

## **E=K**

Brian Weatherson

First Posted: August 29, 2003; Last Revised: December 1, 2009

### *Abstract*

Timothy Williamson has argued that our evidence is what we know. This implies that anything we come to know by inference instantly becomes part of our evidence, and that all of our evidence is true. I argue that neither of these implications is correct. I conclude by noting a rival theory of evidence, one based on a suggestion Jerry Fodor makes in *The Modularity of Mind*, is not vulnerable to the criticisms I make of Williamson, nor to the criticisms he makes of traditional theories of evidence.

In chapter 9 of *Knowledge and its Limits*<sup>1</sup> Timothy Williamson argues that our evidence includes all and only the propositions we know. This view has a number of striking consequences. It entails, among other things, that all evidence is propositional, that all evidence is true and that our evidence is massively redundant. The first two consequences should be obvious. The last consequence follows from the fact that for most propositions  $p$  that we know, we know lots of propositions that entail or at least inductively support  $p$ , and lots of propositions that  $p$  entails or at least inductive supports. After arguing that two of these consequences of the theory that evidence is knowledge (E=K) are implausible, I will sketch a theory of evidence that doesn't have the undesirable consequences.

### *1. Redundancy*

Here's an argument that Williamson uses against the theory that our evidence includes all of our justified true beliefs (200-1). Assume we see that  $p$  and infer  $q$  from it using probabilistic reasoning. In this case, says Williamson, we do not know  $q$  even if we justifiably truly believe it. There's an interesting

---

<sup>1</sup> All page references until section 3 to this book.

asymmetry in such a case between  $p \wedge \neg q$  and  $q \wedge \neg p$ . Although we (truly, justifiably) believe each of these to be false, we would say that  $q \wedge \neg p$  is inconsistent with our evidence while  $p \wedge \neg q$  is not. If evidence included all justified true beliefs, this asymmetry would not be supported. Williamson concludes, rightly, that our evidence does not include all of our justified true beliefs.

Unless we are willing to embrace a very strong form of scepticism, the same argument can be used against E=K. Assume we see that  $p$  and infer  $q$  from it using inductive reasoning. Unless induction can never provide knowledge (a very strong sceptical claim) it must be possible in such cases to come to know  $q$ . But even in those cases the asymmetry between  $p \wedge \neg q$  and  $q \wedge \neg p$  remains. Intuitively, in such a case  $q \wedge \neg p$  is inconsistent with our evidence while  $p \wedge \neg q$  is not. But this is inconsistent with E=K, since both  $p$  and  $q$  are known and hence, according to E=K, part of our evidence.

The intuitions here may be stronger with a particular example. On Thursday night Trinity watches the new hit movie *The Moroccan Falcon*. She sees that the hero dies in the final scene. Her friend Magdalen is going to see the movie Friday night. If Trinity tells Magdalen that when she saw the movie the hero died in the final scene, Magdalen might well complain that the movie has been spoiled because she now know how it ends. It wouldn't really do for Trinity to argue that she is only reporting how the movie ended *on the Thursday screening*, and that for all anyone knows it might end differently on Friday. The mere possibility that the movie has alternate endings hardly defeats Magdalen's claim that she now know how the movie will end.<sup>2</sup>

Let  $p$  be the proposition that the hero dies in the final scene of the Thursday screening of *The Moroccan Falcon*. Let  $q$  be the proposition that the hero dies in the final scene of the Friday screening of *The Moroccan Falcon*. Consider Trinity's relationship to the propositions  $p \wedge \neg q$  and  $q \wedge \neg p$ . I've argued that she knows that each of these are false. So according to E=K the falsity of each conjunction is part of

---

<sup>2</sup> Not only that, but Magdalen now knows how the movie will end on Saturday and Sunday and Monday as well.

Whether this is a counterexample to the occasionally heard claim that induction on a single case is never sufficient for knowledge is left as an exercise for the interested reader.

her evidence.<sup>3</sup> Intuitively,  $q \wedge \neg p$  is inconsistent with Trinity's evidence since she saw that  $p$  is true, while  $p \wedge \neg q$  is consistent with her evidence, though not with what she knows. So  $q$  is not part of her evidence even though she knows it, so  $E \neq K$ .

