

Evidence as Knowledge Indicators

This chapter puts forth a novel view of evidence in terms of knowledge indicators, and it shows that it is superior to its competition in that it can account for the epistemic impermissibility of resistance cases, as well as for the effect that resistance to evidence has on doxastic justification. Very roughly, knowledge indicators are facts that enhance closeness to knowledge: a fact e is evidence for S that p is the case if and only if S is in a position to know e and e increases the evidential probability that p for S .

7.1 Knowledge Indicators

In the previous chapter, I have argued that evidence resistance is an instance of input-level epistemic malfunctioning of our cognitive systems. Input-level malfunctioning is a common phenomenon in traits the proper functioning of which is input dependent, such as our respiratory systems. Since our cognitive systems, I have argued, are systems the proper functioning of which is input dependent, we should expect the failure at stake in resistance cases. I have also argued that, since pieces of evidence are pro tanto, prima facie justification-makers, they are the proper inputs to our processes of belief formation. When we have enough evidence and our belief-formation cognitive capacities are otherwise properly functioning, the resulting belief is epistemically justified. In turn, when our belief-formation capacities either fail to take up justification-makers that they could have easily taken up or they take them up but fail to output the relevant belief, they are malfunctioning.

The question that this chapter purports to answer is: how should we understand one's evidence such that we predict its normative impact on our properly functioning cognitive systems? Or, in other words, how should we understand one's evidence such that our account thereof predicts the epistemic impermissibility of resistance to evidence for cognitive systems that have generating knowledge as their epistemic function?

To lay my cards right on the table, the answer I will offer will make use of the notion of a knowledge indicator: on my view, evidence consists of knowledge indicators, which enhance closeness to knowledge by enhancing evidential probability. In turn, for any system S with a function F , since S ought to fulfil F , it is plausible that S ought to enhance closeness to F fulfilment. If so, our cognitive systems should take up pieces of evidence because they enhance closeness to function fulfilment (i.e. they enhance evidential probability and thereby closeness to knowledge).

Here is, in more detail, how I think about these things: evidence consists of facts. They can be facts about the world around us or mere facts about a subject's psychology. My having a perception as of a table in front of me is a psychological fact; it (pro tanto, prima facie) supports the belief that there is a table in front of me. So does the fact that there is a table in plain view in front of me.

In my view, evidence consists of facts that are knowledge indicators: facts that one is in a position to know and that increase one's evidential probability (i.e. the probability on one's total body of evidence) of p being the case. The fact that I see that there is a table in front of me is a piece of evidence for me that there is a table in front of me. It is a knowledge indicator, in that it raises the probability on my evidence that there is a table in front of me, and I'm in a position to know it.

Not just any psychological facts will constitute evidence that there is a table in front of me: my having a perception as of a table will fit the bill in virtue of having the relevant indicator property. Perceptions are knowledge indicators; the fact that I have a perception as of p is a fact that I am in a position to know, and that increases my evidential probability that p is the case. The fact that I wish that there was a table in front of me will not fit the bill, even if, unbeknownst to me, my table wishes are strongly correlated with the presence of tables: wishes are not knowledge indicators, for they don't raise my evidential probability of p being the case (although they may, of course, raise the objective probability thereof). For the same reason, mere beliefs, as opposed to justified and knowledgeable beliefs, will not be evidence material; they lack the relevant indicator property.

Here is the view in full:

Evidence as knowledge indicators: A fact e is evidence for one for a proposition p if and only if one is in a position to know e and one's evidential probability that p is the case conditional on e is higher than one's unconditional evidential probability that p is the case.

Or, slightly more formally, and where P stands for the probability on one's total body of evidence:

Evidence as knowledge indicators: A fact e is evidence for p for S iff S is in a position to know e and $P(p/e) > P(p)$.

Let's unpack the view further. What is it for me to be in a position to know e ? Plausibly, a certain availability relation needs to be instantiated. On my view, availability has little to do with the limits of my skull. Evidence may consist of facts 'in the head' or facts in the world. Some facts – whether they are in the head or in the world, it does not matter – are available to me; they are, as it were, 'at hand' in my (internal or external) epistemic environment. Some – whether in the head (think of justified implicit beliefs, for instance) or in the world, it does not matter – are not thus available to me.

