

Resistance to Evidence as Epistemic Malfunction

This part of the book develops a full positive epistemology: an account of the epistemic normativity of evidence resistance, in conjunction with novel accounts of epistemic oughts, evidence, defeat, and permissible suspension. This first chapter argues that resistance to evidence is an instance of epistemic malfunction. It first puts forth a normative picture according to which the epistemic function of our cognitive systems is generating knowledge, and epistemic norms drop right out of this function. Second, it shows how this picture accommodates epistemic obligations, which, in turn, explain the normative failure instantiated in cases of resistance to evidence. According to this view, cognitive systems that fail to take up easily available evidence and defeat instantiate input-level malfunctioning. 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 argue, are systems the proper functioning of which is input dependent, we should expect the failure at stake in resistance cases.

6.1 Epistemic Oughts

Let us start with the elephant in the room: the epistemic impermissibility of evidence resistance implies that there are such things as epistemic oughts to govern our practices of belief forming, updating, and maintaining. As such, any epistemology that is able to predict epistemic impermissibility in resistance cases will be an epistemology that is able to incorporate epistemic oughts.

This is not trivial. It is not trivial, first and foremost, methodologically: until very recently, normative work in epistemology has, for the most part, been negative, in that it has concerned itself with restricting what we are

permitted to, for example, believe, assert, or use as a premise in reasoning. Investigations into epistemic oughts are thin on the ground.¹

It is also not trivial normatively: permissions are not easily turned into obligations (henceforth I am following the literature in deontic logic² in using ‘oughts’ and ‘obligations’ interchangeably: nothing hinges on it – I am employing a notion of obligation that merely maps onto an ought). Just because I’m permitted to believe things about objects located towards the periphery of my visual field, it does not follow that I ought to do so – after all, a lot of things are happening in my visual field; I cannot possibly be expected to form beliefs about all of them. Furthermore – and going back to methodological difficulties – some permissions endorsed by traditional epistemological frameworks simply don’t speak at all towards the kind of obligation I’m breaching when I’m evidence resistant. We have seen this problem surface clearly with permissions to believe in virtue of knowing: this permission is silent when it comes to obligations to update. Similarly, think of classical process reliabilism: even if we grant its champions that we are permitted to believe the outputs of our reliable belief-formation processes, the view remains silent on the norms I’m breaking when my processes fail to deliver outputs that they should deliver.

What we need is an epistemological framework that has the resources to be naturally extended to incorporate obligations to update.

Another, more well-known difficulty has to do with the nature of these epistemic oughts: most people think voluntarism about belief is false. Even bold voluntarists would likely accept that voluntarism is false about the vast majority of doxastic phenomena. Notably, people have worried about non-voluntarism being incompatible with epistemic oughts for belief: if ought implies can, the thought goes, and since I cannot believe at will, I cannot be subject to norms obliging me to do so either.

Now, a lot of ink has been spilled on the implausibility of an unrestricted ‘ought-implies-can’ principle for normativity in general, and on putting forth more or less successful restrictions thereof (e.g. see Ryan 2003 for excellent work on this topic). One thing that has become clear in recent years in epistemology is that, while we don’t want an individual-based ‘ought-implies-can’ to constrain our normativity – after all, we don’t want the sexist cogniser to come out as justified in his sexist beliefs just

¹ See Kornblith (2001), Fricker (2007), Chrisman (2008), Feldman (2008), Goldberg (2016, 2017), Lackey (2019), Jenkins-Ichikawa (2020, forthcoming), Kelp (2022), and Simion (2023a) for exceptions. In putting this distinction in terms of positive versus negative epistemology, I follow Jenkins-Ichikawa (2020).

² See, for example, McNamara and Van De Putte (2022).

because he can't believe otherwise due to his sexism – we do want general facts about limitations pertaining to adult human cognitive architecture to play a role in restricting our normative claims. Norms governing ideal epistemic agents are just not very informative, nor is there a straightforward way to go from ideal to non-ideal epistemological theorising; for instance, it's not clear what the obligations of an ideal agent who knows that they know for all of their pieces of knowledge imply for people like me, who can hardly come anywhere close to such achievements (nor do we particularly want to).

