

E = K, but what about R?

(for Maria Lasonen-Aarnio and Clayton Littlejohn, eds. *The Routledge Handbook of Evidence*, to appear)

Timothy Williamson

1. *Educated common sense and formal theory*

Holocaust deniers lack evidence that the Holocaust never took place. Even on what evidence they have, their claim is not very probable. Such observations should be commonplace. They tell against *subjective* theories of evidence, on which your evidence consists of your internal mental states, or your perceptual and intellectual seemings, or what you are certain of, or what you take for granted, or something like that. For Holocaust denial may do very well by all such subjective standards, if the denier is sufficiently consistent and blinkered. Instead, most of the relevant evidence is *publicly available*. There are better ways of deciding what happened in history than by introspection.

Thinkers with no capacity at all to recognize the difference between good and bad evidence in various areas are in no position to start theorizing philosophically about evidence: they have too little to test their theories against. Fortunately, many people do have such a capacity. We may call it *educated common sense*. But we should not expect to get far in theorizing about evidence by educated common sense alone. For it often relies on *heuristics*, quick and dirty ways of thinking, reliable enough for ordinary purposes but capable of sometimes leading us astray.¹

Here is an example. We routinely speak and think in terms of *weighing* the evidence, assessing the *balance* of the evidence, and so on. Sometimes the evidence for a hypothesis *outweighs* the evidence against, sometimes it is the other way round, and sometimes the evidence is *evenly balanced*. Thus we unreflectively use the metaphor of a pair of scales to structure our thinking about evidence. It is often helpful to do so. The metaphor has structural consequences. In particular, it implies that putting together two pieces of evidence for a hypothesis yields even stronger evidence for that hypothesis. But that is a fallacy, as probability theory soon shows. For even if the probability of a hypothesis h on evidence e_1 is very high, as is the probability of h on evidence e_2 , the probability of h on the conjoined evidence $e_1 \& e_2$ may still be very low. For instance, suppose that a number N has been chosen at random from the integers 1, 2, 3, . . . , 100; each value for N is equally likely. We receive evidence about N from a completely reliable device. Let h be the hypothesis that N is odd, and e_1 the proposition that N is one of 1, 3, 5, . . . , 49, 50 (so N is either 50 or both odd and less than 50). Then the probability of h conditional on e_1 is 25/26, so e_1 makes excellent evidence for h . Similarly, let e_2 be the proposition that N is one of 50, 51, 53, 55, . . . , 99 (so N is either 50 or both odd and over 50). Then the probability of h conditional on e_2 is also 25/26, so e_2 too makes excellent evidence for h . But the conjoined evidence $e_1 \& e_2$ entails that N is 50, which is inconsistent with h . Thus the probability of h conditional on e_1

& e_2 is 0, so the conjunction is decisive evidence *against* h . In such cases, the metaphor of weighing the evidence leads us wildly astray. Of course, in a toy model like this, we can easily see what is going wrong and correct accordingly. But in more realistic cases, where the evidence is richer, more complicated, and messier in its bearings, we are in more danger of committing the fallacy unawares.²

Thus we should be wary of relying on educated common sense in laying down structural constraints on evidential relations. Here is another example. Common sense feels strongly that nothing can be evidence for itself: that would be *circular*! Now consider a standard definition of the evidence-for relation in a probabilistic framework: e is evidence for h just in case the probability of h conditional on e exceeds the unconditional probability of h ; in other words, e raises the probability of h . To be more precise, we need to say what probability distribution is relevant, and what makes e evidence at all, but for now we can bracket those issues. In the special case when e has probability 1, the probability of h conditional on e is just the unconditional probability of h , so e is not evidence for h . Thus if e is evidence for anything, e has probability less than 1. But the probability of e conditional on itself is trivially 1. Thus, by the definition, if e is evidence for some hypothesis, e is evidence for itself.³ Consequently, the standard probabilistic definition of evidence-for clashes with common sense. If we side with common sense, we lose the systematic advantages of the probabilistic approach, for which common sense has no comparably powerful and precise substitute to offer. Indeed, common sense has a long track record of obstructing progress in the development of fruitful theoretical frameworks by digging in its heels against pseudo-anomalies in logic and basic mathematics which are really just limiting cases of a simple general rule. Common sense did not want 0 to be a number; it did not want a contradiction to entail everything; it did not want an axiom to have a one-line proof consisting of just the axiom itself. In each case, common sense was blinded by superficial differences to underlying similarities. The limiting case of evidence as evidence for itself is another such case where formal theory is a better guide.⁴

