

## **The New Evil Demon Problem and the Nature of Evidence**

Thomas Kelly  
Princeton University

### **1. Introduction**

Stewart Cohen's New Evil Demon Problem is an enduring contribution to philosophy, one whose significance for the theory of knowledge transcends its traditional role in debates about reliabilism. In this chapter, I explore its significance for debates about the nature of evidence. More specifically, I want to explore the challenge that the New Evil Demon Problem poses for so-called "externalist" accounts of evidence, the paradigm of which is Timothy Williamson's (2000) account of evidence as knowledge.

To lay my cards on the table: like Williamson, I think that paradigmatic evidence in the sciences and elsewhere includes publicly known facts about the external world, evidence that we would not have in a hypothetical scenario in which we are being radically deceived about the true character of our world. In this sense, I am externalist about evidence. On the other hand, I think that the New Evil Demon Problem provides the best reason for doubting the kind of picture that I favor, and I am dissatisfied with the kinds of things that externalists tend to say about it. In short: I think that externalists about evidence need a better story about the New Evil Demon Problem.<sup>1</sup>

---

<sup>1</sup> To be fair, no attempt is made here to critically survey other externalist stories. For a sophisticated recent approach, see Lord (forthcoming).

Ultimately, I'll suggest that the search for an improved story leads to a kind of *inclusive pluralism* about evidence. The pluralist account that I favor retains the idea that our evidence includes our knowledge but also incorporates elements of more traditional accounts of evidence as well. In a broad sense of 'evidence' of particular interest to philosophers, evidence is that which justifies belief, or that which tends to justify belief.<sup>2</sup> Given the radical diversity of topics about which we can have justified beliefs, and the extremely wide range of circumstances in which it's possible for us to hold such beliefs, we should be open to the possibility that even radically different things can play the justifying role of evidence.

## 2. Evidence as Knowledge

Williamson's account of evidence as knowledge (1997, 2000) is perhaps the most prominent and influential account of evidence in contemporary philosophy. In my view, its prominence is well-deserved, and its influence has been salutary. Indeed, if forced to choose among theories of evidence that are currently on offer, I choose Williamson's: by my lights, it's all things considered the best account of evidence that philosophers have devised so far. I cannot possibly hope to vindicate that sweeping comparative judgment here. In lieu of attempting to do so, let me note a few of its features that I believe speak strongly in its favor, respects in which it's preferable to at least many of its historically and currently prominent rivals.

(1) First, as Williamson emphasizes, the view that evidence consists of known facts fits well with much of our thought and talk about evidence outside of philosophy—in the sciences, in legal contexts, and in everyday life. For example, as evidence of climate change, climate scientists cite the known facts that the planet's average surface temperature is rising, that the

---

<sup>2</sup> See Kelly 2006/2018, esp. §1 for reflections on the theme.

oceans are growing warmer, that the Greenland and Antarctic ice sheets have decreased in mass, that glaciers are in retreat around the world, and that the amount of spring snow cover in the Northern Hemisphere has decreased over the past five decades, among others.<sup>3</sup> Similarly, in summarizing the evidence that the defendant committed some crime, a prosecuting attorney might remind the jury of the fact the defendant's DNA was found at the crime scene, or the fact that his fingerprints were discovered on the murder weapon. If I'm wondering whether my wife is home, I might treat the fact that her car is parked in the driveway as evidence that she is. In such cases and countless others, what is cited or treated as evidence are known facts, or putatively known facts. In contrast, many traditional philosophical accounts of evidence do not align nearly as well with the way in which people think and talk about evidence outside of the philosophy seminar room.<sup>4</sup>

(2) Moreover, notice that if we take our practices outside of philosophy at anything like face value, those practices tell strongly against the idea, common to much traditional and current philosophical thinking about evidence, that what counts as evidence is restricted to what can be observed or perceived, or more generally, that something is not evidence if it's arrived at via inference.<sup>5</sup> When a prosecutor cites the fact that *the defendant's DNA was on the murder*

---

<sup>3</sup> <https://climate.nasa.gov/evidence/>

<sup>4</sup> On this disconnect, see my (2008).

<sup>5</sup> Some examples: According to Neta (2018), your evidence is the set of facts that you are in a position to know non-inferentially. According to Goldman (2008:86), your evidence consists of those propositions that you are non-inferentially, propositionally justified in believing. According to Maher (1996), something is evidence if and only if it is "known directly by experience" (158). A major theme of Comesaña (2020) is that "inferentially acquired beliefs are not part of our evidence" (188). Carnap (1950) describes evidence as observational knowledge. According to Quine and Ullian (1970): "Such actual evidence as any one of us does have is in the end the direct evidence of the senses. Likewise such evidence as there is and ever was, collectively, for the whole overwhelming edifice of science, has consisted only in the direct evidence of many peoples' senses" (21-22).

*weapon* as evidence, the fact that he cites is an empirical discovery, but it's not something that anyone observes (nor is it the content of an 'apparent perception' or some nonfactive surrogate). Rather, that a given person's DNA is present at a given location or on a particular object is something that is discovered as the result of an inferential process involving biological and chemical analysis. Nevertheless, it is an absolute paradigm of forensic evidence.

Similarly, the fact that the earth's oceans have grown warmer over some specific interval of time is an empirically discovered fact, but it is not something that one observes or perceives, in the sense of "observation" or "perception" that are at issue when philosophers discuss those notions. Rather, the fact that the earth's oceans have grown warmer over some interval of time is discovered via an inferential process, one that involves comparing temperatures at earlier times with temperatures at later times. Nevertheless, that *the earth's oceans have grown warmer over time* is a paradigm of scientific evidence, and a good philosophical theory of evidence shouldn't suggest or tell us otherwise. This liberalism fits well with Williamson's account, on which *anything* that is known is evidence for the person or group who knows it, irrespective of its content, and even if it's arrived at via an inferential process.

(3) The view that evidence is knowledge also vindicates another important feature of paradigmatic evidence: its potential publicity. If the attentive members of a jury all know that the defendant's fingerprints were found on the murder weapon, then they literally share a piece of evidence that suggests that the defendant is guilty, and that shared evidence can play a role in their collective deliberations. Similarly, if it's known among climate scientists that the earth's oceans are growing warmer over time, then the community of climate scientists can assess the significance of this fact with respect to the hypothesis of global warming. By contrast, on many traditional accounts of evidence, according to which a believer's evidence consists of his private

experiences or the like, there is no possibility of genuinely shared evidence. At best, your evidence might be similar to or resemble the evidence possessed by another person (as in a case in which your private experiences resemble his private experiences), even though there is no actual overlap between his evidence and your evidence. But the possibility of shared or common evidence has often been thought to be important for objective inquiry.<sup>6</sup> The view that  $E=K$  vindicates this possibility, while many rival accounts fail to do so.

So much for my brief survey of some of the notable virtues of  $E=K$ . What about potential weaknesses?