Keeping all of this in mind, several views in the literature attempt to offer accounts of epistemic ought that bypass the voluntarism objection; I will not run through all of them in great detail here. However, one thing that transpires from even a brief sketch of the literature is that the endeavour of incorporating epistemic oughts in one's epistemology is faced by a strength dilemma: make the source of the ought too thick, normatively, and it seems incompatible with non-voluntarism. Make it too thin and it seems to fail to capture intuitions of 'normative oomph' when it comes to the epistemic (i.e. that epistemic norms are, in a significant way, '*heavier*' than mere conventions). To see the dilemma, let's look at a couple of classic proposals.

According to Richard Feldman, epistemic obligations are what he calls 'role oughts':

There are oughts that result from one's playing a certain role or having a certain position. Teachers ought to explain things clearly. Parents ought to take care of their kids [. . .]. Incompetent teachers, incapable parents [. . .] may be unable to do what they ought to do. Similarly, I'd say, forming beliefs is something people do. That is, we form beliefs in response to our experiences of the world. Anyone engaged in this activity ought to do it right. In my view, what they ought to do is to follow their evidence (rather than their wishes or fears). I suggest that epistemic oughts are of this sort – they describe the right way to play a certain role. (Feldman 2000, 677)

Feldman's account is the paradigmatic case of a theory that is affected by the lack of 'normative oomph' worry: after all, role oughts can – and often are – in an important sense uninteresting, or even bad oughts. Role oughts generated by the role of mafia boss, for instance, seem too normatively thin to constitute the right kind of model for epistemic normativity.

A similar worry affects Hilary Kornblith's view: according to Kornblith, a desire for truth, together with the corresponding hypothetical imperatives it generates, is all we need for explaining epistemic oughts. That's all there is to epistemic normativity: not much more 'normative oomph' than this is needed to explain the data we need to explain.

I share Kornblith's naturalism as well as, to a large extent, his 'normative oomph' scepticism. I think, however, that a desire-based picture remains unsatisfactory on precisely the desiderata Kornblith and I both endorse: naturalistic friendliness. To see why, note that desire is also, plausibly, normatively constrained: there are such things as more and less fitting desires. In turn, hypothetical imperatives are only as thick, normatively, as the desire triggering them is fitting. Ideally, we want to explain this datum in a naturalistically friendly fashion: we want to be able to predict – in a naturalistic normative framework – why some norms are thicker than others, and we want, ideally, to get the result that the epistemic is on the thicker side of things than, again, norms generated by mafia-boss-characteristic desires. In other words, we should be able to predict that nature cares more about the epistemic than about the mafia. Kornblith's response is that, as opposed to other desires, a desire for truth is pretty much implied by a desire for anything else, since whatever else one might be trying to achieve, truth will be needed in order to make the right decisions to the aim of getting it. This instrumental way of looking at things doesn't help much with a mafia boss-type issue, though: after all, on this instrumental picture, a desire for truth will be just as fitting as the desires it is instrumental to fulfilling; in this, the view fails to beef up the normative oomph-ness of truth-generated norms.

Finally, Matthew Chrisman ventures to explain belief oughts in a non-voluntarist-friendly fashion by pointing out a classic distinction between norms governing actions, or 'ought-to-dos' – which may be subject to some variety of the ought-implies-can principle – and 'rules of criticism', governing states, or 'ought-to-bes', which, Chrisman thinks, are not thus constrained. Here is Chrisman:

In developing an account of robustly normative claims about what someone ought to believe, my strategy is to treat these as adverting to a species of state norms. For instance, the claim, 'You ought to believe you are reading this text right now' could be understood to be an instance of the general form, 'S ought to be in doxastic attitude A towards proposition p under conditions C'. Then, the crucial observation is that some such normative claims seem to be true, just like 'Clock chimes ought to be disposed to strike on the quarter hour', 'The beds ought to be made by 8am every morning', 'A child ought to be able to tie their shoes by age seven' [. . .]. Yet these true normative claims don't presuppose that their subjects be capable of voluntarily following the rule. (Chrisman 2008, 362)

The worry for Chrisman's account is, once more, one having to do with 'normative oomph': state norms are clearly not always normatively

thick – indeed, several of the examples given are paradigmatic examples of social standards of correctness. Is there any interestingly thick normative sense in which beds that are unmade by 8 a.m. are defective? Not really. If so, we need to find extra normative resources to account for the epistemic.