The way to defeat a theory of evidence is not by appeals to educated common sense but by producing a *better* theory of evidence. It may be better partly because it can explain more of what we pre-theoretically know about evidence. It may also be better partly because it is formally simpler and stronger. In particular, an apt challenge to any theory of evidence is to show how it integrates with a probabilistic framework. That is not to make the inappropriately Procrustean demand that all serious thinking about evidence should be done in probabilistic terms. The point is just that much (though far from all) of our most rigorous, sophisticated, and subtle thinking about evidence *is* already done in such terms, so a theory of evidence which fails to engage properly with that framework is off the pace. But there is no question of reducing the theory of evidence to the mathematics of probability theory: what constitutes evidence is a question of epistemology, not of mathematics.

In short, the appropriate methodology for theorizing on this topic, as on so many others, is *abductive*. One aspect of this methodology is that a theory of evidence can regiment the term 'evidence' as it sees fit, as long as it remains in contact with the general topic. For theoretical purposes, we want our key theoretical terms to cut the subject matter at its natural joints, wherever they are, and locating them is a task for the theory itself. We

should not expect 'evidence' as a theoretical term to trace every nook and cranny of the word as ordinarily used in English.

The rest of this chapter concerns one specific theory of evidence, proposed in *Knowledge and its Limits* (Williamson 2000) as doing justice to the demands of both formal theory and educated common sense. The present aim is not to compare it explicitly with its rivals (discussed in other chapters of this book) but to explain the resources it brings to such comparisons. The next section summarizes the theory and some considerations in its favour, mainly in their original form. Sections 3 and 4 develop the account in response to some salient challenges.

2. $E = K$

The job description for evidence includes three general tasks. First, evidence *rules out* some hypotheses by being inconsistent with them. Second, in *inference to the best explanation* (closely related to abduction), evidence is what the best hypothesis best explains. Third, in *probabilistic confirmation*, evidence is what the probability of the hypothesis is conditional on.

All three tasks require evidence to be *propositional*, true or false in various circumstances. What is inconsistent with a hypothesis is propositional: its truth-condition is disjoint from the theory's. What a hypothesis explains is propositional: the hypothesis explains *why* or *how* its truth-condition obtains. What the probability of a hypothesis is conditional on is propositional: the probability of h conditional on e is normally defined as the probability of the conjunction of h and e divided by the probability of e , where both h and e are treated as 'events' in the probability space, which are in effect truth-conditions. Of course, the term 'evidence' is often applied more widely: for example, to material evidence in a court of law, and by some philosophers to mental states. But when it comes to the three tasks of evidence just described, what does the work is not the non-propositional evidence itself but some *proposition* about it. Thus for theoretical purposes it is more perspicuous to restrict the term 'evidence' to the propositional.

Presented with an account of what constitutes evidence, one should ask: which propositions does it propose as evidence? Unfortunately, many accounts fudge this elementary question. For example, philosophers who speak of 'intuitions as evidence' often leave it unclear whether the proposition supposed to be in one's evidence is the proposition one intuits or the proposition that one intuits it.

To serve the three tasks, what is needed is not just a proposition but a *true* proposition. That a hypothesis h is inconsistent with the evidence rules out h only by entailing that h is false; but if the evidence were false, it would be inconsistent with some true hypotheses. Inference to the best explanation and probabilistic confirmation would also be more problematic if the evidence to be explained or conditionalized on were false. Of course, in applying the principle that all evidence is true, one must be careful to distinguish the proposition h (or e) from the proposition h^* (or e^*) that h (or e) holds *approximately*, otherwise one can easily generate fallacious 'counterexamples'.

Clearly, the question is not just what evidence is ‘out there’, but what evidence a given agent *has*—which true propositions they have access to in the relevant epistemic sense. The salient form of access is simply *knowledge*. Thus your total evidence is just the totality of truths you know. This is the equation $E = K$. Since you know p only if p is true, this immediately explains why evidence has to be true.

If you do not know p , p is not available to serve as part of your evidence (though it may seem to be). Conversely, if you do know p , p is available to serve as part of your evidence. Although someone might be tempted to restrict your evidence to truths you know in some special, supposedly basic way—for instance, by sense perception, or by introspection—the case of evidence for the Holocaust already suggests that no such restriction is faithful to the actual role of evidence in our cognitive lives. The publicly available evidence for the Holocaust consists of all sorts of historical facts.