By my lights, the most serious challenge to  $E=K$ , and to externalist theories of evidence more generally, is the challenge of coming up with a satisfying story about the New Evil Demon Problem.<sup>7</sup> In normal circumstances, you have a great deal of knowledge about the external world. For example, perhaps you know, among much else, that the planet's average surface temperature is rising, that the oceans are growing warmer, that glaciers around the world are in retreat, and so on. In the demon world, you know none of these things, because none of them is true. More generally, in the demon world, you know almost nothing about the external world. On the view that all and only knowledge is evidence, you not only have more evidence in normal circumstances than in the demon world, you have *vastly* more evidence. Suppose that in normal circumstances, you believe that global warming is occurring on the basis of your evidence. In these circumstances, your belief is rational. In the demon world, you also believe that global warming is occurring, because you're misled by how things appear. On the assumption that only knowledge is evidence, it seems that your belief will count as irrational in the demon world, for

---

<sup>6</sup> For an overview of this theme, with further references, see my (2006/2018), especially §4.

<sup>7</sup> Here I concur with Littlejohn's (forthcoming) assessment.

you seem to lack any empirical evidence for what is clearly an empirical hypothesis. And it seems that the same holds, *mutatis mutandis*, for countless other beliefs that you would hold about the external world. The Williamsonian view that all and only knowledge is evidence seems to imply that the beliefs of even a maximally scrupulous and conscientious believer are hopelessly irrational in the demon world. But this apparent consequence contradicts the intuitive thought that sufficiently scrupulous and conscientious believers might believe rationally in the demon world, even though their rational beliefs would be false.

What should a proponent of E=K say about the New Evil Demon Problem?

Consider what we can call:

**COHEN'S VERDICT:** In the evil demon scenario, those who are taken in by the demon's deceptions are rational in believing as they do, and their beliefs are rational beliefs.

An initial choice point for a proponent of E=K is whether to deny Cohen's Verdict or to try to accommodate it within the E=K framework. Suppose that the proponent of E=K opts to deny Cohen's Verdict. In that case, they will go on to offer some kind of error-theoretic explanation of its plausibility, or for the relatively widespread inclination to judge that it's true. For example, they might suggest that, insofar as we're tempted to endorse Cohen's Verdict, perhaps what we're actually picking up on is the fact that the thinker's false beliefs are *excused*, or the fact that those beliefs are *the result of a good cognitive disposition operating in unfavorable circumstances*, and this tempts us to mistakenly conclude that the beliefs are rational even though they are not.<sup>8</sup>

---

<sup>8</sup> For views of this kind, see especially Williamson (forthcoming) and Littlejohn (forthcoming), although both put the point in terms of 'justified belief' rather than 'rational belief.' For a forceful critique of the appeal to excuses, see Comesaña (2020), chapter 5.

On some radical views, the falsity of Cohen's Verdict follows from a more general picture on which *only true beliefs can be fully rational*.<sup>9</sup> Like Cohen and Comesaña (2013, forthcoming, Comesaña 2020), I believe that there are compelling reasons to reject this view.<sup>10</sup> Less radically, a theorist might deny Cohen's Verdict while allowing that some rational beliefs are false: perhaps believers in more normal circumstances sometimes have beliefs that are rational and false, although believers in the demon scenario aren't in a position to have such beliefs. Although I take this more moderate view to be more plausible, I believe that we should accept Cohen's Verdict. I'll offer some reasons for accepting it in the next section. In the rest of this section, I want to consider an alternative response to the New Evil Demon Problem that a proponent of E=K might offer, one that attempts to accommodate Cohen's Verdict within the E=K framework.

How might a proponent of E=K try to accommodate Cohen's Verdict? In *Knowledge and Its Limits*, Williamson appeals to the idea that even in the bad case, in which the appearances misleadingly suggest that p is true and one falsely believes that p, one still has genuine evidence for p. For even though one does not know the false proposition p, one still knows the weaker proposition that *it appears as though p*.<sup>11</sup>

---

<sup>9</sup> See Williamson (2013). Compare Williamson (forthcoming) and Littlejohn (forthcoming), which argue that only true beliefs can be justified.

<sup>10</sup> Although I think that the case has already been prosecuted well by others, I attempt to pile on a bit in §3 below.

<sup>11</sup> "In unfavorable circumstances, one fails to gain perceptual knowledge, perhaps because things are not the way they appear to be. One does not know that things are that way, and E=K excludes the proposition that they are as evidence. Nevertheless, one still has perceptual evidence, even if the propositions it supports are false...If perceptual evidence in the case of illusions consists of true propositions, what are they?"

"The obvious answer is: the proposition that things appear to be that way" (2000:198).

Applying this idea to the New Evil Demon Problem, a proponent of E=K might offer speeches like the following:

Consider the scrupulous believer in the demon world, who believes that global warming is occurring, in a way that intuitively seems rational. (For example, she is not like someone who believes in global warming simply because otherwise people in her social circle will make fun of her for not believing it.) Given that she's in the demon world, propositions such as *the earth's average surface temperature is rising*, or *the oceans are growing warmer*, are not evidence, because none of those propositions are true. However, it doesn't follow that she has no evidence for her belief that global warming is occurring, for she knows or is in a position to know facts such as the following: *it appears as though the earth's average surface temperature is rising*, *it appears as though the oceans are getting warmer*, and so on. And these known facts about the way things appear are genuine evidence for the proposition that global warming is occurring, because the probability that global warming is occurring conditional on a fact such as *it appears as though the earth's average surface temperature is increasing* (etc.) is higher than the unconditional probability that global warming is occurring. So given the believer's knowledge of these appearance facts, her belief that global warming is occurring turns out to be rational after all. And the same is true, *mutatis mutandis*, for countless other beliefs that she holds about the external world.

Although believers in the demon world have little or no knowledge about how the external world actually *is*, they have lots of knowledge about how the external world *appears to be*. And this body of appearance knowledge will support many (false) propositions about what the external world is like. So E=K is perfectly consistent with the idea that the scrupulous believer in the evil demon world is justified in many of her beliefs about the external world.

In my view, this is the best response to the New Evil Demon Problem available to the proponent of E=K. Does it work?

Here is an objection. The response depends crucially on the claim that a believer in the demon world will have a great deal of propositional knowledge about how things appear. However, this is a questionable assumption. Notably, some philosophers hold that, although we can of course form beliefs about how things appear when prompted to do so, this is relatively rare. Rather, in typical cases of belief formation, we simply form beliefs about the way the

world is, without bothering to form beliefs about how it appears to be. Here, for example, are John Pollock and Joseph Cruz:

We rarely have any beliefs at all about how things appear to us. In perception, the beliefs we form are almost invariably beliefs about the objective properties of physical objects—not about how things appear to us (1999:61).

If this is correct, and if knowledge requires belief, then, contrary to the line of thought on offer, we actually have relatively little knowledge of how things appear to us. And the same is true of someone in the demon world. So the proponent of E=K is not in a position to accommodate Cohen's Verdict after all.<sup>12</sup>

This objection isn't decisive. In particular, there is room to push back on the idea that we are typically ignorant or lack knowledge of the way things appear.<sup>13</sup> But at a minimum, a proponent of E=K who endorses this line of thought seems committed to a strong empirical claim, that

---

<sup>12</sup>I first floated this objection in Kelly (2006). Williamson discusses the objection in some detail in his (2023:34-36). There, he moves away from the view suggested in *Knowledge and Its Limits*, according to which the deceived thinker in the bad case has evidence that justifies her false beliefs, albeit weaker evidence than she has in the good case. According to the more recent view, the deceived thinker's beliefs are *not* supported by her evidence; rather, what we are responding to, in judging those beliefs favorably, is our awareness of the fact that they are the result of *good cognitive dispositions* (albeit dispositions that are operating in unfavorable circumstances). For reasons that I explain in the next section, I think that it's a mistake to claim that the deceived thinker's beliefs are not supported by her evidence. In short, and especially because I don't claim that the objection presented in this section is decisive, I believe that the earlier view suggested in *Knowledge and Its Limits* is the better view for a proponent of E=K to pursue, and that the more recent view suggested by Williamson (which abandons the attempt to accommodate Cohen's Verdict) is a step in the wrong direction. For further discussion, see below.