The general lesson to be learnt, I think, is that we need a model for epistemic normativity that, at the same time, circumvents voluntarism worries and is normatively thick enough to account for the intuitively non-conventional nature of the epistemic. Otherwise, if the model put forth is too normatively thin, the suspicion will be that it only circumvents voluntarist worries *in virtue* of its normative thinness – and thus that an extensionally more adequate, normatively thicker incarnation thereof will fail to do so.

I think that the view we are after should be naturalistic, exhibit prior normative plausibility, and be generalisable to other normative domains but also, at the same time, have enough ‘normative oomph’ to explain the intuitive categoricity of epistemic normative constraints. In the next sections, I will develop a view that purports to meet all of these desiderata: on this account, epistemic normativity is etiological functionalist normativity.

6.2 Functions and Norms

In traits, artefacts, and practices alike, functions generate norms.³ There is such a thing as a properly functioning heart, a properly functioning can opener, and a proper way to make coffee. If that is so, when we are interested in a particular type of norm governing a particular type of trait, it is helpful to first identify its function.⁴

On the etiological theory of functions,⁵ functions turn on histories that explain why the item exists or operates in the way it does. Take my heart;

³ See also, for example, Millikan (1984), Graham (2012), and Sullivan-Bisset (2017).

⁴ I am not arguing for the functionalist picture here (I have done so many times (e.g. Simion 2019a, 2020, 2021a) and at considerable length); I assume it, develop the corresponding views, and show that they nicely explain the data that competing views have been struggling with. I also do not make the corresponding descriptive claim: the claim that reasoning is a functional kind (i.e. a kind that only affords a functionalist analysis and not a dismantling analysis). In a nutshell: for my purposes here, I only subscribe to normative functionalism, not metaphysical functionalism.

⁵ Defended by people like Ruth Millikan (1984), Karen Neander (1991), and David J. Buller (1998). The etiological theory of functions is, by far, the most widely endorsed view in the literature. Its main competitor is the ‘systemic’ theory of functions, notably defended in Cummins (1975). Systemic functions describe how something works or operates – what it does – as a part of a larger system. Functions, in this sense, are the causal role capacities of parts that contribute to some capacity of the containing system. Systemic functions are widely believed to lack normative import, which is what explains, to a large extent, both the popularity of the competing, etiological account and why the latter is thought to be much better suited for applications to normative domains like epistemology.

plausibly, tokens of the type pumped blood in my ancestors. This was beneficial for my ancestors' survival, which explains why tokens of the type 'heart' continue to exist. As a result, my heart acquired the etiological function (henceforth also e-function) of pumping blood. Acquiring an etiological function is a success story: traits, artefacts, and actions get etiological functions of a particular type by producing the relevant type of benefit. My heart acquired a biological etiological function by generating biological benefit. Through a positive-feedback mechanism – the heart pumped blood, which kept the organism alive, which, in turn, ensured the continuous existence of the heart – our hearts acquired the etiological function of pumping blood.

Importantly, while aetiology does require some history of beneficial effects, it does not require an awful lot of it; what it all amounts to is explaining the existence/continuous existence of a trait through a longer or shorter history of positive feedback:

Functions arise from consequence etiologies, etiologies that explain why something exists or continues to exist in terms of its consequences, because of a feedback mechanism that takes consequences as input and causes or sustains the item as output. (Graham 2014, 35)

Functions can be of different sorts: there are biological functions, aesthetic functions, social functions, etc. In contrast to the Graham/Millikan view, my account takes functions to be typed by the corresponding benefit. As such, if a trait produces a benefit of type B in a system, the function thereby acquired will be a function of type B. The heart's function to pump blood is a biological function in virtue of the fact that the produced benefit is also biological (i.e. survival). The function of art is an aesthetic function in virtue of the fact that the produced benefit is an aesthetic benefit. Now, of course, aesthetic benefit might, and often will, also result in biological benefit. This, however, in no way renders the function at stake a biological function. What is important to keep in mind is that the benefit that is *essential* to aesthetic function acquisition is the aesthetic one. The fact that biological benefit is also associated with the latter is a mere contingent matter of fact. Here is, then, the full etiological account to be employed here:

E-function: A token of type T has the e-function of type B of producing effect E in system S iff (1) tokens of T produced E in the past, (2) producing E resulted in benefit of type B in S/S's ancestors, and (3) producing E's having B-benefitted S's ancestors contributes to the explanation of why T exists in S.