Such an account of evidence is easy to integrate with a probabilistic framework. The probability of a hypothesis h on your evidence is simply the probability of h conditional on the conjunction of all the truths you know. Unlike views which allow false evidence, this view does not have to deal with the possibility of inconsistent evidence, probabilities conditional on which are ill-defined, for truths are always mutually consistent.

That all evidence is true does not mean that evidence supports only true hypotheses. Of course, the evidence for a false hypothesis cannot be logically conclusive, since only truths follow deductively from truths. But a falsehood can still be highly probable on someone’s evidence, for the conditional probability of a falsehood on a truth can be high. For example, when someone is skilfully framed for a crime she did not commit, the false hypothesis that she is guilty is highly probable on true evidence as to where fingerprints were found, what witnesses said, and so on. True evidence can be *misleading*.

One salient consequence of $E = K$ is that we are often in no position to know whether a given proposition belongs to our evidence, for we are often in no position to know whether we know something. For example, someone sees a mule cleverly painted to look like a zebra in the zoo; he falsely believes that he saw a zebra; indeed, he falsely believes that he *knows* that he saw a zebra, and that the proposition that he saw a zebra is part of his evidence; for all he knows, the proposition that he saw a zebra *is* part of his evidence. Thus sometimes a proposition is *not* part of one’s evidence, even though one is in no position to know that it is not part of one’s evidence (*negative introspection* fails for evidence). More subtly, it can be argued, sometimes a proposition *is* part of one’s evidence, even though one is in no position to know that it is part of one’s evidence (*positive introspection* fails for evidence). In these ways, it is not transparent to one’s evidence what one’s evidence includes.⁵

The risk of error is not the only reason for the non-transparency of evidence. Consider an animal with perceptual evidence (that is, knowledge) streaming in through various sense modalities. Virtually none of that evidence is of the form ‘*my evidence does not include p*’. If it happens to include some evidence incompatible with p , a clever animal may be able to deduce ‘*my evidence does not include p*’ (because evidence is always true), but if p is independent of all its actual evidence, how is it to recognize that? Even if it somehow manages to make its evidence satisfy positive introspection, and to derive all the logical consequences of combined elements of its evidence, negative introspection may still

fail. For none of the propositions obtained in those ways need entail anything special about p . The animal would need a further capacity to survey all its evidence and register the absence of p . That might be an extremely demanding cognitive task, especially if the evidence from different sense modalities comes in different formats (verbal, visual, olfactory, ...) and its entailments are hard to compute. Even if it has some capacity to survey its evidence, it will not be *instantaneous*—certainly it is not in the human case. The animal may have to wait for a signal that p has been found. If no signal comes within, say, five minutes, the animal may give up and assume '*my evidence does not include p* ' (human memory sometimes works like that). But while the animal is waiting for a signal, for all it knows the signal will still come, and so its evidence does not yet include '*my evidence does not include p* ', even if in fact its evidence does not include p . At that time negative introspection fails for the animal's evidence. Should we dismiss such creatures as *irrational*? In doing so, we risk dismissing ourselves as irrational. Of course, humans *are* irrational to some extent, but epistemology becomes uninteresting if we idealize away from too many of the limitations inherent in real-life knowledge.

The non-transparency of evidence can easily look problematic. For if you are in no position to know what your evidence includes, you may well be in no position to know how probable a proposition is on your evidence, and so be in no position to be guided by your evidence on how to treat that proposition. But guidance on how to treat propositions is just what we want evidence for. Thus the result might be used as an objection to $E = K$.

However, non-transparency is an objection to $E = K$ only if some other approach does better. But, arguably, other approaches to evidence end up facing the same problem. For example, suppose that one's evidence consists of those truths one currently introspects. It is often hard to know just which truths one currently introspects. Inner life is an unclear, fleeting, elusive affair—at least in my case (perhaps the reader is luckier). I find it on the whole much easier to assess what I observe of my external physical environment than what I introspect of my own internal mental world. Going internalist does not solve the problem of non-transparency.

Of course, to say that non-transparency is everyone's problem is not yet to solve it. How can we live with non-transparent evidence? The next section gets to grips with that challenge.

