

*Scepticism as Resistance to Evidence*

The view of evidence, defeat, and suspension put forth here delivers the result that paradigmatic scepticism about knowledge and justification is an instance of resistance to evidence. This chapter argues that this result is correct. In order to do that, I look at extant neo-Moorean responses to purported instances of failure of knowledge closure (Pryor 2004, Williamson 2007) and warrant transmission and argue that they are either too weak – in that they concede too much to the sceptic – or too strong – in that they cannot accommodate the intuition of reasonableness surrounding sceptical arguments. I propose a novel neo-Moorean explanation of the data, relying on my preferred account of defeat and permissible suspension, on which the sceptic is in impermissible suspension but in fulfilment of their contrary-to-duty epistemic obligations.

**11.1 Two Neo-Mooreanisms**

Moore sees his hands in front of him and comes to believe that HANDS: ‘hands exist’ based on his extraordinarily reliable perceptual belief-formation processes. Moore’s belief is warranted, if any beliefs are: Moore is an excellent believer. Indeed, Moore knows that hands exist. In spite of his laudable epistemic ways, Dretske (1971) thinks Moore shouldn’t feel free to do whatever it pleases him to do with this belief, epistemically speaking; in particular, Dretske thinks that, in spite of his warranted belief that HANDS, Moore should refrain from reasoning to some propositions he knows to be entailed by HANDS, such as WORLD: ‘there is an external world’. He thinks that this is an instance of closure failure for knowledge: we don’t always know the stuff that we know our knowledge to entail. In better news, conversely, that’s why the sceptic is wrong to think that my not knowing that I’m not a brain in a vat implies that I don’t know any of the ordinary things I take myself to know.

Wright (2002, 2003, 2004) agrees: Moore shouldn't reason to WORLD from HANDS. However, that's not because closure fails, but because the stronger principle of warrant transmission fails: the problem here, according to Wright, is not that we sometimes fail to know the stuff that we know is entailed by what we know. Rather, the issue is that the warrant Moore has for HANDS fails to transmit to WORLD. Compatibly, though, Moore may still be entitled to believe WORLD on independent grounds. If Moore is entitled to believe HANDS, then perhaps he must also be entitled to believe WORLD. But it doesn't follow that his warrant to believe WORLD is his warrant to believe HANDS. Rather, it may be that Moore needs to be independently entitled to believe WORLD to begin with if he is to be entitled to believe HANDS.

Many philosophers are on board with rejecting at least one of these principles – be it merely warrant transmission or closure as well. At the same time, since closure and warrant transmission constitute a bedrock of our epistemic ways – indeed, they are crucial vehicles for expanding our body of knowledge – one cannot give them up without a working restriction recipe: if closure and/or warrant transmission don't hold unrestrictedly, when do they hold? It is fair to say that the jury is still out on this front, and a satisfactory restriction recipe does not seem to be within easy reach.<sup>1</sup>

That being said, several philosophers take the alternative route of resisting the failure claims altogether and thus fully dismiss the data: according to them, closure and warrant transmission are too important theoretical tools to be abandoned on grounds of misguided intuitions. They reject the intuition that something fishy is going on in Moore's argument and argue that scepticism is just an instance of cognitive malfunction: the sceptic's cognitive system malfunctions in that it fails to get rid of their unjustified sceptical beliefs in favour of the justified Moorean conclusion. I call these people 'radical neo-Mooreans'. Here is Williamson:

Our cognitive immunity system should be able to destroy bad old beliefs, not just prevent the influx of bad new ones. But that ability sometimes becomes indiscriminate, and destroys good beliefs too. (Williamson 2007, 681)

I like radical neo-Mooreanism a lot. The majority reaction to this move, however, is that it is less than fair to the sceptic; indeed, this view

<sup>1</sup> But see Kelp (2019) for my favourite proposal.

(intuitively unfairly) categorises scepticism, without qualification, in the same normative boat with other epistemic malfunctions, such as wishful thinking. It is undeniable, though, that in the case of the sceptic, but not in the case of the wishful thinker, we think that there is something reasonable – even if not quite right – about their resistance to Moore's argument. This intuitive difference cries out for an explanation.

At the other side of the neo-Moorean spectrum, we find concessive neo-Mooreans (e.g. Pryor 2004, 2012); these philosophers accept both closure and transmission in Moorean inferences and try to come up with alternative explanations of the data (i.e. with an alternative account of what is intuitively amiss with Moore's argument). In the next section, I look closer at the concessive neo-Moorean explanation of this datum.

