We remain confident that we have described a viable and vital research project, formal and informal.

## Acknowledgments

Thanks for comments are due to Robert Audi, Earl Conee, John Hawthorne, Scott Sturgeon, Mike Huemer, Kevin McCain, John Matheson, and Evan Fales.

## References

- Plantinga, A. (1993) *Warrant and Proper Function*. Oxford: Oxford University Press.  
Wolterstorff, N. (2004) Reid on common sense. In *The Cambridge Companion to Thomas Reid*, ed. T. Cuneo and R. van Woudenberg. Cambridge: Cambridge University Press.