3. *What is R?*

Since the problem of non-transparency was introduced with reference to sceptical scenarios, let us consider them further.

A sceptical scenario is usually characterized as a bad case by contrast with a corresponding good case. For example, in the good case, the thinker sees '*there is blood on the knife*' in the normal way, and thereby knows and believes '*there is blood on the knife*'. In the bad case, there is no bloodied knife, but a Demon causes the thinker to have a visual experience as of seeing '*there is blood on the knife*'; the total experience is perfectly indiscriminable 'from the inside' from the total experience in the good case; she believes falsely '*there is blood on the knife*', without seeing or knowing that there is. Internalists

typically insist that the thinker has *exactly the same evidence* in the good and bad cases, and conclude that the belief '*there is blood on the knife*' has exactly the same level of justification in the two cases. Given $E = K$ instead, the evidence is not exactly the same in the two cases. In the good case, the thinker's evidence includes the proposition '*there is blood on the knife*', since she knows that there is. In the bad case, by contrast, the thinker's evidence does not include the proposition '*there is blood on the knife*', since she does not know that there is. Thus there is scope to deny that the belief '*there is blood on the knife*' has exactly the same level of justification in the two cases.

Thomas Kelly (2016) uses this example to raise a problem for $E = K$. He begins by distinguishing two different 'intuitions':

The first intuition is that a thinker in the bad case has *exactly* the same evidence as a thinker in the good case. Perhaps abandoning this intuition is not much of a cost (if it is any cost at all). A different intuition is the following: when a thinker in the bad case takes his experiences at face value and forms beliefs about the external world in the usual manner, those beliefs are not simply unreasonable, in the way that they would be if, for example, the thinker adopted those same beliefs on a whim, or in the absence of any reason to do so at all. Abandoning this intuition would seem to be a much heavier price to pay.

Kelly goes on to argue that the second intuition is hard to reconcile with $E = K$ in the example, since in the bad case the thinker may know nothing to give significant support to the proposition '*there is blood on the knife*'. He considers true fallback propositions such as '*it appears that there is blood on the knife*' and '*my experience is as of there being blood on the knife*', but he points out that, on some plausible views, in normal cases of perception, a thinker who does not suspect that something is amiss will not even consider such fallback propositions, but will simply go straight ahead and form beliefs about the external world, such as '*there is blood on the knife*'. Presumably, if she does not even consider the fallback propositions, she does not come to believe them either. But then, since knowledge entails belief, she does not know the fallback truths. Therefore, given $E = K$, they do not belong to her evidence. Thus $E = K$ threatens to eliminate any evidence for the proposition '*there is blood on the knife*'. But, as Kelly says, 'Intuitively, this belief is at the very least better justified than it would be in the absence of the relevant visual experience'. Even if some beliefs can be justified without evidence, the false belief '*there is blood on the knife*' is hardly a good candidate for that status.

One might try to respond to Kelly by being more liberal in ascribing knowledge, arguing that in some relevant sense the thinker does implicitly know '*it appears that there is blood on the knife*' or '*my experience is as of there being blood on the knife*', even though all she explicitly formulates is '*there is blood on the knife*'. But a young child might be able to think simple thoughts about the external world such as '*there is blood on the knife*' while still lacking the semantic sophistication to think meta-cognitive thoughts such as '*it appears that there is blood on the knife*' or '*my experience is as of there being blood on the knife*': to ascribe implicit knowledge with the latter contents to the child is quite a stretch. So let us just concede to Kelly that the thinker does not know any of the relevant fallback truths; thus by $E = K$ she lacks evidence for her belief '*there is blood on the knife*'. What then?

When formulating the key intuition in the quoted passage, Kelly describes the thinker's beliefs as 'not simply unreasonable'. Like the distinction between rationality and irrationality, the distinction between reasonableness and unreasonableness can be understood in two contrasting ways. We can focus either on the proposition thought or on the agent thinking it. If we focus on the proposition, we naturally wonder what evidence there is for it. If instead we focus on the agent, we naturally wonder what cognitive dispositions they are exercising. Of course, the two questions are not completely independent of each other. To a first approximation: good cognitive dispositions typically generate knowledge, or at least beliefs well supported by the thinker's evidence; bad cognitive dispositions typically generate beliefs poorly or not at all supported by the thinker's evidence, and so not knowledge. Nevertheless, the two questions are far from equivalent. In the bad case where there appears to be blood on the knife, nothing is wrong with the thinker's cognitive dispositions; the relevant ones are just those involved in acquiring ordinary visual knowledge. However, the thinker is exercising those good cognitive dispositions in unfavourable circumstances, where they unluckily generate beliefs unsupported by his evidence. There is also a converse phenomenon, which Maria Lasonen-Aarnio (2010) calls 'unreasonable knowledge': sometimes, the thinker exercises bad cognitive dispositions in favourable circumstances, where they luckily generate knowledge (see also Lasonen-Aarnio 202X, 202Y).⁶