Here are, for starters, some paradigmatic cases that illustrate what I'm talking about: if there is a table right in front of me, but I'm not paying attention to it, there is evidence for me that there is a table in front of me. If, unbeknownst to me, you put a new table in the other room, the fact that you put it there is *not available* to me: it is not evidence *for me*. Similarly, if I have some mental state that is so deeply buried in my psychology that I can't access it, it is not evidence for me.

As a first approximation, my notion of availability will track a 'can' for an average cogniser of the sort exemplified (e.g. with the relevant kind of cognitive architecture, social and physical limitations, etc.).

Here is some theory about this. First, there are *qualitative* limitations on availability: we are cognitively limited creatures. There are of *types* information that we just cannot access or process: the fact that there is a table in front of me is something that I can easily enough access. Your secret decision to put the table in the other room is not something I can easily access. There are also types of support relations that we cannot process: the fact that your car is in the driveway is evidence *for me* that you're home. But it's not evidence for my three-year-old son, Max, to believe that you're home. Max belongs to a variety of epistemic agents that are not sophisticated enough to process¹ the support relation into a belief that you are home. Evidence is not available to you if the kind of epistemic

¹ What is the relation between processing the support relation and knowledge indicators as probability enhancers? Is one supposed to be able to form probability beliefs in order to count as being able to process the support relation? The answer is 'no': merely treating an indicator as such is enough; awareness of its being one is not needed, nor is awareness of what makes a fact into an indicator. 'Treating' is a lowbrow affair: I can treat my cat as a friend without believing that she is.

agent that you are cannot access or process the particular variety thereof at stake (henceforth also *qualitative availability*).

There are also *quantitative* limitations on my information accessing and processing. The fact that there's a table somewhere towards the periphery of my visual field – in contrast of it being right in front of me, in plain view – is not something I can easily process: I lack the power to process everything in my visual field: it is too much information (henceforth also *quantitative availability*). My cognitive limitations make it such that the facts available to me are only a subset of what is going on in my visual field. More on this later.

The 'can' at stake here will be further restricted by features of the social and physical environment: we are supposed to read the newspaper on the table in front of us, but not the letter under the doormat. That's because we can't read everything, and our social environment is such that written testimony is more likely to be present in the newspaper on the table than under the doormat (henceforth also *environmental availability*).

In sum, for a fact to be such that I am in a position to know it, it needs to be at hand for me in my epistemic environment: at hand qualitatively (it needs to be the *type* of thing a creature like me can access and process), quantitatively (it needs to belong to the quantitatively limited subset of facts that a creature like me can access and process at one particular time), and environmentally (it needs to be easily *available* in my – internal or external – epistemic environment; i.e. in my mind or in my physical and social surroundings).

I take this availability relation to have to do with a fact being within the easy reach of my knowledge-generating cognitive capacities. A fact *e* being such that I am in a position to know it has to do with my having a properly functioning knowledge-generating cognitive capacity that can take up *e*:

Being in a position to know (BPK): S is in a position to know a fact *e* if S has a cognitive capacity with the function of generating knowledge that can (qualitatively, quantitatively, and environmentally) easily uptake *e* in cognisers of S's type.

A few crucial clarifications about this account: first, note that BPK is a sufficiency claim. It is not necessary that *e* is available to me in order for me to be in a position to know *e*: I can also come to know *e* via taking up facts that increase my probability for *e*.

Second, note that BPK is a restricted ought-implies-can: agent obligations imply capacities in the kind of cogniser that they are. This opens the account to a mild generality problem, of course: how to individuate the relevant type of cogniser? Stable, constitutive features will matter: cognitive architecture, inherent social and physical limitations. Fleeting, contingent

features will not (i.e. mere cognitive ‘furniture’): biases, previously held beliefs, wishes, among others. The advantage of the view is that, in restricting ‘ought implies can’ to types of cognisers, the account will predict that biased cognisers are in breach of their epistemic obligations: they may be unable to, for example, believe women because of bias, but cognisers with their cognitive architecture can, and therefore they should too.