While intuition here provides a *prima facie* reason for thinking  $q$  is not part of Trinity's evidence, it doesn't provide much insight into why  $q$  isn't part of her evidence. One tempting thought is that propositions that are first believed on the basis of inference cannot be part of one's evidence. The picture here is that evidence is a kind of foundation, and the structure that is built on that foundation does not include any evidence.<sup>4</sup> This would be a rather hasty generalisation from what the case tells us. And it seems that there are two interesting kinds of counterexamples to the generalisation, though it is not clear that either undermine the intuitive picture. The first counterexample is where an agent comes to know  $q$  by inference, then forgets that  $q$ , then later sees that  $q$ . At the last stage  $q$  is clearly part of her evidence, but  $q$  is something she first came to know by inference. The second counterexample is more contentious. If one comes to know  $q$  by inference, ceases to *consciously* believe  $q$ , and later retrieves  $q$  from memory for use in another inference, it is arguable that  $q$  then is part of one's evidence.<sup>5</sup>

Neither of these counterexamples threatens the following weaker generalisation: at the time one comes to know  $q$  solely on the basis of inference,  $q$  is not part of one's evidence. This captures the

---

<sup>3</sup> Actually, she probably knew each of these conjunctions before watching the Thursday screening, but that won't be relevant to the following argument.

<sup>4</sup> Note that the structural relations in this picture need not be *justifying* relations, as they would be on a foundationalist picture. They might simply be causal relations. A coherentist could think that our evidence causes us to have beliefs, and these beliefs are in turn justified by their relations to other beliefs. So while the picture here is most naturally compatible with a kind of foundationalism, it does not *require* foundationalism.

<sup>5</sup> On the view I'll sketch below, whether  $q$  is part of one's evidence then depends on empirical facts about the working of memory. As far as I can tell the empirical evidence is that memory is not a source of evidence as I'm going to construe evidence, but the important point is that whether it is a source of evidence is something to be settled empirically, not by philosophical fiat.

intuitive idea that evidence is the *ground* for what we know, and is consistent with our intuitions about the Trinity/Magdalen case, without being vulnerable to these counterexamples. This generalisation has the consequence that in a chain of reasoning our evidence is not massively redundant in that it usually does not include what is inferred at every step, even if we know that each of these intermediate steps are true.

## 2. *Truth*

If  $E=K$  then all evidence is true. In cases of illusion, however, it seems plausible to say that our evidence is false. When it looks like there's a tiger behind that tree, our evidence is that there is a tiger behind that tree. (That's why we are justified in running away. Fast.) If there is no tiger, then we cannot know that there is a tiger, so if  $E=K$  it cannot be part of our evidence that there is a tiger. So it seems  $E=K$  is inconsistent with intuition here.

Williamson (198-200) is aware of this objection, and answers it by saying that in such a case our evidence is that there *appears* to be a tiger behind the tree. Presumably this is what justifies our running away, since it makes the probability of there actually being a tiger rather high, and hence the expected utility of not running rather low. There are, however, two problems with this response.

First, as Williamson notes, a creature that lacks the concept APPEARS cannot know that there appears to be a tiger behind the tree, and hence that proposition cannot be part of its evidence. Williamson denies that this creature gets any (relevant) evidence in this case. (Perhaps it knows there is no elephant behind the tree, but that's hardly relevant in the situation.) This makes it hard to explain our intuitions about the creature's actions.

Assume, to make the case as clear as possible, that the creature is generally a fairly reliable perceiver of predators and prey. I'll assume this creature has few concepts; it has PREDATOR, and it has some ostensive directional concepts (like THAT WAY) and it has the concept OPPOSITE. Finally, I'll assume that it has the desire not to be eaten hardwired, so it doesn't have to do any processing to move from *predator that way* to *run opposite way*. Which, in this case, it does.

This seems to be a fairly justified reaction in the situation. But it's a platitude that potentially costly actions based on no evidence are not justified. And since running certainly involves some costs, and according to  $E=K$  the creature has no evidence to support the run, according to  $E=K$  the reaction is *not* justified. This seems like a fairly serious strike against  $E=K$ .

Does this objection to  $E=K$  generalise to objections against other factive theories of evidence? It seems like it should. The argument against  $E=K$  is as follows:

1. The creature's (relevant) evidence is either that there is a tiger or that there appears to be a tiger.
2. If  $E=K$ , the creature's evidence is not that there is a tiger (Since evidence is true)
3. If  $E=K$ , the creature's evidence is not that there appears to be a tiger (Since evidence must be conceptualised)
4. The creature has some (relevant) evidence. (Since something justifies the run)

C.  $E \neq K$

It seems that many theories other than  $E=K$  can be substituted into that argument without losing soundness. Any theory on which all evidence is true will support premise 2. And any theory on which the creature must have the concept APPEARS to have evidence that there appears to be a tiger will support premise 3. But if all it takes to have a concept is to have states representing propositions that include that concept as a constituent, and if having one's evidence include the proposition  $p$  involves representing  $p$ , then premise 3 will be true on any theory of evidence. I'm inclined to think both antecedents here are true, so the argument generalises to an argument against any factive theory of evidence, but defending those claims would take us too far afield.