In the bad case, our positive assessment of the thinker's cognitive dispositions naturally leaches into our assessment of the belief they generate: it has *good origins*. But we should not confuse that positive feature of the belief with a quite different positive feature, which it lacks: support from the thinker's evidence. Thus proponents of $E = K$ can meet Kelly's challenge to explain why the thinker's beliefs in the bad case 'are not simply unreasonable, in the way that they would be if, for example, the thinker adopted those same beliefs on a whim, or in the absence of any reason to do so at all'. For beliefs adopted 'on a whim, or in the absence of any reason to do so at all' are paradigms of beliefs with *bad origins*, generated by cognitive dispositions which typically generate beliefs not supported by the thinker's evidence. The mistake is to assume that the difference in reasonableness must be specifically a difference in evidential support, rather than in the quality of the underlying cognitive dispositions. That is why this section is entitled 'What is R?'; 'R' stands for 'reasonableness' (or, to similar effect, 'rationality').

For completeness, we can also briefly consider the epistemic status of the thinker's belief '*there is blood on the knife*' in the *good* case, where by hypothesis it constitutes perceptual knowledge, and so perceptual evidence, given $E = K$. As the case has been developed, that knowledge does not rest on evidence such as '*it appears that there is blood on the knife*' or '*my experience is as of there being blood on the knife*': those fallback propositions are no more believed and no more known in the good case than they are in the bad case. Given the richness of vision, the thinker's visual knowledge in the good case is presumably not exhausted by the single proposition '*there is blood on the knife*', but that does not mean that her visual knowledge '*there is blood on the knife*' is somehow based on other visual evidence of a radically different kind. Given $E = K$, the knowledge '*there is blood on the knife*' is not strictly unsupported by evidence, because it is itself evidence, and so trivially has probability 1 given the thinker's evidence, but that is not a (feeble) attempt to

understand knowledge in terms of evidential relations. Rather, the point is that there is a limit to how much heavy lifting we should expect evidential relations to do in epistemology: yet again, knowledge comes first.

This treatment of good and bad cases makes essential use of norms on cognitive dispositions. Such norms require further consideration.

4. *Anti-agnostic norms*

The cognitive dispositions at issue in section 3 were dispositions to acquire knowledge or belief. As a first approximation, the dispositions were qualified as ‘good’ or ‘bad’ according to the quality of the cognitive states they were dispositions to generate—specifically, the evidential status of the beliefs, with knowledge as the best status. That may suggest that norms on dispositions to believe are derivative from norms on belief. However, there is more to it than that.

On one natural understanding, a norm on belief is broadly *functional* in nature: beliefs which violate it are *defective*. This approach has the advantage of applying straightforwardly to the beliefs of unsophisticated cognizers such as young children and non-human animals. After all, such creatures act on their knowledge of their external environment, and on their beliefs about it; they are just as vulnerable as adult humans to many sceptical scenarios. Now, for a norm of non-defectiveness on belief to be violated, there must first *be* a belief: without a belief, there is nothing to be defective. Analogously: for the norm of promise-keeping to be violated, a promise must first have been made: without a promise, there is nothing to be broken. One can comply with the norm of promise-keeping simply by not making any promises. Similarly, one can in principle comply with a norm of non-defective belief simply by not having any beliefs. Indeed, that is what extreme sceptics recommend. If one has no beliefs, one has no beliefs with a defective evidential status. If the only norm in play is one banning defective belief, by far the simplest way of complying with it would be to avoid believing altogether. By that standard, Kelly’s thinker’s belief ‘*there is blood on the knife*’ is unjustified in the good case as well as the bad one, since even in the former it constitutes an unnecessary risk of defectiveness.

Of course, to avoid believing altogether is psychologically impossible for normal humans, and for other animals which naturally form beliefs. But a determined sceptic could probably avoid forming a particular belief, such as ‘*there is blood on the knife*’. Thus believing becomes a frailty of our lower nature.