Third, it is important to distinguish between being in a position to know and being in a position to *come* to know²: I am in a position to know that there is a computer in front of me; I am not in a position to know what is happening in the other room. I am, however, in a position to come to know the latter. Roughly, then, the distinction will, once more, have to do with epistemic availability: if all that needs to happen for me to come to know e is that my relevant cognitive capacities take up e and process it accordingly, then I am in a position to know e . If more needs to be the case – I need to open my eyes, or turn around, or go to the other room, or give you a call – I am in a position to come to know e but not in a position to know it. For now, I have not made any claim about the epistemic import of being in a position to come to know. Compatibly, being in a position to come to know might also, in some cases, deliver epistemic oughts: some cases of normative defeat and failure of evidence gathering are cases in point (e.g. see Lackey 2008, Goldberg 2016, 2017) See the next chapter for a discussion of this phenomenon.

Finally, and crucially, note that quantitative limitations on being in a position to know will make it so that I can only take up a limited number of the $e_1, e_2, e_3 \dots e_n$ facts that lie within reach with my knowledge-generating capacities. What facts go in my body of evidence in these cases? Which are the ones I am in a position to know, and which are the ones I am merely in a position to come to know (by changing focus, etc.)? On the account defended, in these cases, I will shoulder an epistemic obligation to take up a subset of $e_1, e_2, e_3 \dots e_n$ that is as large as my quantitative take-up limitations. Therefore, my body of evidence will only include the relevant subset that a creature with my cognitive architecture can (quantitatively) take up at one time. When looking straight at my computer, my visual field is populated with very numerous facts, such that taking them all up exceeds my quantitative take-up limitations. I am only under an obligation to take up a quantitatively manageable subset of facts.

The crucial question that arises is: which is the set that takes normative primacy and thereby delivers my body of evidence? *Availability rankings* will

² Many thanks to Ernie Sosa and Matt McGrath for pressing me on this.

deliver the relevant set, on my view: the most easily available subset of facts that I can take up delivers the set of evidence I have. In the case of visual perception, for instance, these are the facts located right in front of me, in the centre of my visual field, which are the brightest, clearest, etc. – in general, those facts that are most easily available to the cognitive capacities of a creature like me.

Tim Williamson (in conversation) worries that there will be cases in which too many facts (too many for my quantitative limitations) will have the same availability ranking. I see the worry (although I suspect it can be alleviated for most cases by our relation to space, time, complexity, brightness, etc.). Maybe the easiest case to imagine along these lines is the case of very simple arithmetical truths. In these cases, other normative constraints will have to decide the relevant set: I will have an all-things-considered obligation to attend to a particular range of simple arithmetical truths, and, among these, the most easily available will constitute my evidence, in virtue of them delivering the corresponding epistemic obligation to take them up.³

With the account fully unpacked, let's move on to checking how it fares on accommodating the resistance data.

7.2 Evidence and the Impermissibility of Resistance

Here are, first and foremost, a few theoretical virtues of this view of evidence. First, it is naturalistically friendly, in that it situates the epistemic normativity of epistemic oughts to believe within an etiological functionalist picture of normativity: epistemic oughts to believe have to do with the proper function of our cognitive capacities, just like biological oughts to take up oxygen have to do with the proper function of our respiratory systems.

Second, the view enjoys high extensional adequacy. In line with intuition, it predicts that there is evidence for the Gettierised victim that there is a sheep in the field: the fact that they have a perception as of a sheep is a fact that they are in a position to know and that raises their evidential probability that there is a sheep in the field.

Also, there is evidence for the (recently envatted)⁴ brain in the vat (BIV) for *p*: 'there is a tree in front of me' when they have a perceptual experience as of a tree, since that is a fact that they are in a position to know and that raises their evidential probability that there is a tree in front of them.

³ Thanks also to Matt McGrath for many discussions on this topic.

⁴ In line with content externalism, I take the BIV that has been envatted from birth to not have Earth contents and thereby no beliefs about Earth entities. Compatibly, my view predicts that they are justified to believe whatever they believe about vat entities.

There is no evidence for Norman the clairvoyant that the President is in New York: clairvoyant experiences are not evidential probability raisers when one is ignorant of the reliability of clairvoyance.

Finally, and most importantly for our purposes, it is easy to see that, when plugged into REEM, this view of evidence delivers the straightforward resistance intuition and thus explains that subjects in Cases 1–7 from Chapter 1 are in breach of their obligation to believe for failing to take up available evidence. Recall REEM:

Resistance to evidence as epistemic malfunction (REEM): A subject *S*'s belief-formation capacity *C* is malfunctioning epistemically if there is sufficient evidence supporting *p* that is easily available to be taken up via *C* and *C* fails to output a belief that *p*.