Note that etiological functions are successes. They explain the continuous existence of the trait that bears them because that is so. The etiological economic function of knife-producing economic systems is not just that of producing knives; it's to produce good, sharp knives. To see this, note that the positive-feedback loop that is presupposed by etiological function acquisition – the trait produces the effect and the effect benefits the system and thereby contributes to the explanation of the continuous existence of the trait – presupposes a history of success. What contributes to the explanation of the continuous existence of knife-producing economic systems is their producing good, sharp knives. Blunt ones would plausibly not have done the trick.

Just as the economic function of knife-producing economic systems is to produce good, sharp knives, the epistemic function of our belief-forming systems is to produce good beliefs. Mere belief, then, is a failure on the part of our cognitive system to fulfil its epistemic function just as blunt knives are failures on the part of knife producers to fulfil their economic function.

Note that functions will also come with associated norms: these prescribe the right way to proceed in order to reliably enough⁶ fulfil the function in question under normal conditions. Because its function contributes to the explanation of its very existence, the trait in question *ought* to perform in a way that is associated with likely enough function fulfilment. Your heart will be properly functioning when it functions in the way that reliably enough delivers function fulfilment under normal conditions: it will beat at a rate between 40 and 100 beats/minute, which, under normal conditions (i.e. conditions similar to those present at the moment of function acquisition), reliably leads to pumping blood in your circulatory system.

In my view, generating knowledge is the function of our epistemic practice of inquiry; in turn, cognitive processes, in virtue of their being the central mechanisms engaged in our practice of inquiry, inherit this function of generating knowledge. Further, norms governing moves in inquiry – such as beliefs, suspensions, withholdings, credences, assertions, or pieces of reasoning – will drop out of this knowledge-generating function.

Moves in practices generally aim – either directly or indirectly – to fulfil the function of the practice. The difference between direct and indirect aiming lies with achievability: for some moves in practices, the function of

⁶ How reliable is reliably enough? The threshold varies with the type of normativity, the type of trait, and evolutionary needs. The heart (biologically) ought to be very reliable for a species' continuous existence; in contrast, sperm cells only need to fulfil their function very rarely.

the practice is only indirectly rather than directly reachable. Cardiologist consults are moves in the practice of medicine, and they aim directly at fulfilling the goal of the practice of medicine: curing diseases. In turn, some moves aim at this final practice goal only indirectly while aiming directly at intermediate goals: performing electrocardiogram (ECG) tests aims (directly) at informing the cardiologist as to how well the patient is doing, which, in turn, aims directly at their diagnosing the patient correctly and, further, at curing their disease. In this, ECG tests aim indirectly at the function of the practice of medicine: curing diseases. They aim at making progress towards it. Similarly, baking cakes is a move in the practice of cooking, and it aims directly at fulfilling the function of the practice: producing tasty, nourishing food. My getting the flour off the shelf aims indirectly at the general function of the practice by aiming directly at adding flour to the cake mix. It aims at making progress towards producing tasty, nourishing food. And so on.

Similarly, on the epistemic normative picture I favour, moves in the practice of inquiry – that is, all epistemically significant states and actions – will be governed by norms borne out by this central knowledge function of the practice. Moves in inquiry, that is, will aim either directly (plausibly: beliefs, assertions, reasonings) or indirectly (credences, suspensions, withholdings) at the aim of the practice of inquiry. The difference, once more, will lie with goal achievability: beliefs, assertions, and conclusions of reasonings can be knowledgeable in a way in which things like credences, suspensions, and withholdings cannot. Belief formation aims directly at the aim of the practice of inquiry (knowledge), whereas, at the same time, credence, withholding, and suspension aim at knowledge indirectly: they are transitional attitudes, in the sense in which these are attitudes held en route to knowledge but are not in the running for knowledge.