<sup>13</sup> As noted by Williamson (2023:35).

human beings are more or less constantly forming and maintaining beliefs about how things appear, at least if knowledge requires belief.<sup>14</sup>

On the assumption that it's an empirical question how often human beings form the relevant kind of appearance beliefs, we can abstract away from that empirical question by considering a being who resembles us as closely as possible but who is *stipulated* not to form such beliefs, unless prompted to do so. (That is, we can consider a being who is stipulated to be exactly as Pollock and Cruz take us to be.) We can then compare two believers, each of whom is trapped in the demon world. The two believers hold all of the same beliefs about how their world is, almost all of which are false. The difference between the two is this. The first believer doesn't bother to form any beliefs about how things appear to him, although his beliefs about how things are are *causally sensitive* to the way they appear, in intuitively rational ways. For example, when it appears as though there is a red wall immediately ahead, he forms the false belief that there is a red wall ahead, and changes in the way things appear inspire changes in his beliefs about how things are, in ways that are intuitively rational. By contrast, the second believer regularly forms beliefs not only about how things are but also about how they appear: for example, when there appears to be a red wall immediately ahead, she forms both the false belief that there is a red wall ahead and the weaker, true belief that *it appears as though there is a red wall ahead*. In short: Although both believers lack knowledge of how things are, the second believer, but not the first, has a great deal of knowledge about how things appear.

A proponent of E=K who endorses the response under consideration here seems committed to saying that while the second believer might very well be rational in holding many of her

---

<sup>14</sup> Thus, Conee and Feldman object to Williamson on the grounds that "The hypothesis that all who have deceptive perceptual evidence have some suitable doxastic attitude toward all these appearance propositions is a highly doubtful empirical claim" (2008:104).

beliefs about the external world, the first believer is not; for the second believer has evidence that tends to justify her empirical beliefs, while the first believer has no such evidence. But that seems like the wrong result. Suppose that their respective beliefs about how things are are *equally* causally sensitive to the way things appear to be. Given that the first individual's beliefs about how things are reflect the way things appear just as perfectly as the second individual's beliefs, the mere fact that he doesn't bother to form beliefs about how things appear doesn't seem to prevent his beliefs about how things are from enjoying whatever epistemic status is enjoyed by the second believer's beliefs about how things are.

Again, this is not decisive.<sup>15</sup> Still, the considerations rehearsed here seem to me to do enough to explore whether a better externalist story about the New Evil Demon Problem might be available.

### 3. The Evil Demon Scenario as a Case of Misleading Evidence

When philosophers disagree about evidence, their disagreements can seem extreme even compared to the usual standards of philosophy. In ethics, consequentialists and deontologists disagree sharply at the level of theory, but when they apply their theories to cases, they generally end up counting many of the same actions as morally wrong, and many of the same actions as morally right. Contrast the disagreement between Williamson, according to whom your evidence consists of all and only the propositions that you know, and the late David Lewis, whose considered view was that your evidence consists of your perceptual experiences and

---

<sup>15</sup> At this point, I think that the best move for the proponent of E=K is to maintain that, when a thinker's beliefs about how things are are sufficiently causally sensitive to (and therefore, informed by) the way they appear, she ipso facto *does* count as knowing that that's how they appear, even in the absence of beliefs about the appearances. That is, they should deny that knowledge entails belief.

apparent memories (1996:424). No perceptual experience or apparent memory is itself a piece of propositional knowledge. Thus, if Williamson's theory is correct, *nothing* that Lewis' theory counts as evidence is actually evidence, and vice versa. It's as if the consequentialist and the deontologist offered us theories of moral wrongness on which none of the same actions are morally wrong.

Given the vast differences in the way philosophers think about evidence, is there any common ground at all?

Here is a claim that I think would and should be conceded by everyone, regardless of their other theoretical commitments:

(1) Evidence can be misleading. ✓

In what follows we will understand misleading evidence in the usual way, according to which

E is **misleading evidence** with respect to hypothesis H if and only if:

- (i) E is genuine evidence for H, i.e., E satisfies the requirements for being evidence for H, *whatever* those requirements are; and
- (ii) H is false.

Here is a paradigm:

**AN IMPROBABLE SEQUENCE:** You watch closely as a coin is flipped fifty times in a row, in an apparently fair way. Each time, the coin lands heads. In fact, the coin is fair. You thus have misleading evidence that *the coin is biased in favor of heads*.

How strong can misleading evidence be? Given that misleading evidence is by definition evidence for what is false, evidence that's misleading can never be *conclusive* evidence.

However, it seems that misleading evidence can be arbitrarily strong, short of conclusiveness. If you observe a fair coin land heads fifty times in a row, your evidence for the false proposition

that the coin is biased is very strong; if you observed it land heads a hundred or a thousand or a million times in a row, your misleading evidence would be even stronger.

Given that people can rationally believe things on the basis of strong evidence that falls short of conclusiveness, we should accept that believers can believe rationally on the basis of misleading evidence when that evidence is strong enough. When you observe the coin land heads enough times in a row, it's rational for you to believe that the coin is biased, even though that proposition is false. It's not simply that your evidence makes it rational for you to invest a *high degree of credence* in the proposition that the coin is biased, or to believe the weaker, qualified proposition that it's *likely* that the coin is biased—although those things are certainly true as well. Rather, in whatever sense we have all-or-nothing beliefs, it will be rational for you to believe that the coin is biased, once you've observed enough flips. After all, at any given point in the process, we can ask what all-or-nothing attitude(s) towards the unqualified proposition it would be rational for you to take. It's not as though you would still be rationally obligated to suspend judgment about whether the coin is biased, after seeing it lands heads a million times in a row. Hence, *pace* the doubts of Williamson and others, we should accept:

(2) False beliefs can be fully rational or justified. ✓

Consider next the phenomenon of deception through the intentional provision of misleading evidence. An obvious example is lying testimony that is credible. Here is another paradigm:

**SUCCESSFULLY FRAMING AN INNOCENT PERSON.** Although the actual murderer knows that the butler is completely innocent of the crime, he wants to convince others—particularly those in positions of power and authority—that the butler did it. He pursues this goal by producing misleading evidence that incriminates the butler. For example, he arranges for the butler's fingerprints to be found on the murder weapon, for the butler's DNA to be discovered at the scene of the crime, and for the search history on the butler's computer to show multiple recent visits to the webpage "How can I dispose

of a dead body without any trace?”<sup>16</sup> The incriminating evidence is so strong that even those who are initially strongly disposed to believe that the butler is innocent are ultimately persuaded by it.