Anna's testimony in Case 1; media testimony, Dump's statements, etc., in Case 2; the scientific testimony in Case 3; the perceptual experience as of a table in Case 4; the partner's behavioural changes in Case 5; the fact that the Black students raise their hands in Case 6; and the incriminating fingerprints, etc., in Case 7 all constitute facts that are indicators of knowledge in virtue of being evidential probability enhancers that the subjects in these cases are in a position to know. These indicators of knowledge are easily available to creatures such as our protagonists: the subjects in Case 1–7 are members of a type of cogniser that hosts cognitive capacities with the function of generating knowledge that can easily take up these facts. Since they fail to do so, their cognitive capacities are malfunctioning, just like their lungs would be were they to be disinclined to take up the right amount of easily available oxygen. The account predicts that these subjects are all exhibiting resistance to evidence (by REEM) and are in breach of their obligation to believe (by OTB).

To see just how efficacious a view like mine is in accounting for evidence resistance and obligations to update, it will be useful to compare my account to $E = K$ once more. In *Knowledge and Its Limits*, Williamson considers an account of evidence in terms of being in a position to know, and he dismisses it based on the following rationale:

[. . .] suppose that I am in a position to know any one of the propositions p_1, \dots, p_n without being in a position to know all of them; there is a limit to how many things I can attend to at once. Suppose that in fact I know p_1 and do not know p_2, \dots, p_n . According to $E = K$, my evidence includes only p_1 ; according to the critic, it includes p_1, \dots, p_n . Let q be a proposition which is highly probable given p_1, \dots, p_n together, but highly improbable given any proper subset of them; the rest of my evidence is irrelevant to q .

According to $E = K$, q is highly improbable on my evidence. According to the critic, q is highly probable on my evidence. $E = K$ gives the more plausible verdict, because the high probability of q depends on an evidence set to which as a whole I have no access. (Williamson 2000, 189)

Two things about this: first, note that, in virtue of the quantitative limitations that my account imposes on being in a position to know, the view does not suffer from the problem Williamson points to here. Indeed, given that there is a limit to how many things I can attend to at once, it is only the most available subset that I can attend to that is part of my body of evidence.

Even more importantly, I submit that once we put flesh on the bones of Williamson's case, my view, and not $E = K$, gives the intuitively right prediction. Here it goes:

FRIENDLY DETECTIVE 2: It's highly probable that John killed the victim given that (p_1) John is a butler, (p_2) John is a very nice guy with an impeccable record, and (p_3) the only butler who's a very nice guy with an impeccable record was seen stabbing the victim. Friendly Detective is told p_1 , p_2 , and p_3 but can't get himself to believe p_3 because of wishful thinking, and he believes John didn't do it based on p_1 and p_2 .

FRIENDLY DETECTIVE 2 is an instance of Williamson's case. It is easy to see, however, that it is $E = K$ that delivers the counterintuitive result here: according to $E = K$, the detective is justified to believe John didn't do it. My view disagrees, and it scores on extensional adequacy.

Going back to the high societal stakes of evidence resistance: crucially, real-world, high-stakes cases of climate change denial and vaccine scepticism will sometimes be diagnosed by this account of evidence as evidence resistance. This will happen in cases of cognisers who have easily available evidence that climate change is happening and that vaccines are safe but fail to take it up and update their beliefs accordingly. It is compatible with this account, however, that this is not always the case: not all evidence rejection is evidence resistance. Sometimes, cognisers inhabit an epistemic environment heavily polluted with misleading evidence against the reliability of scientific testimony and public policy: if reliable testifiers in one's community testify that not- p : 'climate change is not happening', and one has every reason to trust them (say, because they have an exceptional track record of reliability as testifiers – although they get it wrong on this particular occasion), it can happen that one justifiably rejects evidence for p due to being in a position to know 'heavier' (albeit misleading) evidence against p . Note, however, that these cases – cases of justified evidence rejection – will be fairly specific cases epistemically that, while

they may happen in fairly isolated communities, the more one has access to evidence for p , the less justified their evidence rejection will be.