### 11.2 Against Concessive Neo-Mooreanism

According to Jim Pryor (2004), while Moore is right to reason from HANDS to WORLD, he wouldn't be very convincing were he to do so in conversation with a sceptic. The problem behind the intuitive fishiness of Moore's reasoning pattern is pragmatic, not epistemic: it is lack of dialectical force, not lack of warrant transmission, that's triggering the uneasiness intuition. In the cases of alleged failure of closure and/or transmission, warrant transmits, but the argument fails dialectically due to psychological higher-order defeat.<sup>2</sup> The sceptic about WORLD will not be convinced by Moore's argument in its favour from HANDS. Here is Pryor:

For a philosopher with such beliefs [i.e. sceptical beliefs], it'd be epistemically defective to believe things just on the basis of her experiences – even if those experiences are in fact giving her categorical warrant to so believe. (Pryor 2012, 286)

Why would it be thus epistemically defective? According to Pryor, the sceptic's unjustified sceptical beliefs rationally obstruct them from believing based on Moore's argument via psychological defeat. In particular, Pryor thinks that Moore's argument gives the sceptic propositional justification for the conclusion, but it fails to generate doxastic justification due to the psychological defeat generated by the sceptic's previously acquired sceptical beliefs. Since the sceptical beliefs are not justified, according to

<sup>2</sup> To my knowledge, the first to have introduced the category of psychological (or doxastic) defeat is Jennifer Lackey (e.g. 2006, 438). For excellent recent work on defeat, see Brown and Simion (2021).

Pryor, they don't defeat the propositional justification generated by Moore's argument. They do, however, rationally obstruct the sceptic from justifiably believing the conclusion of Moore's argument, and in this they defeat the sceptic's doxastic justification.

The point, then, in a nutshell, is that even though it transmits warrant, the Moorean argument fails to convince the rational sceptic in virtue of the conflict between the Moorean claims and the sceptic's previously held beliefs. The sceptic has propositional justification but does not have doxastic justification for HANDS and WORLD.

In what follows, I will take issue with this claim at several junctures. First and foremost, though, it is worth clarifying what exactly the content of the sceptical beliefs that allegedly do the defeating work here is. I want to start off by noting that it is implausible to think that the sceptical belief at stake in the literature is (or should be) something like non-WORLD: 'the external world does not exist'. After all, what we are talking about – and the philosopher that is worth engaging with – is a reasonable sceptic who, for example, believes in underdetermination (i.e. who thinks that, for all they know, they may well be a brain in a vat), not someone who is anxiously fully confident that they're a brain in a vat. The reasonable sceptic that is worth engaging with thinks that, for all the evidence that they have, there may well be no external world. If so, the reasonable sceptic will, at best, have a 0.5 credence that non-WORLD, or else they will suspend belief on the issue. Not much will hang on this below, but since I am interested in being maximally charitable to concessive neo-Mooreanism, I will, for the most part, discuss the reasonable sceptic rather than the maximally anxious sceptic in what follows. Everything I will say, though, will apply *mutatis mutandis* to the anxious sceptic as well.

Now here is a widely endorsed thesis in philosophy: justification is normative. The following is an attractive way of capturing this thought: one's *phi*-ing is *prima facie* practically, morally, epistemically, etc., justified if and only if one *prima facie* practically, morally, epistemically, etc., permissibly *phis*. Plausibly enough, then, one's belief that *p* is epistemically justified if and only if one epistemically permissibly believes that *p*. Justifiers are considerations that support belief, in that, if all else goes well (i.e. proper basing, no defeat, good processing, etc.), enough justifiers render a belief epistemically permissible.

Where does defeat fit within this picture? Just like justification, defeat is a normative category, in that it affects the permissibility of belief. Unlike justification, however, its function is to counter rather than support believing. If justifiers support belief – they contribute to rendering it

permissible – defeaters contribute to rendering it impermissible. It is plausible, then, to think that defeat is the arch-enemy of justification: if justification is normative with a positive valence – in that it renders belief permissible – (full) defeat is normative with a negative valence, in rendering belief impermissible. In reason terms, if you wish, justifiers are normative reasons for belief, whereas defeaters are normative reasons against believing.

Now let's go back to Pryor's account of what goes on in the exchange between Moore and the sceptic. Recall that, according to Pryor, even though Moore's argument does provide the sceptic with propositional justification, it fails to provide them with doxastic justification, in virtue of their unjustified sceptical beliefs defeating the latter but not the former. As such, according to Pryor, the sceptic's belief that HANDS (and WORLD) based on Moore's argument would be rendered unjustified via defeat.