In this case, the murderer’s project is to get people to believe that the butler did it by making it rational for them to believe that the butler did it and irrational for them not to believe this. His project involves *exploiting* the rationality of those he intends to dupe and relies on the fact that he can count on them to form certain beliefs (including false beliefs) when those beliefs are overwhelmingly likely to be true given their evidence. To the extent that the framer is successful in producing compelling evidence of the butler’s guilt, rational people will respond to it by coming to believe that the butler is guilty. When that happens, both the believer and their beliefs count as rational.

The plan of misleading people by manufacturing and planting misleading evidence presupposes that they are the kind of people whose beliefs will be rationally influenced by their evidence, not only in general or in some dispositional sense, but on *this* particular occasion—with respect to *these* particular pieces of evidence, and with respect to *this* particular question. Suppose instead that those the framer seeks to mislead are the sorts of people who arrive at their beliefs about which suspects committed which crimes irrationally, perhaps on the basis of racial and ethnic animus. In that case, the strategy of planting misleading evidence is not a good one; he might be better off calling attention to the butler’s ethnicity or race.

What is the relationship between the framing case and the evil demon scenario? I submit that the best way to think of the evil demon is as an agent who is engaged in a global version of the same project that the framer pursues in the framing case. The framer’s goal is to get people to

---

<sup>16</sup> <https://www.quora.com/How-can-I-dispose-of-a-dead-body-without-any-trace-Its-for-a-story-that-Im-writing>

believe something false about one particular issue: the identity of the person who committed the crime. The framer pursues this goal by arranging things so that *what people have to go on* in thinking about that issue strongly suggests what is false. He exploits their rationality by feeding them intentionally misleading evidence. The demon's goal is to get people to believe something false about *many* issues. The demon pursues that goal by arranging things so that *what people have to go on* in arriving at their beliefs about these many issues strongly suggests what is false. But what they have to go on is their evidence. In short, the demon, like the framer, is in the business of exploiting rational people by feeding them misleading evidence: the essential mechanism is the same in both cases.

If I'm correct in holding that the best way to think about the demon scenario is as a globalized version of the framing case, and the two are alike with respect to their essential epistemological features, then certain consequences immediately follow. First, we should endorse Cohen's Verdict:

- (3) Cohen's Verdict: In the demon scenario, those who are taken in by the demon's deceptions are rational in believing as they do, and their beliefs are rational beliefs. ✓

For in the framing case, those who come to believe that the butler is guilty are rational—not merely in some dispositional sense, but on this particular occasion—as is their belief that the butler is guilty. Second, we should also accept the conclusion that

- (4) The demon's victims have significant evidence for their beliefs, evidence that is strong enough to rationalize or justify those beliefs. ✓

for this is the mechanism of deception in the framing case.

The conclusion that we would have significant evidence in the demon scenario puts pressure on the thought that *only* knowledge is evidence, for we would have little knowledge in the demon scenario. However, the conclusion that we would have significant evidence in the demon scenario is perfectly consistent with externalism about evidence, as well as with the more specific externalist thesis that *everything* that we know is part of our evidence. For the conclusion that we would still have significant evidence in the demon scenario is perfectly consistent with the core externalist thought that we have significantly *more* evidence in normal circumstances. In §2, I praised Williamson's account for its liberalism or inclusivity: notably, for example, it allows for a piece of non-inferential knowledge such as *the oceans are growing warmer* to count as evidence. It's worth asking whether a satisfying externalist response to the New Evil Demon Problem can be had by liberalizing things further. Even if paradigmatic evidence consists of facts that we know, perhaps our evidence is not exhausted by our knowledge.

What else might count as evidence? In a broad sense of evidence particularly of interest to philosophers, evidence is that which justifies belief, or that which tends to justify belief. Given that understanding, the question then becomes: what else might justify belief, besides propositional knowledge? Given the nature of the New Evil Demon Problem, a natural place to look is to the kinds of non-factive mental states and events that internalists about evidence take to exhaust our evidence.<sup>17</sup>

---

<sup>17</sup> I first proposed a more inclusive account of evidence along these lines in Kelly (2006):

According to the phenomenal conception of evidence, *only* one's experiences can serve as evidence. According to Williamson's conception of evidence as knowledge, one's experiences are excluded from counting as evidence—at best, one's evidence includes whatever propositions about one's experiences that one knows. Even if one abandons the phenomenal conception of evidence, however, one might hold on to the idea that one's evidence *includes* one's experiences, inasmuch as one's experiences can and often do make some difference to

Indeed, in principle, an externalist about evidence can allow that the believer in the demon scenario has every bit as much evidence as any internalist thinks. For there is nothing in externalism as such that would prevent the externalist from recognizing as evidence whatever it is that the internalist takes to be evidence. For example, if the best version of internalism takes our evidence to include (non-factive) perceptual experiences, then it's open to the externalist to endorse that idea. What the externalist is committed to denying is that our evidence in normal circumstances is exhausted by what the internalist takes to be evidence. But by the internalist's own lights, the evidence that's available in the demon scenario is sufficient to make it fully rational for the scrupulous believer to believe as she does. So it's open to the externalist to recognize as evidence things that, by the internalist's own lights, are sufficient to return the verdict that the scrupulous believer and their beliefs are fully rational in the bad case.

In short: *an externalist can accommodate Cohen's Verdict so long as the internalist can.*

Is it objectionable that the externalist will go on to claim that the believer in normal circumstances has significantly more evidence than she would in the demon world, even though on current assumptions her evidence in the demon world suffices for the rationality of her ordinary beliefs? No, for the relevant epistemic phenomenon here is one that we find in everyday life, even before we begin to think about evil demon scenarios. The phenomenon is

---

what one is justified in believing, regardless of whether one forms beliefs about those experiences themselves. A view of evidence that is more liberal than either Williamson's or the phenomenal conception might thus take one's evidence to include both one's experiences and one's knowledge, on the grounds that the beliefs of a rational thinker will exhibit direct sensitivity both to what he knows and to the experiences that he undergoes.

Schellenberg (2016a) develops and defends a view of perceptual evidence that allows for both factive and phenomenal evidence.

this: when one is in favorable circumstances, the fact that one's rational beliefs are rational is often overdetermined, in the sense that, even if one had significantly less evidence than one actually does, it would still be rational to believe as one does. For example, plausibly, many of the things that it's rational for you to believe about your family members, or about your job, are things that it would still be rational for you to believe, even if you only had access to some of the evidence to which you in fact have access.<sup>18</sup>

Of course, even though a suitably liberalized version of externalism can accommodate Cohen's Verdict, it doesn't follow that the externalist ultimately escapes the New Evil Demon Problem. For it might be that the best version of externalism comes to grief with respect to the New Evil Demon Problem, even if externalism as such does not. For example, even if it is consistent with externalism about evidence to allow that (e.g.) one's perceptual experiences are evidence, it is not consistent with all versions of externalism. Notably, for example, it is not consistent with E=K. So the New Evil Demon Problem might be thought to provide a good objection to externalism about evidence in virtue of posing a good objection to the best version(s) of externalism. Moreover, if an externalist allows some non-factive mental states to count as evidence simply to parry the New Evil Demon Problem, this might look disreputably *ad hoc*, and out of the spirit of her account of evidence.