Now, all of this tells us that the account put forth is extensionally adequate: the view gets the resistance cases right. That is an important theoretical virtue of the view, and, as we have seen, it singles it out in the epistemic normative landscape.

That being said, extensional adequacy is not explanatory adequacy: even if thinking of evidence in terms of evidential probability increasers that one is in a position to know delivers the result that there is evidence for the subjects in Cases 1–7 that they fail to take up, the question as to why they should have done so remains open. One task remains, then, for the theorist of evidence resistance: explaining the normative force exercised by available evidence on our properly functioning cognitive systems. Or, in other words, explaining why, given the account of evidence proposed, it is epistemically impermissible for cognitive systems that have generating knowledge as their epistemic function not to take up easily available evidence.

Here it goes: some evidence I take up with my belief-formation machinery, whereas some I fail to take up, although I should. What grounds this ‘should’, in my view, is proper epistemic functioning.⁵ Because they are knowledge indicators, pieces of evidence are justification-makers: they are the proper inputs to our processes of belief formation that have generating knowledge as their function, and when we have enough thereof, and the processes in question are properly functioning in all other ways, the resulting belief is epistemically justified

Since evidence for S that p , on my account, consists of facts that enhance closeness to knowledge that p for S by enhancing S 's evidential probability for p , our cognitive systems are malfunctioning if they fail to take up easily available evidence, in virtue of thereby failing to take up opportunities for enhancing closeness to knowledge. Since for any system S with a function F , S should fulfil F , and it is plausible that S should enhance closeness to F fulfilment, and since the function of our cognitive systems is to generate knowledge, our cognitive system should take up enhancers of closeness to knowledge. Our cognitive systems should take up pieces of evidence because they enhance closeness to function fulfilment (i.e. they enhance evidential probability and thereby closeness to knowledge).

In turn, when our belief-formation capacities either fail to take up knowledge indicators that they could have easily taken up or they take

⁵ See, for example, Millikan (1984), Graham (2012), and Simion (2019a, 2023a).

them up but fail to output the corresponding belief, they are malfunctioning. A subject *S*'s belief-formation capacity *C* is malfunctioning epistemically if *S* has sufficient evidence supporting *p* that is available to be taken up via *C* and *C* fails to output a belief that *p*.

7.3 Infallibilism: Evidence and Knowledge

Before moving on, I would like to address an important worry that has been put forth in recent literature for views of evidence like the one defended in this chapter (i.e. knowledge-centric views of evidence).

Most contemporary epistemologists are fallibilists: they think that you can know a proposition *p*, even if your evidence does not entail that *p*. In recent work, Jessica Brown (2018) offers a thorough defence of fallibilism against knowledge-centric views of evidence, or what I will dub 'new infallibilism'. More specifically, her central aim is to show that epistemologists who also want to be non-sceptics and want to endorse a non-shifty view of knowledge attributions should be fallibilists rather than new infallibilists. To this end, Brown argues that there is reason to think that fallibilism compares favourably with new infallibilism when it comes to evidence and evidential support. Perhaps most importantly, Brown identifies and takes issue with three key commitments of the new infallibilist's view of evidence, to wit:

The factivity of evidence: If *p* is part of one's evidence, then *p* is true.

The sufficiency of knowledge for evidence: If one knows that *p*, then *p* is part of one's evidence.

The sufficiency of knowledge for self-support: If one knows that *p*, then *p* is evidence for *p*.

Brown argues against all three of these claims. Since fallibilists can avoid these commitments, the thought goes, fallibilism scores points against new infallibilism.

The account of evidence I defended in this chapter implies all of the claims above. As such, if Brown is right, my account is in trouble, alongside its E = K Williamsonian cousin.

However, I think that there are ways to be an infallibilist that survive Brown's excellent arguments. Thus, in what follows, I will explore ways in which new infallibilism can resist both Brown's case against infallibilism and her fallibilist response to at least some of the data points that have been thought to favour the new infallibilism.

Let's start by looking at Brown's argument against the sufficiency of knowledge for evidence (i.e. the claim that if one knows that p , then p is part of one's evidence). Brown's key idea is to appeal to citable evidence. She points out that one cannot felicitously cite p when queried about one's evidence for p , not even if one knows that p (Brown 2018, 49–50). But given that knowledge is sufficient for evidence, it is hard to see why this should be the case.