The problem with this picture is that it's not clear how an unjustified belief can have defeating force to begin with. To be clear, I am not claiming that we do not often resist information that we are presented with because of our previously held unjustified beliefs. Indeed, we often resist information presented to us for bad reasons (e.g. due to wishfully believing that it is not true; think, for instance, of cases of resistance to evidence due to partisanship in virtue of friendship, cases of people in abusive relationships who refuse to acknowledge the abuse, etc.). The question at stake when it comes to defeat, though, is not one concerning the *possibility* of resistance to evidence but of permissibility: since justification and defeat are normative, they can only be instantiated in cases in which permissibility is at stake. Cases of wishful thinking are paradigmatic cases in which the hearer is, to use Pryor's term, 'obscured' from believing information that is presented to them due to their wishes. Clearly, though, wishful thinking cases are impermissibility cases: the hearer should not, as a matter of fact, resist the testimony in question, even though they do. Again, to follow Pryor's terminology, these are cases in which the believer is not 'rationally obscured' from forming said beliefs but merely 'obscured'. Or, to put it in reason terms, their unjustified, wishful thinking-based beliefs are mere *motivating* reasons for resisting testimony but not *normative* reasons.

If all of this is so, the question that arises is: is the sceptic being 'rationally obscured', as Pryor would have it, from adopting a belief based on Moore's testimony by their previously held unjustified sceptical beliefs or rather, just like the wishful thinker, merely 'obscured' from so doing?

Since defeat is a normative category, and since, by Pryor's own stipulation, the sceptic's sceptical beliefs are unjustified, it would seem as though they do not qualify as justification defeaters proper, but rather as mere motivating reasons for resisting Moore's argument. The non-normative cannot defeat the normative: motivating reasons cannot outweigh normative reasons normatively. Just because I wish really hard to steal your purse, it does not follow that it is permissible to steal your purse: my motivating reasons, no matter how strong, in favour of stealing cannot outweigh the normative reasons against stealing, since they don't factor into the overall permissibility calculus to begin with.

Why, then, is it intuitive and, according to Pryor, right to think that, once one has adopted a belief that non- $p$  (or a doubt about whether  $p$ , or a 0.5 credence that non- $p$ ), it would be importantly epistemically defective to adopt a subsequent belief that  $p$ ? Take the following standard case of higher-order defeat: I come to believe that the walls in your studio are white but illuminated by a red light to look red. Subsequently, upon arriving at your studio, it seems problematic for me to adopt the belief 'the wall in front of me is red' based on my corresponding perceptual experience as of a red wall. Why is this so? In particular, why is it that, even if we stipulate that my initial belief that the wall is white and illuminated to look red is unjustified, it would seem that, now that I hold it, I shouldn't just trust my perceptual experience?

Maybe the answer to this question has something to do with the order in which the beliefs have been acquired; that is, maybe a difference in extant doxastic states is an epistemologically significant difference. Indeed, Pryor himself alludes to an answer along these lines. According to him, were the sceptic to believe based on Moore's testimony that HANDS, and thereby WORLD, their belief would be irrational because it would not cohere with their previously held sceptical beliefs. According to Pryor, since irrationality precludes justification, were the sceptic to believe what Moore says, their belief would also be unjustified:

I will count a belief as rational when it's a belief that none of your other beliefs or doubts rationally oppose or rationally obstruct you from believing. [...] A rational commitment is a hypothetical relation between your beliefs; it doesn't 'detach'. That is, you can have a belief in  $P$ , that belief can rationally commit you to believe  $Q$ , and yet you be under no categorical requirement to believe  $Q$ . Suppose you believe Johnny can fly. This belief rationally commits you to the belief that someone can fly. If you're not justified in believing that Johnny can fly, though, you need not have any justification for the further belief. You may even have plenty of evidence and be fully justified in

believing that no one can fly. But your belief that Johnny can fly still rationally commits you to the belief that someone can fly. Given your belief about Johnny, if you refrain from believing that someone can fly, you'll thereby exhibit a rational failing. (Pryor 2004, 363–364)

Since rational failings are incompatible with justification, Pryor takes it that this hypothetical type of normativity that he associates with rationality – of the form ‘if you believe that  $p$ , then you are rationally committed to believing that  $q$ ’ – will affect the permissibility of belief tout court: were the sceptic to believe what Moore tells them, their belief would be irrational – since they are antecedently committed to believing the opposite – and thereby unjustified.