However, I think that we have good reasons to prefer a more liberal version of externalism independently of anything having to do with the New Evil Demon Problem. Indeed, I think that we have reasons to prefer a more liberal version of externalism even before we consider the more

---

<sup>18</sup>A related objection is that, once the externalist allows that the evidence available to us in the demon scenario (or more generally, in 'the bad case') is sufficient for rational believing (as the internalist thinks), externalism about evidence has become *pointless*: evidence that satisfies internalist constraints can do all of the work that needs doing, so there is no reason to recognize anything else. For discussion of this objection, see §4.

mundane cases of apparently rational false beliefs that are often used to put pressure on externalist views of evidence and reasons in the literature. In the next section, I look in some detail at a mundane case of knowledge that I think gives us good reason to think that our evidence is not exhausted by our propositional knowledge.

#### 4. Externalism Liberalized: Beyond E=K, Towards Inclusive Pluralism about Evidence

Consider the epistemology of headaches. When I have a headache, I typically know that I do, and my belief that I have a headache is rational. Suppose that in such circumstances we ask: what *causes* me to believe that I have a headache? Here the natural thing to cite is the headache itself: the event or state whose onset leads me to take up the belief in the first place, and which I'm trying to bring to an end when I take an aspirin. (Contrast an unusual case in which I don't have a headache but believe that I do because I've been hypnotized to believe this.)

Suppose that we next ask: What *justifies* my belief that I have a headache? Here again an obvious candidate is the headache itself, a certain token event or state that causes the belief.<sup>19</sup> Of course, many philosophers will reject that answer on the grounds that a headache is simply not the right kind of thing to justify a belief, and they will offer philosophical considerations to support that claim. But before attending to some of those considerations, we should take note of just how commonsensical this answer is.<sup>20</sup> Indeed, one consideration that lends this answer

---

<sup>19</sup> Of course, many subtle questions could be asked about the ontology or metaphysics of headaches. I intend to skirt such questions. On the view taken here, in whatever sense it can be true to say "The headache causes my belief that I have a headache" it can also be true to say "The headache justifies my belief that I have a headache," while holding the designation of "the headache" fixed.

<sup>20</sup> Notably, for example, Williamson, who explicitly denies that sensations or experiences can justify beliefs, admits that "on a broad interpretation of 'evidence'...it is natural to say that my evidence that I am getting a cold includes various sensations" (2000:194).

plausibility is the way in which whether it's rational for me to believe that I have a headache seems to closely track the presence or the absence of the headache itself: generally speaking, when I have a headache, it's rational for me to believe that I do, and when I don't have a headache, it would not be rational for me to think that I have one.

The idea that it's the headache itself that justifies the belief is the kind of answer that will be enthusiastically endorsed by proponents of a kind of old-fashioned foundationalism according to which all of our knowledge of the empirical world ultimately rests on a foundation provide by acquaintance with our own private sensations. But we shouldn't despise that answer simply because the relevant kind of old-fashioned foundationalism is false, or because the epistemology of headaches provides a poor model for the epistemology of the sciences, as it surely does.<sup>21</sup>

As noted above, in a broad sense of evidence particularly of interest to philosophers, evidence is that which justifies belief, or that which tends to justify belief. Of course, we should not uncritically assume that every rational belief is made rational by the believer's evidence.<sup>22</sup> For example, my belief that  $1+1=2$  is (I assume) a rational belief, but it's doubtful that the right way to understand its rationality is in terms of my having evidence for its truth; rather, I simply recognize that it's true. In contrast, my rational belief that global warming is occurring is a paradigm of a belief that is rationalized by evidence. Could my belief that I have a headache be relevantly similar to my belief that  $1+1=2$ , something that I simply recognize as true, without its being justified or rationalized by anything beyond itself?

The problem with that thought is that, in the case of my belief that I have a headache, there really does seem to be a mental event or state that is distinct from the belief that justifies me in

---

<sup>21</sup> Although see Kriegel (2024).

<sup>22</sup> For discussion of this point, see Kelly (forthcoming).

holding it. Of course, there are all sorts of questions that a philosopher of mind might ask about how exactly that state should be understood; as suggested above, for our purposes we can pick it out as that event or state that I'm trying to bring to an end when I consume some pain-killing medication. On anyone's view, my being in a state of pain plays a *causal role* in my coming to believe that I have a headache, and the persistence of my believing that I have a headache will in the usual case be causally sensitive to the persistence of the pain itself. The only question is whether the occurrence and persistence of the pain also plays a justificatory role with respect to my belief that I'm in pain, or *merely* a causal one. Intuitively, the case for a justificatory role seems strong. After all, there are numerous events that play a causal but not a justificatory role with respect to my current belief that I have a headache. For example, there is the event of my travelling last night to the bar at which I consumed too many drinks. But the pain that I currently feel seems to contrast with such events precisely in that it, but not they, play a justificatory role with respect to my belief that I currently have a headache. But what justifies belief is evidence. Therefore, the pain that I currently feel is evidence. But the pain that I currently feel is not itself a piece of propositional knowledge. Therefore, not all evidence is propositional knowledge, even if paradigmatic evidence is.

By my lights, a paradigmatic context in which it's natural to speak of evidence is a scientific context, or a legal one, or the context of some criminal investigation, or in cases involving ordinary inductive reasoning in everyday life. In such contexts, the picture of evidence as propositional knowledge holds up at least relatively well (although as the internalist will remind us, there are certainly challenging cases). By contrast, it's difficult to imagine a non-artificial case in which someone wonders about what their evidence is that they have a headache, and this naturally encourages the thought that it's simply a mistake to understand the case as one that

involves evidence. The case for thinking of the feeling of pain as evidence only comes into view when we ask the question, “What, if anything, justifies my belief that I have a headache?” For various reasons, that is a question that we are unlikely to ask outside of the epistemology seminar room. But it is not a bad question for that reason.

Are there reasons to doubt the seemingly plausible thought that it’s the feeling of pain that justifies me in believing that I have a headache, and therefore deserves to be counted as evidence?

On some views, an experience of pain has no propositional content, and some philosophers hold that only something with propositional content can be evidence.<sup>23</sup> For example, in *Knowledge and Its Limits*, Williamson argues that evidence must consist of propositions, on the grounds that “Only propositions which we grasp serve the central evidential functions of inference to the best explanation, probabilistic confirmation, and the ruling out of hypotheses” (2000:196).

Suppose that we concede the quoted claim, at least for purposes of argument.<sup>24</sup> Consider the case of probabilistic confirmation. The logical positivist tradition, which inaugurated the

---

<sup>23</sup> The line of thought goes back at least to Wilfrid Sellars (1956) and plays a significant role in the history of 20<sup>th</sup> century epistemology. To be clear, the approach to the New Evil Demon Problem suggested in this paper is compatible with the view that all evidence is propositional. For example, it’s compatible with the view that our evidence in the demon world consists of perceptual experiences with propositional content. However, because I believe that broadly Sellarsian considerations cannot bear the kind of argumentative weight that’s often placed on them in this kind of context, the issue is worth considering.