Just like the biological functions of our hearts generate etiological biological norms, the epistemic functions of our cognitive systems generate etiological epistemic norms governing our belief formation.

In previous work (e.g. Simion 2019a, 2021a), I have argued extensively that knowledge is the etiological function of our cognitive systems. This makes my epistemological approach, in a crucial sense, a knowledge-first epistemological approach: epistemic normative constraints, on this view, drop out of the knowledge function of our cognitive systems.

A few things by way of clarification are in order: in the more than twenty years since Timothy Williamson's (2000) seminal work putting forth the first defence of a full knowledge-first epistemological framework, the knowledge-first research programme has generated an impressive

amount of high-quality work and very promising results across epistemological subfields (e.g. epistemic justification, defeat, evidence, epistemic normativity, social epistemology, know-how, the nature and normativity of inquiry) and also at the intersection of epistemology with philosophy of language (e.g. the nature of speech acts, contextualism), mind (e.g. the nature of mental states), and moral and political philosophy (e.g. blame, trust, responsibility, political discourse).

Against this backdrop, however, this book zooms in *only* on what I take to be the core normative claims of the knowledge-first programme: that knowledge is the central epistemic value, and that thereby central normative notions in epistemology – such as justification, evidence, and defeat – are to be analysed in terms of knowledge. This view, in contrast to its extant knowledge-first competition, analyses epistemic normative categories in terms of knowledge without requiring further theoretical commitments (e.g. to the non-analysability of knowledge or to knowledge being a mental state in its own right). These questions remain open.

The key claim that I endorse is that generating knowledge is the function of our cognitive systems. To see why this is plausible, very briefly, note that knowledge meets the conditions for an e-function: it is plausible that it has been generated by our cognitive systems and those of our ancestors in the past (after all, knowledge is readily available in our environment), that this benefitted our organisms (e.g. by informing us about the presence of predators and the whereabouts of food), and that the fact that knowledge benefitted us in this way contributes to the explanation of why cognitive systems continue to exist in individuals like us. So generating knowledge is at the very least *one* function of our cognitive processes. Is it the main function? Note that what determines the proper level of generality for main function individuation is the T-value of the relevant T-function. The main biological function of the heart, for instance, maps onto its most valuable biological contribution: its main function is not ‘pumping blood and making a ticking sound’, but neither is it merely ‘pumping’, for instance. Plausibly, that is because if the heart pumps orange juice in our circulatory system, that’s not very valuable for our survival. I submit that knowledge is more valuable than any lesser epistemic standing⁷: that much is very widely accepted in the literature. It is easy to see that, if I am right about main function individuation, the distinctive value of knowledge thesis together with e-function deliver the

⁷ See Kelp and Simion (2017) for an account of the value of knowledge as commodity.

result that the main epistemic function of our cognitive processes is generating knowledge.

Functions provide us with a straightforward way to identify the norms governing a particular trait or activity: first, if we are interested in identifying a norm of type T of the relevant trait, we start off by identifying its type T function. Once that is done, we look at the way in which it is reliably fulfilled under normal conditions. That will give us its conditions for proper function and, correspondingly, the content of the norm of type T we were looking for (Simion 2019a).

In what follows, I employ the functionalist machinery in investigating the epistemic normativity of belief. Before we move on, however, let us go back to the strength dilemma generated by voluntarist worry and the question of ‘normative oomph’: the good news is that, clearly, norms of proper function do not imply any voluntarist claim: just like my heart is governed by norms of proper function about blood pumping, my cognitive capacities are governed by norms of proper function about belief forming.

Is functionalist normativity going to generate a normatively thick enough model for epistemic normativity? Let’s go back to norms and practices in the domain of mafia. Here is a worry one might have: mafia practices may well continue to exist because they achieve the corresponding ‘values’ internal to the domain of mafia. If so, on a functionalist picture, we can get norms out of these functions: norms that regulate proper ways of going about achieving the domain-specific values that the domain of mafia is organised around. But do the resultant ‘norms’ have any normative *oomph* at all?