That line of thought concedes far too much to scepticism. Suspension of all belief is not an ideal. It is pathological even on the functional approach with which we started. Agnosticism is dysfunctional. Animals *need* knowledge of their environment on which to act. For both predators and prey, it is a matter of life and death. Since animals need knowledge of their environment, they need a cognitive system to produce that knowledge, and plenty of it. The system has to be *open to knowing*. The inevitable downside of such a system is that it will be *open to believing falsely* in unfavourable circumstances, of which sceptical scenarios are just an extreme case. Your cognitive system would be defective if it did *not* dispose you to believe ‘*there is blood on the knife*’ in the bad case. What the argument of

section 3 really requires is a norm of non-defectiveness on whole cognitive systems, not just a norm of non-defectiveness on individual beliefs.⁷

A good cognitive system is positively disposed to gather knowledge, not just negatively disposed to gather *only* knowledge. There is a trade-off between the two dispositions: too high a threshold means gathering too little knowledge, too low a threshold means gathering too much non-knowledge. A good system finds an appropriate mean between the two extremes.

Given $E = K$, to gather knowledge *is* to gather evidence. Thus a good cognitive system is positively disposed to gather evidence. This point helps with cases in the literature of agents who manifest unconscious bias in evidence-gathering: they behave in ways which make them unlikely to acquire evidence that would challenge their prejudices. Their decision-making may be reasonable given their evidence, but their evidence constitutes a severely biased sample (see Miracchi 2019 and Lasonen-Aarnio 202Y).

In an extreme case, all the agent's beliefs constitute knowledge. For example, in a job search, he may pay much more attention to the CVs of candidates with the desired race and gender profile than to the CVs of other candidates, with the result that he acquires much more knowledge about the former candidates than the latter. Naturally, some of the knowledge may tell in favour of the candidate, some may tell against. But in a large field, with many well-qualified candidates, probably some candidate with the desired profile will have a good enough CV for his evidence about them to be more favourable than his much sparser evidence about any of the candidates without the desired profile. The agent may even come to *know* that his evidence favours that candidate with the desired profile more than it favours any other candidate, and accordingly decide to vote for that candidate. That process does not *have* to involve any belief which falls short of knowledge—though of course it is likely to.

A norm of non-defectiveness on individual beliefs will not pick up what has gone wrong in such a process. A norm of evidence-gathering will help. The agent was insufficiently open to acquiring knowledge from the CVs of candidates without the desired profile. Given $E = K$, the agent was insufficiently open to acquiring evidence from those CVs. 'Openness' here is not a merely passive quality; it involves a disposition to *seek* knowledge. Sometimes the seeking is as easy as just reading the page in front of you; sometimes, of course, it is much harder, and involves long and arduous searching.

However, a norm of knowledge-gathering must be constrained by feasibility. One cannot search simultaneously for knowledge relevant to all the questions one is capable of asking. One must have priorities, and those priorities will often be practical rather than epistemic. In the job search case, the priority was high because an important decision was at stake. A secretary with no role in the decision-making, idly leafing through the CVs with the same levels of attention as the decision-maker, does not count as equally negligent, because their priorities are different, for practical rather than epistemic reasons. Even the decision-maker may not count as negligent until the moment of decision. For, having looked through the CVs and recognized that his present evidence favoured one of the candidates, he could have decided that it was time to read through all the CVs carefully. Instead, he simply formed the intention to vote for that candidate. That was negligent, because it was based on inadequate evidence-gathering, but its inadequacy was *as a basis for an important*

practical decision. If he also formed the belief '*This is the best candidate*', then that belief did not constitute knowledge and was based on epistemically inadequate evidence, but he could have jumped to the intention without forming the belief. Thus we cannot expect a norm of knowledge-gathering to do all the normative work in such cases, without reference to practical considerations.

Still, we should not go to the opposite extreme either, and expect to absorb the epistemic into the practical. Cognitive systems have considerable autonomy, in part for good practical reasons. One often does not know in advance what knowledge is going to come in handy. When our cognitive systems are working properly, knowledge streams into us during waking hours. Animals are naturally curious: knowledge is comparatively cheap to acquire and store, and too little of it is usually worse than too much. Knowledge-seeking is not confined to scholars: the nosy and the prying enjoy it too. Not all evidence-gathering takes the dignified or pompous form associated with the phrase '*intellectual virtue*'. A non-defective cognitive system is constantly on the look-out for cheap knowledge: what requires cultivation is the taste for *expensive* knowledge.