<sup>24</sup> For the record, I doubt that the quoted claim is true. Consider the idea that only propositions can serve the “central evidential function” of inference to the best explanation. This idea depends in turn on the following argument: “...evidence is the kind of thing that hypotheses explain. But the kind of thing that hypotheses explain is propositional. Therefore, evidence is propositional” (p.195). However, we often seek to explain (the occurrence of) token events. Indeed, some of the best philosophical work on explanation simply takes for granted that token events are among the paradigms of things that we seek to explain by offering hypotheses. (Consider, for example, the opening sentence of David Lewis’ seminal paper “Causal Explanation”: “Any particular event that

systematic study of confirmation within philosophy, often took confirmation to be a relation between *sentences*: so-called “protocol sentences” or “observation reports” on the one hand, and the hypotheses that they confirm on the other, hypotheses that were similarly understood to be sentences in some language. On this picture, it’s the job of the protocol sentences to *report* the evidence. This naturally suggests the possibility of distinguishing between the evidence itself and the statements that report that evidence, while taking the statements to stand in the relevant probabilistic relationships to the hypotheses. Similarly, even if we dispense with sentences, and assume with Williamson that the kind of probabilistic relationships that are of primary relevance for confirmation theory hold only between *propositions*, this is perfectly compatible with the idea that it’s a certain token event or experience which justifies (and therefore, is evidence for) my belief that I have a headache. For we can consider the proposition that *reports* the experience and ask about the probability of various hypotheses on that proposition. In this way, a view on which some experiences are themselves evidence for beliefs can be fully integrated into a probabilistic framework, even on the assumption that the relevant kind of probability holds

---

we might wish to explain stands at the end of a long and complicated history” (1986: 214)). According to the suggestion endorsed in the main text, the evidence for my belief that I have a headache is a certain token event. On this view then, what’s taken to be evidence is *exactly* the type of thing that we often seek to explain when we offer hypotheses. So the requirement that evidence must be the type of thing that we seek to explain provides no reason whatsoever to doubt the idea that token events can play the role of evidence; on the contrary, the explanatory requirement fits very well with that idea. Of course, a proponent of Williamson’s argument might suggest that all of our talk of explaining events both in and outside of philosophy should be reinterpreted so that it’s really always propositions that are the proper explananda, but as far as I can see there is no reason to accept that suggestion.

Williamson’s arguments for the conclusion that all evidence consists of propositions are criticized effectively by Neta (2008: 95-98). After criticizing Williamson’s arguments, Neta offers another argument for the same conclusion, an argument that he characterizes as “more compelling, though still far from conclusive” (p.96). Because the view that I ultimately endorse in this paper is consistent with the view that all evidence is propositional, I won’t explore the limitations of Neta’s argument here.

exclusively between propositions, as Williamson thinks. Analogous remarks apply to the two other “central evidential functions” mentioned by Williamson.<sup>25</sup> However, on the suggested picture, it’s not *the proposition that I’m having a headache* that justifies my belief that I have a headache; rather, it’s the feeling of pain that justifies the belief.

*What justifies me in believing that I have a headache is not some piece of information or proposition that I know.*<sup>26</sup>

In contrast, in many other cases, my evidence includes propositions (pieces of information) to which I stand in some cognitive relationship. Again, the most plausible view is that my evidence for global warming includes information such as the proposition that *the earth’s glaciers are in retreat*, or the proposition *that the community of climate scientists generally agree that global warming is occurring*. In fact, I think it’s extremely implausible that my evidence for the claim that global warming is occurring is exhausted by my current experiences.<sup>27</sup> On the other hand,

---

<sup>25</sup> Compare Neta’s critique of Williamson’s argument that evidence must consist of propositions because only propositions can rule out hypotheses in virtue of being inconsistent with them: “Let’s grant that only propositions can be inconsistent in the relevant sense: Why should we allow, though, that there is an inconsistency between hypothesis and evidence itself, rather than an inconsistency between hypothesis and one or another statement of the evidence? Why not say, instead, that the hypothesis that I feel just fine is ruled out by my sore throat, and so the hypothesis is inconsistent with a particular statement about the evidence, namely, that my throat is sore?” (2008: 97-98).

<sup>26</sup> Contrast: if you know that I had more than  $n$  drinks last night, and that, whenever I have more than  $n$  drinks, I have a headache the following day, then what justifies you in thinking that I have a headache is in fact some information or proposition that you know. And of course, I might also know the same information. So we can construct cases in which what justifies me in believing that I currently have a headache is my possessing certain information. But it would be a mistake to assimilate the usual case, in which I simply believe that I have a headache in response to the headache, to the same model.

<sup>27</sup> Perhaps the best piece of evidence that I have that global warming is occurring is my knowledge that the climate scientists generally agree that it is. But my knowledge of this fact is not a matter of my current experiences; rather, I know it on the basis of preservative memory. Compare Williamson: “The evidence for the proposition that the sun is larger than the earth is not just my present experiences or degrees of belief” (2000:193).

there is nothing similarly implausible about the idea that what justifies me in believing that I have a headache is my experience of the pain. Nor should it be surprising that my best evidence that global warming is occurring is a very different type of thing than that which justifies my belief that I currently have a headache.

If my evidence that global warming is occurring includes the facts that the glaciers are in retreat, or that climate scientists generally agree that it's occurring, then your evidence might include these facts as well: the evidence that I have for the hypothesis might quite literally also be a part of your evidence. In contrast, if what justifies me in believing that I have a headache right now is the pain that I currently feel, then (science fiction cases aside) there is no possibility that you literally share my evidence. Although you might know that I have a headache, what justifies you in thinking this will be quite different from what justifies me. Inasmuch as my belief is justified by evidence, that evidence is private evidence. But this seems exactly right. Prior to any particular theory, the epistemology of my knowing that I have a headache seems very different from how you know this. It is a virtue of a hybrid account, on which token sensations as well as known propositions can count as evidence, that it readily allows for both public evidence and private evidence. By contrast, some traditional accounts (e.g., accounts on which our evidence ultimately consists in private mental states or events) do not do justice to the public character of much of our evidence, while accounts that see all evidence as propositions that can be grasped by more than one person run the opposite risk, of making too little private.<sup>28</sup>

In the usual case, one's beliefs about whether one has a headache will be causally sensitive to the presence and absence of the headache. This causal sensitivity is not mediated by other

---

<sup>28</sup> Compare some of the awkward epistemological consequences of behaviorism (as in the old joke: "I see that you're doing well! Now, how am I?"), which was itself in part a historical reaction to views that had an implausibly private conception of evidence in psychology.

beliefs: typically, coming to believe that one has a headache is a more or less immediate response to the painful state itself. In the same way, a person's believing that things *are* a certain way (in the external world) might be causally sensitive to the way things appear, even if one has no beliefs to the effect that things appear that way. (If things ceased to appear in that way, then one would no longer believe that things are that way, and so on.) When things visually appear to one in a certain way, one is in a certain state, a state that one might be in even if one lacked a belief that things are as they appear. (Recall here the familiar example: a person who is already familiar with 'the bent spoon illusion' will not hold the false belief that the partially submerged spoon is bent, even though it still appears to him as though the spoon is bent, in a way that it does not appear to the blind person.)

If we allow that a feeling of excruciating pain can justify (and therefore, count as evidence for) my belief that I have a headache, then I think we should also allow that non-doxastic perceptual states can justify (and therefore, count as evidence for) the belief that things are the way those states represent them as being. It follows that a thinker in the demon world does have evidence for his false beliefs about how the world is, even if he lacks beliefs (and therefore, knowledge) about how things appear.

Of course, the externalist who seeks to accommodate Cohen's Verdict is not committed to this or any other specific story about what our evidence is in the demon scenario. As emphasized above, the externalist is free to coopt the best internalist story (whatever it is) about what justifies our ordinary beliefs in the demon scenario, and in this way accommodate Cohen's Verdict.