Also note that if the argument in section 1 against the massive redundancy of evidence worked, then we should be suspicious of claims that *our* evidence includes such closely related propositions as *there's a tiger* and *there appears to be a tiger*. Presumably one of those really is part of the evidence, and the other is inferred. I'm inclined to think that the former really is part of the evidence and the latter is

inferred – our evidence is in the first place about the world not about ourselves. But then it's implausible that in cases of illusion our evidence is really about appearances. If there seem to be tigers behind both the left and right trees, it seems our evidence about the two trees' dangers should be alike. If evidence is both sparse and (inevitably) true then our evidence about one tree is *there's a tiger behind it* and about the other is *there appears to be a tiger behind it*, which is implausible.

### 3. An Alternative Theory

Much of Williamson's discussion of evidence consists of arguments against the phenomenal conception of evidence and its philosophical supports. Williamson argues that evidence does not entirely consist of sense data, phenomenal states or anything similar. He argues that the reasons for thinking it does, in particular the theses that our evidence must be infallible or that we must know our evidence, are indefensible. I agree with these arguments. Indeed, I think Williamson does not go far enough. I think that even in cases of illusion our evidence is about the world not about our impression of it, and that evidence can not only be fallible, it can be false.

The attacks on the phenomenal conception of evidence provide indirect support for E=K by knocking out its most prominent rival. If there is no other theory of evidence on the market, E=K might win by knockout. I doubt Williamson intends to argue *for* E=K this way, because such an argument would have two major lacunae. First we need a premise that there is a theory of evidence, and second we need the premise that such a theory is on the market. In one sense I think the first missing premise here is false. We might be able to say something substantive about what evidence actually is, but there is no (simple) substantive theory of evidence with as much generality as E=K. Williamson is agnostic about whether E=K is *a priori*, but he says it is a necessary truth. I suspect our best theory of evidence will be contingent.

David Lewis's work on defining theoretical terms provides a useful framework for thinking about what evidence is. In summary form, we first isolate the concept EVIDENCE by rounding up the platitudes involving evidence, then we do some empirical work to find out what actually satisfies most of

those platitudes. (See Lewis 1970, Jackson 1998 for more details.) Some of the platitudes are that evidence is what grounds knowledge, that action based on no evidence is *prima facie* irrational, that evidence is normally a guide to the truth, that evidence is not always redundant in the sense set out above and that in practice much of our evidence is perceptual. What might satisfy these platitudes?

That's an empirical question and the history of philosophers trying to answer empirical questions is dubious at best. But let's end with some speculations about how it might go. If the relevant parts of cognitive science don't turn out happily, it might turn out that evidence is a massively disjunctive concept. *p* is part of our evidence iff we see that *p* or hear that *p* or smell that *p* or, etc. A disjunctive account like this may satisfy many of the platitudes, but it would still be unfortunate not only because disjunctive theories are in general unattractive, but because we have few constraints on where to stop the list. To choose one salient example, is what we remember part of our evidence on this view of evidence? It's not only hard to say, it's even hard to know at this stage what empirical results would settle the matter.

It is possible that cognitive science will end up delivering us a prettier account of evidence. Let's assume, for the sake of discussion, that Fodor's theories about how modular the mind is are basically correct<sup>6</sup>. On Fodor's picture, most (if not all) input systems are modular. That is, they are 'domain specific' (47), their operation is 'mandatory' (52), their operation is typically inaccessible to central processors (56), they are fast (61), they are informationally encapsulated (64) and they have 'shallow' outputs (86). In (2000) Fodor stresses the idea that modules are informationally encapsulated, so it's worth focussing on that for a bit.<sup>7</sup>

---

<sup>6</sup> See Fodor 1983, 2000 for details. All page references from now on to Fodor 1983 unless otherwise stated. It's possibly worth noting that Fodor's theory here is supported by an impressively small number of cognitive scientists. Fodor's defences seem fairly compelling to me, but the (vast) majorities are in other camps.