One would think that the case of epistemic norms is different. For instance, the fact that S said that *p* and that S is very reliable on *p*-related matters seems like a reason to believe that *p* has normative force – I can’t just set it aside in the way I can set aside the orders from my crime boss (e.g. Fricker 2007). These cases *feel* different in a way that needs to be explained.

On functionalism, indeed, it can be the case that, for example, efficient killing is a domain-specific value in the domain of mafia, which, in turn, generates corresponding (domain-internal) functionalist norms. It’s a completely different question, however, if the domain itself is, for example, valuable *simpliciter* to begin with – and I take it that the domain of mafia is not. If so, the normative force of its norms will be restricted to the domain of mafia.

I take it to be empirically plausible that doing well epistemically is, at least for the most part, good for us biologically. If so, the domain of the

epistemic, in contrast to the domain of mafia, will itself be valuable; if so, its internal functional norms will bear 'normative oomph' in a way in which norms of the domain of mafia do not. Importantly, though, I can afford to stay neutral on *the extent* to which the epistemic is good or bad: for 'normative oomph', I just need the fairly weak claim that it is generally good for us. This is important because it gives my view the flexibility to adapt to empirical results that purport to show that, at times, epistemic unreliability co-varies with biological benefits.

Compatibly, though, and plausibly, the domain of epistemology is valuable for our survival in a way that generates thick normative constraints. My functionalism thus does have the capacity to distinguish different kinds and strengths of normative force. On my view, epistemic norms have (1) domain-bound normative force, in that they promote knowledge, which is the value around which the domain is organised, and (2) non-domain-bound normative force, in that 'the epistemic' is a domain that is (empirically plausibly) valuable for our survival.⁸

6.3 Resistance to Evidence as Epistemic Malfunction

On my account, the main function of our belief-formation systems is to generate knowledge. In turn, they are properly functioning just in case they work in a way that is normally conducive to generating knowledge. When that happens, the beliefs they generate are justified.

I dub my view of justification 'knowledge-first functionalism': the account is functionalist in that it follows Millikan (1984), Burge (2010), and Graham (2012) in taking the epistemic normativity of belief to drop out of the epistemic function of our cognitive capacities. It is knowledge-first epistemological in that, unlike traditional, truth-first functionalism, it unpacks the function at stake in terms of knowledge. Here is a more precise formulation of the view:

⁸ Is it really knowledge, though, that is good for our survival rather than simply true beliefs? I have argued extensively in previous work (e.g. Simion 2016, 2020, Kelp and Simion 2017) for the distinctive value of knowledge over and above true beliefs that fall short of knowledge. In my view, knowledge as a commodity is distinctively valuable in that it is *our way* of correctly representing the world around us in virtue of its ready availability, just like water is our way of quenching our thirst in virtue of its ready availability: in a wide range of areas, knowledge is widely and readily available, just like water is widely and readily available. If so, just like water is a valuable commodity to us because it is our way of quenching our thirst, knowledge is valuable to us because it is our way of representing the world.

Knowledge-first functionalism (KFF): A belief is *prima facie* justified if and only if it is generated by a properly functioning cognitive capacity that has the etiological function of generating knowledge.

On this knowledge-centric picture, good belief is knowledgeable belief, while justified belief – belief that is permissible by the epistemic norm of belief – is belief generated by a properly functioning cognitive capacity that has the etiological function of generating knowledge. The standards for proper functioning are thus natural normative standards, and they are constitutively associated with promoting knowledgeable beliefs.

So far, we have been solely talking in moderate normative terms: we have taken justification of belief to supervene on epistemic permissibility and, in turn, epistemic permissibility to have to do with the proper function of our knowledge-generating belief-formation systems. How does resistance to evidence fit in this picture? After all, an epistemic subject being resistant to evidence seems to be a matter of obligation breach rather than permissibility breach: one is resistant to evidence insofar as there is evidence one should take up but one fails to do so.