At this point, one might object that externalism about evidence has become *pointless*. What is the point of insisting that our evidence in normal circumstances includes publicly known facts about the external world (or more generally, evidence that we wouldn't have in the demon

scenario), if it would still be rational for us to hold our ordinary beliefs in the absence of this evidence, on the basis of evidence that's available to us in the demon scenario?

Among other things, this Pointlessness Objection assumes that the only epistemic difference that having stronger evidence would make is to what it is rational or justified to believe on its basis. But there is no reason to accept that assumption. For example, it's plausible that having stronger evidence (e.g., evidence that includes known facts about the external world) puts us in a position to *know* things that we would not be in a position to know on the basis of weaker (less inclusive) evidence, even if the valence of the evidence is the same in both cases with respect to the propositions at issue.

To illustrate: suppose that some scientists who are thoroughly acquainted with the best evidence that global warming is occurring not only rationally believe that it's occurring, but *know* that it is. In explaining how they are in a position to know that such a far-flung hypothesis is true, it's natural to cite the fact that their evidence includes various facts that have been vetted by and are now publicly known within the scientific community, including facts about the retreat of the glaciers, the decrease in mass in the Greenland and Antarctic ice sheets, and so on. Even if a scientist genuinely knows that global warming is occurring in this way, it hardly follows (and indeed, does not seem particularly plausible that) that she would still be in a position to know this, if her evidence did not include any such facts, but was limited instead to her non-factive mental states, or the weaker kinds of evidence that would also be available to her in the demon scenario.

Of course, in the evil demon scenario, a scientist is not in a position to know that global warming is occurring because it's *false* that global warming is occurring. The relevant comparison is rather between (1) a possible world in which it's true that global warming is

occurring and the believer's evidence includes publicly known and scientifically vetted facts such as those mentioned above, and (2) a possible world in which it's true that global warming is occurring but for some reason the believer's evidence includes no such facts, but only evidence that would also be available to her in the demon scenario. The point is that, even if the believer's evidence is strong enough to know in world (1), it does not follow (and indeed, does not seem especially plausible that), her evidence is also strong enough for her to know in world (2).

More generally, the following claim is extremely plausible: the capacity of the empirical sciences to sometimes deliver full-fledged theoretical knowledge depends on the existence of observational knowledge of the external world. For example, a world in which there is no experimental knowledge is not a world in which we would know the theoretical proposition that *there are negatively-charged subatomic particles*. Given that our observational evidence is sometimes strong enough to support full-fledged theoretical knowledge, it seems that such observational evidence should include full-fledged knowledge as well: nothing weaker will do. The fact that we think that the empirical sciences sometimes do deliver full-fledged theoretical knowledge is thus another reason for thinking that our evidence sometimes consists of knowledge. However, in a demon world in which everything merely appears as though there are negatively-charged sub-atomic particles, it would still be rational for us to believe this theoretical proposition, even though in that world we would have no observational or experimental knowledge.<sup>29</sup>

---

<sup>29</sup> As noted above, Schellenberg (2016a) develops and defends a view of perceptual evidence on which one has both factive and phenomenal evidence in the good case but only phenomenal evidence in the bad case. For Schellenberg, this difference primarily manifests itself in terms of which propositions one is justified in believing in the good case as opposed to the bad case. This opens the account up to the charge that it underestimates the epistemic credentials of the believer's beliefs in the bad case. (See McGrath 2016 for the charge and Schellenberg 2016b for a reply). In contrast, on the view taken here, the fact that our evidence sometimes includes full-fledged knowledge is primarily relevant to what we're in a position to know on its basis, as opposed to

Finally, one might wonder about how the inclusive view of evidence sketched here deals with *doxastic* justification or rationality.<sup>30</sup> Consider the following objection:

You say that the evil demon victim has enough evidence for rational belief, but that that evidence would be insufficient for knowledge even if the beliefs that it supports were true. On the other hand, you say that his counterpart in normal circumstances has knowledge in part because they have *factive* evidence that the victim lacks. But do you think that there is a difference between the victim and the normal subject regarding *doxastic* justification? Do they base their beliefs only on their common evidence, and that is why they are both *doxastically* justified? That would mean that the normal subject does not base his beliefs on the *factive* evidence, and if it's the *factive* evidence that accounts for his knowledge, this would be a problem. Or does the normal subject base his beliefs on the *factive* evidence? That would mean that there is a difference in the basis of the beliefs for the victim and the normal subject. One could say this, but surely that's a cost of the view.

I reply: given a *purely psychological characterization* of the bases of their beliefs (i.e., a characterization that abstracts away from facts about which underlying beliefs are known), there is no difference between the basis on which the normal subject believes and the basis on which the deceived subject believes. But that's consistent with the claim that the normal subject bases his belief on knowledge while the deceived subject does not. Let me explain.

Consider once again the would-be climate scientist in the demon world. There, it appears as though *the planet's average surface temperature has been rising*, and the would-be climate scientist believes this. Suppose that her belief that *climate change is occurring* is appropriately sensitive both to the appearances and to her belief that the planet's average surface temperature has been rising. (If it did not appear that the average surface temperature has been rising, she

---

which propositions we're justified in believing; for genuine evidence that falls short of knowledge will often be enough to license rational or justified belief so long as no countervailing or defeating evidence is present.

<sup>30</sup> Thanks to Juan Comesaña and Matt McGrath for asking about this, and for the objection discussed here.

would not believe that climate change is occurring; and it's also true that if she did not believe that the average surface temperature has been rising, then she would not believe that climate change is occurring.) This appropriate sensitivity is enough for her belief that climate change is occurring to be doxastically justified.<sup>31</sup>

All of this is true for the normal subject as well. It does not follow that the normal subject bases her belief that climate change is occurring on something weaker than knowledge. For in her case, the proposition that the planet's average surface temperature has been rising is something that she knows, and its status as such is what puts her in a position to know the climate change hypothesis on its basis.

## 5. Conclusion

What then should we say about Cohen's New Evil Demon Problem, in terms of what it suggests about the nature of evidence?

Suppose that the Philosophy Oracle informs us that Cohen's original judgment about the demon scenario was right all along: those who are taken in by the demon might be fully rational in believing as they do. Should this revelation lead us to conclude that when the climate scientists tell us that the evidence for global warming includes publicly known facts about the earth's surface temperature they speak falsely, or in a way that must be carefully reconstructed by the epistemologist in order to come out true? No! For what we've learned from the oracle is simply that the true story about evidence in the demon scenario (whatever it is) is consistent with

---

<sup>31</sup>At least, on the assumption that being propositionally justified in believing that *the planet's average surface temperature is rising* is sufficient for being propositionally justified in believing that *climate change is occurring*. If necessary, we could complicate the example by bringing in more propositions that appear to be true in the demon world and that support the climate change hypothesis (e.g., propositions having to do with the apparent retreat of the glaciers, and so on). But that would not alter the essential picture.

the rationality of the deceived, and there is no reason to think that that story should make us take what the climate scientist says at anything other than face value.

Some philosophers treat it as a kind of hard datum for philosophy that a radically deceived thinker in the demon world and his counterpart in normal circumstances are equally well-justified in everything that they believe. The philosopher then infers from this that the counterparts must have the same evidence in the two worlds. As Williamson emphasizes, this is a severe restriction on the kinds of things that can count as evidence, a restriction that immediately rules out many plausible views, such as the view that the evidence for global warming includes the known fact that the earth's glaciers are in retreat.