<sup>7</sup> For definiteness I will take modularity to mean information encapsulation plus most of the other features Fodor lists here. As Fodor notes (2000: 56) there have been a lot of different things meant by 'modularity' over the years, so we may as well just stipulate a meaning and move on rather than try and work out *a priori* what it must mean.

What information our visual system, or our language processing system, delivers to us is quite often insensitive to how much we know. We could have been reading all day about the Müller-Lyer illusion, yet when we see a good illustration of it one of the lines will still look longer than the other. This is, one hopes, not typical of our cognitive systems. Once you've learned that affirming the consequent is a sin, the premises *If p then q* and *q* don't even seem to provide the basis of an argument for *p*.<sup>8</sup> The general point is that what information we possess can be brought to bear on what inferences we're prepared to draw. As Fodor stresses (at length) this feature is really crucial to our ability to perform abductive inference, since that requires appreciation of all the relevant factors. Often enough, taking account of all the relevant information at our disposal *is* what we (approximately) do in standard reasoning. But it isn't what we, or at least our visual system, does in visual processing. You might have the information that you're looking at a computer screen so clues about depth are to be ignored, but your visual processor doesn't have access to that information so it doesn't ignore the clues and you (again) see one of the lines as longer than the other. If the visual system wasn't so restricted in the information it had access to, maybe it wouldn't be so fooled.<sup>9</sup>

On the basis of reflections like these, Fodor concludes that not only are our input systems modular, our inferential systems are not modular. He suggests that this distinction means that the outputs of the modules might play the role of 'observations' in empiricist philosophy of science. I think that's an

---

<sup>8</sup> Or at least I hope they don't. If I'm wrong here then a lot of us are probably wasting our time teaching logic classes.

<sup>9</sup> On the other hand, maybe it would be so slow that you, or more likely your distant ancestors, would be eaten by tigers. It's not an adaptation to be disposed to think about whether there is a tiger behind the tree. Or at least it wasn't back when that kind of thing was relevant. For all I know it could be adaptive now, at least if the right kind of conspecifics are attracted to that kind of thing.



excellent suggestion.<sup>10</sup> Doing away with the empiricism altogether, I propose that the outputs of *reliable* modules constitute our evidence. (I mean to use an extensional conception of reliability here: a module is reliable just in case it usually outputs truths.) The restriction to reliable modules is to satisfy the platitudes that evidence grounds knowledge and that evidence is a guide to the truth. Defending this would take us well outside our main topic, but I think this proposal captures what is right about process reliabilism, and does so in a way that is invulnerable to the generality problem.

Let's close by noting how this Fodorian conception of evidence does with the problems above. In the case of illusion, the visual system still outputs a claim about the world, albeit a false claim. So this matches our conclusion that what makes for an illusion is that our evidence is false. In the case of the movie, it is possible that our visual system delivers the output *the hero dies* in the context of seeing the movie on Thursday night.<sup>11</sup> Unless we have a Movie Ending Constancy Predictor module (MECPm), the conclusion about how the movie ends on Friday is not the output of a module. So it is not part of our evidence, as intuition suggests. There's obviously a lot more to be said, but the early evidence is that if Fodor's theories about how modular our minds are are correct, then identifying our evidence with the outputs of our reliable modules captures most of our central intuitions about the nature and role of evidence.

---

<sup>10</sup> He also suggests that this means the observation/theory distinction isn't that important theoretically. But his reason for saying that is just that it no longer corresponds to the infallibly known/inductively risky distinction. That correspondence certainly fails, but there still might be epistemological use for a fallible notion of observation.

<sup>11</sup> Whether this is really possible depends (a) on whether Fodor's right that the outputs of the modules is always shallow and (b) whether *the hero dies* is a shallow content. If the answers are *yes* and *yes*, or *no* and *no*, there's no problem here. I suspect one of those combinations of answers is correct, though I'm ambivalent about which it is. Note I'm also assuming a kind of externalism here - because the visual system delivers that message on Thursday night it delivers a message about Thursday night. I don't assume that it's intrinsic to the message the visual system delivers that it be about Thursday night.

## *References*

Fodor, Jerry (1983) *The Modularity of Mind*. Cambridge, MA: MIT Press.

Fodor, Jerry (2000) *The Mind Doesn't Work That Way*. Cambridge, MA: MIT Press.

Jackson, Frank (1998) *From Metaphysics to Ethics*. Oxford: Oxford University Press.

Lewis, David (1970) "How to Define Theoretical Terms" *Journal of Philosophy* 67:427-446.

Williamson, Timothy (2000) *Knowledge and Its Limits*. Oxford: Oxford University Press.