I would like to propose an account of epistemic obligation according to which what all of subjects in Cases 1–7 from Chapter 1 have in common, epistemically, is that they are in breach of their epistemic ought to believe in virtue of resistance to available evidence. Here it is:

The ought to believe (OTB): A subject *S* has an *epistemic*⁹ obligation to form a belief that *p* if there is sufficient and undefeated evidence for *S* supporting *p*.

Once more, importantly: when I talk of obligation, I am following the literature in deontic logic¹⁰ in using oughts and obligations interchangeably (i.e. I am employing a notion of obligation that merely maps onto an ought). Now, note that OTB, together with a moderate evidentialist assumption that one's belief that *p* is epistemically justified only if there is sufficient and undefeated evidence for *S* supporting *p*, straightforwardly implies that epistemic justification is epistemic obligation (and, more generally, that justifiers are obligations). One might wonder, at this stage: is that right? After all, the vast majority of the literature assumes that epistemic justification is mere epistemic permission.

⁹ Crucially, the obligation at stake is merely epistemic. Compatibly, prudential obligations, for example, might override the epistemic obligation and render it all-things-considered permissible to dismiss information that we are not interested in.

¹⁰ See, for example, McNamara and Van De Putte (2022).

A few things about this: first, I take it that what resistance cases show is that this assumption was wrong all along. Mere epistemic permissions cannot, in virtue of their weak normative force, explain why the main characters in these cases (epistemically) *ought* to take up some evidence that they fail to take up.

Second, as I'm about to argue, there are in-principle theoretical reasons for which we should be suspicious of the thought that epistemic justification is mere epistemic permission. Here it goes: note that defeaters are obligations – when our justification is defeated, by definition, it is impermissible to ignore defeat and hold on to the corresponding doxastic attitude. Since it is impermissible to ignore defeaters, it follows that they constitute epistemic obligations (since it is always permissible to ignore mere permissions). Note, also, that there is such a thing as merely partial defeat: these are garden variety cases in which the epistemic agent needs to weigh their evidence in favour of p against their evidence against p . If Mary tells me that the train comes at 8 a.m. and you tell me that the train comes at 7 a.m., and Mary and you are, for all I know, equally reliable testifiers, my justification to believe that the train comes at 8 a.m. is partially defeated – I should lower my confidence in this being the case, but I don't have sufficient epistemic support to move to outright believing that the train doesn't come at 8 a.m. Similarly, I don't have enough epistemic support to believe or disbelieve what you said either: it is impermissible both to form an outright belief that the train comes at 7 a.m. and to form an outright belief that it does not. Justifiers and defeaters can outweigh each other.

However, if defeaters constitute epistemic obligations, and if defeaters and justifiers can outweigh each other, it follows that justifiers constitute epistemic obligations as well: otherwise, if they constituted mere permissions, they would be normatively inert against defeaters, since permissions are normatively inert against obligations. As such, it seems as though the mere possibility of partial defeat implies that justification maps on to epistemic obligation. Here is the argument unpacked for the reader's convenience:

- (1) Defeaters affect what one is justified to believe.
- (2) If defeaters affect what one is justified to believe, then it is epistemically impermissible to fail to adjust one's doxastic attitudes in light of defeaters.
- (3) It is epistemically impermissible to fail to adjust one's doxastic attitudes in light of defeaters (from (1) and (2)).

- (4) If it is epistemically impermissible to fail to adjust one's doxastic attitude in light of a normative consideration x , then x constitutes an epistemic obligation.
- (5) Defeaters constitute epistemic obligations (from (3) and (4)).
- (6) Defeat can be partial defeat.
- (7) If defeat can be partial defeat, then justifiers can carry normative weight against defeaters.
- (8) Justifiers can carry normative weight against defeaters (from (6) and (7)).
- (9) Justifiers constitute either epistemic permissions or epistemic obligations.
- (10) Permissions cannot carry normative weight against obligations.
- (11) Epistemic justifiers constitute obligations (from (5), (8), (9), and (10)).

What is the source of epistemic obligation? What grounds the epistemic OTB, in my view, is proper epistemic functioning. Pieces of evidence are pro tanto, prima facie justification-makers: they are the proper inputs to our processes of belief formation, and when we have enough evidence and the processes in question are otherwise properly functioning, the resulting belief is epistemically justified. In turn, when our belief-formation processes either fail to take up justifiers that they could have easily taken up or they take them up but fail to output the corresponding belief, they are malfunctioning:

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 .