The normative role of evidence is best understood within a framework of functional norms on cognitive systems. The equation of evidence and knowledge integrates smoothly with such a framework, which also naturally hosts a crucial distinction between evidential and dispositional forms of rationality.

Notes

- 1 Williamson 2020a makes a more detailed case that heuristics play a significant role in generating various philosophical problems, in particular concerning the semantics of conditionals. The term 'heuristic' is intended to recall both the heuristics and biases programme in psychology (going back to works such as Kahneman, Slovic, and Tversky 1982) and the partly contrasting tradition of work on adaptive rationality (see for Gigerenzer, Hertwig, and Pachur 2011), though the exact connections are left open.
- 2 The problem is not restricted to our thinking about evidence. We apply the same metaphor in our thinking about reasons: 'weighing reasons', 'the balance of reasons', etc. Unfortunately, the literature on reasons is less used to applying probability theory as a corrective.
- 3 If e has probability 0, probabilities conditional on e are usually taken to be undefined, but then e is not evidence for anything.
- 4 Resistance to evidence being evidence for itself is central to the critique in Brown 2018 of the account of evidence defended in this chapter and in Williamson 2000.
- 5 The case in Williamson 2000 for the non-transparency of evidence is further developed in Williamson 2014, 2019, 2020b. For some recent dissent see Salow 2018.
- 6 For more discussion of this contrast see Williamson 2017, 202X, 202Y.
- 7 For more discussion of norms on cognitive systems see Williamson 202Y.

References

- Brown, Jessica Brown. 2018: *Fallibilism: Evidence and Knowledge*. Oxford: Oxford University Press.
- Gigerenzer, Gerd, Hertwig, Ralph, and Pachur, Thorsten (eds.). 2011: *Heuristics: The Foundations of Adaptive Behavior*. New York: oxford University Press.
- Kahneman, Daniel, Slovic, Paul, and Tversky, Amos (eds.). 1982: *Judgment under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.
- Kelly, Thomas. 2016: 'Evidence', in Edward Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition),
<<https://plato.stanford.edu/archives/win2016/entries/evidence/>>.
- Lasonen-Aarnio, Maria. 2010: 'Unreasonable knowledge', *Philosophical Perspectives*, 24: 1-21.
- Lasonen-Aarnio, Maria. 202X: 'Virtuous failure and victims of deceit', forthcoming in Julien Dutant and Fabian Dorsch (eds.), *The New Evil Demon*. Oxford: Oxford University Press.
- Lasonen-Aarnio, Maria. 202Y: 'Dispositional evaluations and defeat', forthcoming in Jessica Brown and Mona Simion (eds.), *Reasons, Justification and Defeat*. Oxford: Oxford University Press.
- Miracchi, Lia. 2019: 'When evidence isn't enough: suspension, evidentialism, and knowledge-first virtue epistemology', *Episteme*, 16: 413-437.
- Salow, Bernhard. 2018: 'The externalist's guide to fishing for compliments', *Mind*, 127: 691-728.
- Williamson, Timothy. 2000: *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, Timothy. 2014: 'Very improbable knowing', *Erkenntnis*, 79: 971-999.
- Williamson, Timothy. 2017: 'Ambiguous rationality', *Episteme*, 14: 263-274.
- Williamson, Timothy. 2019: 'Evidence of evidence in epistemic logic', in Mattias Skipper and Asbjørn Steglich-Petersen (eds.), *Higher-Order Evidence: New Essays*, 265-297. Oxford: Oxford University Press.
- Williamson, Timothy. 2020a: *Suppose and Tell: The Semantics and Heuristics of Conditionals*. Oxford: Oxford University Press.
- Williamson, Timothy. 2020b: 'The KK principle and rotational symmetry', *Analytic Philosophy*, forthcoming.
- Williamson, Timothy. 202X: 'Justifications, excuses, and sceptical scenarios', forthcoming in Julien Dutant and Fabian Dorsch (eds.), *The New Evil Demon*. Oxford: Oxford University Press.
- Williamson, Timothy. 202Y: 'Epistemic ambivalence', forthcoming in Nicholas Hughes (ed.), *Epistemic Dilemmas*. Oxford: Oxford University Press.