Note, however, that fallibilists, too, will need an account of when p is part of one's evidence. I can think of a few options here: if p is justified for one/if one believes that p /if one justifiably believes that p , then p is part of one's evidence. Crucially, since knowledge entails justified belief, their view entails the sufficiency of knowledge for evidence, no matter which of these options the fallibilist goes for. This means that in cases in which one knows that p , it is equally hard for fallibilists to explain why one cannot cite p when queried about one's evidence for p . In this way, there is no reason to think that new infallibilism is at a disadvantage here.

Let's move on to another of the claims above: the sufficiency of knowledge for self-support (i.e. that if one knows that p , then p is evidence for p). Why think that new infallibilists are committed to this claim in the first place? Here is Brown:

To see why the infallibilist should embrace the *Sufficiency of knowledge for self-support*, consider [...] knowledge by testimony, inference to the best explanation and enumerative induction. It's hard to see how one has evidence for what's known in these ways which entails what's known without allowing that if one knows that p , then p is part of one's evidence for p . [...] So, it seems that embracing the *Sufficiency of knowledge for self-support* is the best way for the infallibilist to avoid scepticism. (Brown 2018, 43)

I agree that it may be hard for fallibilists to see how one can have the evidence for what is known here unless one subscribes to the sufficiency of knowledge for self-support. However, the same is not true of new infallibilists. Note that, according to new infallibilism, what one's evidence is will turn on worldly states (e.g. on the friendliness of the epistemic environment one finds oneself in). For instance, what is one's evidence for the claim that there is a barn before one may vary depending on whether one is in Normal Barn County or in Fake Barn County. But once this point is properly appreciated, there is little reason to think that testimony, inference to the best explanation, and enumerative induction pose a particularly difficult problem. While data from testimony, inference to the best explanation, and enumerative induction may not entail what is

known, they may do so when conjoined with a sufficiently friendly epistemic environment.

This leaves the factivity of evidence (i.e. p is part of one's evidence only if p is true). Brown relies on a familiar line of objection to this claim. Here is Brown:

As is well-known, this conception of evidence [which combines the factivity of evidence with the sufficiency of knowledge for evidence] is open to the objection that it holds that certain pairs of subjects who are intuitively equally justified in some claim (e.g. a person and her BIV twin), are not equally justified. (Brown 2018, 22)

Brown considers a response on behalf of new infallibilists in terms of blamelessness.⁶ Their key idea is that while BIVs don't believe justifiably, they are nonetheless blameless for their beliefs. At the same time, there is empirical evidence that suggests that we are prone to mistaking cases of unjustified but blameless belief for cases of justified belief, which is why intuition leads us astray in these cases.

According to Brown, this move remains unsuccessful. Her strategy is to look at a number of ways of analysing what blamelessness amounts to and to argue that none of these ways will do the trick for new infallibilists.

Note, though, that while it is true that the particular infallibilists (e.g. Williamson, Littlejohn) that Brown discusses have historically held a view that equates justification and knowledge, this is optional to new infallibilisms. There has been a surge of views in the literature that explain justified belief in terms of knowledge without identifying justified belief and knowledge (e.g. Bird 2007, Ichikawa 2014, Miracchi 2015, Kelp 2018, Schellenberg 2018, Simion 2019a). Champions of these views have argued at great length that these views can allow for agents in bad cases (e.g. BIVs) to be justified. If so, they can successfully explain the intuition at issue here. Crucially, the view of justification defended here is precisely one such view: on this account, BIVs believe justifiably insofar as they employ properly functioning cognitive capacities with the function of generating knowledge – which, by stipulation in the justification-intuition-triggering cases (paradigmatically, of recently envatted BIVs) they do. At the same time, and crucially, this view of justification is entirely compatible with new infallibilism. After all, what is key to new infallibilism is a view about the relation between knowledge and one's evidence.

⁶ See Brown (2018, 70–73).

7.4 Conclusion

On the account defended here, one's evidence consists in facts that one is in a position to know and that increase one's evidential probability that something is the case. In turn, being in a position to know has to do with the variety of cogniser at stake: should one be the kind of cogniser that hosts cognitive processes that are able to pick up the relevant facts from the world, the facts at stake will belong in one's body of evidence.