The proper function of belief-formation capacities, then, on my view, is input dependent: failing to take up the right inputs – whether this occurs by taking up the wrong inputs or by failing to take up the right inputs – is an instance of malfunctioning.

One illuminating analogy here is the proper functioning of the lungs: as opposed to functional traits whose proper function is not input dependent (e.g. hearts), what it is for our lungs to function properly is, partly, for them to take up the right amount of the right stuff (i.e. oxygen) from the environment. Lungs that fail to do so, in environments where oxygen is easily available, are improperly functioning – whether they fail via taking up carbon dioxide or by just failing to take up easily available oxygen.

In contrast, hearts can take up and circulate whatever liquid they are fed, while, at the same time, remaining properly functional. Take your heart and place it in a vat with orange juice: insofar as it continues to pump at the normal rate, your heart is normally functioning – in spite of the fact that now it's pumping orange juice rather than blood. The proper functioning of the heart is not input dependent.

Our cognitive systems do not work like hearts, they work like our respiratory systems; inputs matter for proper function. Our belief-formation capacities can't take up wishes and form beliefs based on them whilst at the same time remaining properly functional. A cognitive system that processes wishes into beliefs is malfunctioning. If this is so, it follows that the proper functioning of our cognitive systems is input dependent.

Similarly, then, just like in the case of our respiratory systems, we should expect our cognitive systems and belief-formation capacities to malfunction in at least two input-dependent ways: via taking up the wrong kind of inputs (e.g. wishes), but also, and crucially for my purposes here, via *failing to take up easily available good inputs* (i.e. *easily available evidence*).

It is important to note that empirical results also overwhelmingly confirm the hypothesis that the proper functioning of our cognitive systems is input dependent (i.e. that our cognitive systems are malfunctioning if they fail to respond to environmental stimuli). Our levels of neuroplasticity¹¹ – the brain's disposition for neuron-level changes in response to the environment – predict the brain's capacity to take up information from the environment: when a cognitive system displays abnormally low levels of structural neuroplasticity,¹² learning in response to novel stimuli from the environment fails to occur at a normal rate. In turn, abnormally low levels of neuroplasticity, generating low responses to environmental stimuli, predict improper cognitive functioning. But if this is so, the proper functioning of our cognitive systems is input dependent: one way in which they can malfunction is by failing to respond to easily available environmental stimuli.

We now have a straightforward explanation of what goes wrong in cases of resistance to evidence: it is an epistemic incarnation of input-level malfunction. Our cognitive systems fail to take up easily available evidence from the environment. Just like respiratory failure is an instance of our

¹¹ For recent work on neuroplasticity in adults, see, for example, Lovden et al. (2010).

¹² Structural neuroplasticity, very roughly, has to do with the brain's capacity to change its physical structure as we learn new things or form new memories. Functional plasticity is the brain's ability to move functions from a damaged area of the brain to other undamaged areas.

respiratory systems failing to take up normal levels of oxygen from the environment, resistance to evidence is an instance of our cognitive systems failing to respond normally to stimuli from the environment.

6.4 Conclusion

The function of our cognitive systems is to generate knowledge, and justification turns on this function: we are justified to believe just in case our cognitive systems work in the way in which they generate knowledge reliably under normal conditions. This is not the full story, though: the proper functioning of our cognitive systems is not limited to the fair treatment of evidence that we pick up from the world – it extends to picking up the relevant evidence when easily available. When our cognitive systems fail to do so, they are malfunctioning and are in breach of epistemic norms.

In the next chapter, I further unpack REEM. I will not take a stance on what the sufficiency threshold stands for. Views will differ on this, and they will also differ on what fixes the threshold in question – whether it's a purely epistemic affair or whether practical and moral considerations may play a role.¹³ My focus from here onwards will be on how to understand evidence, defeat, and permissible suspension in order to make good on REEM and, in turn, on OTB and the resistance intuition.

¹³ I have done extensive previous work on this issue (e.g. see Simion 2021a).