But it is surely not a piece of hard philosophical data that our evidence never includes publicly known facts about the way the world is. Even if one finds that the most attractive view when one reflects on the demon scenario in isolation, the thought should not survive consideration of a wider range of contexts in which the concept of evidence is central. With Williamson, we should reject the suggestion that we have no more evidence in normal circumstances than we would have in the evil demon scenario.

Here is an idea that seems to me to have a much better claim to being a piece of hard philosophical data: when our deceived counterparts in the demon scenario believe in accordance with how things appear to them, they are not irrational, in anything like the way a person who confidently believes an empirical hypothesis in the absence of any evidence for it is irrational. (Nor is the difference merely that they are excused for their irrational beliefs whereas the typical person who confidently believes an empirical hypothesis in the absence of evidence has no excuse.) When our deceived counterparts respond in the natural way to what they have to go on, the beliefs that result are rational, for those beliefs reflect all that they have to go on, and what

they have to go on is their evidence. By adopting a more liberal conception of evidence, one that does not limit a believer's evidence to their knowledge, we can accept the thought that both the deceived and their beliefs are rational, for they have evidence that supports what they believe, and no evidence that tells against it.

If evidence is what justifies belief, then, given the diversity things that we rationally believe, and the diversity of circumstances in which it's possible to believe rationally, we should be open to the possibility that more than one kind of thing can play the role of evidence. In short: we should be inclusive pluralists about evidence. When Williamson tells us that paradigmatic evidence includes known facts about the external world, we should endorse that suggestion, and embrace the externalism about evidence that comes with it. But if an internalist like Lewis tells us that our evidence includes our perceptual experiences, our externalism should not lead us to balk at that suggestion. For the suggestion is both plausible on its face and also allows us to say what many of us would like to say about the deceived and their beliefs in the demon scenario.

Isn't the kind of inclusive pluralism about evidence embraced here a theoretically inelegant and unlovely hybrid? Don't considerations of theoretical elegance favor the satisfyingly simple and straightforward equation  $E=K$ , or for that matter, a view on which a believer's evidence is ultimately exhausted by their current experiences, or the contents of those current experiences, or by the way that things appear to them? As someone who values theoretical elegance, I understand the complaint. But I am moved by the thought that the kind of thing that justifies me in believing that I have a headache is quite unlike the best evidence that climate scientists have found for global warming; and that when the demand for theoretical elegance is taken to be a

compelling reason to deny this, that is a good sign that the demand is being misapplied, in a way that has begun to distort the phenomena.<sup>32</sup>

## REFERENCES

- Carnap, Rudolf (1950). *Logical Foundations of Probability* (Chicago).
- Cohen, Stewart (1984). “Justification and truth.” *Philosophical Studies* 46(3):279--95.
- Cohen, Stewart & Comesaña, Juan (forthcoming). “Rationality and Truth.” In Julien Dutant & Fabian Dorsch (eds.), *The New Evil Demon* (OUP).
- Comesaña, Juan (2020). *Being Rational and Being Right* (OUP).
- Conee, Earl and Feldman, Richard (2008). “Evidence.” In Quentin Smith (ed). *Epistemology: New Essays* (OUP): 83-104.
- Goldman, Alvin (2008). “Williamson on Knowledge and Evidence,” in Greenough and Pritchard (eds.) *Williamson on Knowledge* (OUP):73-91.
- Kelly, Thomas (2006/2018). “Evidence.” In the *Stanford Encyclopedia of Philosophy*.
- \_\_\_\_\_(2008). “Evidence: Fundamental Concepts and the Phenomenal Conception”, *Philosophy Compass*3(5), (Blackwell): 933-955.
- \_\_\_\_\_(forthcoming). “Varieties of Evidentialism.” In McCain and Stapleford (eds.) *Evidentialism at 40* (Routledge).
- Uriah Kriegel. (2024). Knowledge-by-Acquaintance First. *Philosophy and Phenomenological Research*. 1-20. <https://doi.org/10.1111/phpr.13051>
- Lehrer, Keith & Cohen, Stewart (1983). “Justification, truth, and coherence.” *Synthese* 55(2):191-207.
- Lewis, David (1986). “Causal Explanation.” In his *Philosophical Papers*, vol.2 (OUP):214-240.
- \_\_\_\_\_(1996). “Elusive Knowledge.” *Australasian Journal of Philosophy* 74(4):549 – 567.
- Littlejohn, Clayton (forthcoming). “A Plea for Epistemic Excuses.” In Julien Dutant & Fabian Dorsch (eds.), *The New Evil Demon* (OUP).
- Lord, Errol (forthcoming). “Defeating the Externalist’s Demons.” In Dorsch and Dutant (eds.) *The New Evil Demon* (OUP).

---

<sup>32</sup> Thanks to participants in my fall 2023 graduate seminar at Princeton University for helpful feedback, as well as to colleagues who attended a meeting of Princeton’s faculty-work-in-progress lunchtime series at which an earlier version of some of this material was presented. Special thanks to Lidal Dror, Mark Johnston, and Gideon Rosen for follow-up conversations.

- Maher, Patrick (1996). "Subjective and objective confirmation." *Philosophy of Science* 63 (2):149-174.
- McGrath, Matthew (2016). "Schellenberg on the epistemic force of experience." *Philosophical Studies* 173 (4):897-905.
- Neta, Ram (2008). "What evidence do you have?" *British Journal for the Philosophy of Science* 59 (1):89-119.
- \_\_\_\_\_(2018). "Your Evidence is the Set of Facts that Are Manifest to You." In Mitra (ed.) *The Factive Turn in Epistemology* (CUP).
- Pollock, John and Cruz, Joseph (1999). *Contemporary Theories of Knowledge* (Rowman and Littlefield).
- Quine, W.V. and Ullian, J.S. (1978). *The Web of Belief* (McGraw-Hill).
- Schellenberg, S (2016a). "Phenomenal evidence and factive evidence." *Philos Stud* **173**: 875–896  
<https://doi.org/10.1007/s11098-015-0528-8>
- \_\_\_\_\_(2016b). "Phenomenal evidence and factive evidence defended: replies to McGrath, Pautz, and Neta." *Philosophical Studies* 173(4):929-946.
- Sellars, Wilfried (1956). *Empiricism and the Philosophy of Mind*.
- Williamson, Timothy (1997). "Knowledge as Evidence." *Mind* 106: 717-741.
- \_\_\_\_\_(2000). *Knowledge and Its Limits* (OUP).
- \_\_\_\_\_. (2013). Response to Cohen, Comesaña, Goodman, Nagel, and Weatherson on Gettier Cases in Epistemic Logic. *Inquiry*, 56(1), 77–96. <https://doi.org/10.1080/0020174X.2013.775016>
- \_\_\_\_\_(2023). "E=K, But What about R?" In Lasonen-Aarnio and Littlejohn(eds.) *The Routledge Handbook of the Philosophy of Evidence* (Routledge):30-39.
- \_\_\_\_\_(forthcoming). "Justification, Excuses, and Skeptical Scenarios." In Julien Dutant & Fabian Dorsch (eds.), *The New Evil Demon* (OUP).