There are two problems with this normative assessment, though. First and foremost, note that there are two ways of resolving cognitive dissonance due to holding two conflicting beliefs  $B_1$  and  $B_2$ : one can either abandon  $B_1$  or abandon  $B_2$ . Coherence doesn't tell us which one we should choose: it merely tells us that one needs to go.<sup>3</sup> There are two ways of proceeding in cases in which one is presented with information  $B_2$  that runs counter to one's extant belief  $B_1$ : one can resist adopting  $B_2$  or, alternatively, one can abandon  $B_1$ . Again, coherence doesn't recommend any particular course of action: it just tells us that we need to choose between them.

One thing that Pryor could reply at this juncture is: time makes a difference, epistemically. The previously held belief takes precedence over the incoming information; this is what explains why the sceptic is rational to resist Moore's argument.

The question that arises, though, is: why should we think that time is of such devastating epistemological significance? Just because the sceptical belief precedes Moore's testimony temporally, why is it that we should think that it also gets normative priority? After all, consider the following pair of cases (adapted from Jessica Brown 2018)<sup>4</sup>:

- Case 1: A reliable testifier A, who knows that  $p$ , asserts that  $p$ . At the very same time as receiving A's testimony, the hearer also receives contrary testimony from another reliable testifier, B, that not- $p$ .
- Case 2: We slightly change Case 1 so that the testimony from B arrives just a bit later than the testimony from A, but for whatever reason the hearer does not form any belief about  $p$  before the testimony from B arrives.

<sup>3</sup> See also Graham and Lyons (2021) for similar points.      <sup>4</sup> See also Goldberg (2021).

In these cases, the evidentiary and doxastic situation is constant: one testimony item for  $p$  and one against  $p$ , and there is no difference in mental states. Clearly, the time difference will not make any epistemic difference: in both Case 1 and Case 2, the hearer has equally strong evidence for and against  $p$ . They should suspend belief. But now consider:

Case 3: This differs from Case 2 only in the following respect: as a result of receiving A's testimony, the hearer forms the belief that  $p$  before receiving B's testimony.

Note that there is no temporal difference between Case 2 and Case 3. As such, even by the lights of the philosopher who believes that time can make an epistemic difference, there should be no difference in epistemic assessment either. But if there is no epistemic difference between Cases 1 and 2, nor any epistemic difference between Cases 2 and 3, it follows that there is no epistemic difference between Cases 1 and 3 either. If so, what the hearer should do in both cases is suspend rather than give priority to the first belief they formed and dismiss the second.

Let's take stock. We have seen that considerations pertaining to coherence cannot explain why we should think that the sceptic is rational to resist Moore's argument: coherence is indifferent between resisting Moore's argument and abandoning the previously held sceptical belief. We have also seen that time does not make an epistemic difference either. If so, just because a belief is antecedently held, it does not follow that it takes epistemic priority. All of this suggests that the sceptic has no epistemic normative reason to give priority to their sceptical belief and thereby resist Moore's argument.

Furthermore, recall that, on Pryor's view, Moore's argument is justification conferring, whereas the sceptical belief is unjustified. If so, there *is* epistemic normative reason for the sceptic to adopt the conclusion of Moore's argument, and there is *no* epistemic normative reason to hold onto the sceptical belief – albeit, of course, the sceptic may well have a merely motivating reason to do so. All in all, it would seem, the sceptic ought (epistemically) to abandon their sceptical belief and adopt the conclusion of Moore's argument. The concessive neo-Moorean solution to the sceptical puzzle is wrong: while Moore's argument may well often fail to convince the sceptic, this is not because it lacks dialectical power, but rather because the sceptic is epistemically impermissibly resisting its conclusion in virtue of their previously held unjustified sceptical beliefs.



### 11.3 A New Radical Neo-Mooreanism

Let's take stock again: we've seen that radical neo-Mooreanism – claiming that the sceptic's resistance to Moore's argument is an instance of epistemic malfunction – is thought by many to fail to offer a fully satisfactory explanation of the datum, in that it places the sceptic in the same boat with wishful thinkers, epistemically speaking. However, intuitively, we find the sceptic to be reasonable, even if wrong, when they resist Moore's inference.

Concessive neo-Mooreanism does better on this front. According to this philosopher, the intuition of epistemic permissibility concerning the sceptic's resistance to Moore's argument is to be explained in terms of psychological defeat: Moore's argument is warrant conferring but dialectically defective. Alas, on closer investigation, this account was shown to run into normative trouble: given that the sceptical belief is unjustified, it remains unclear why the sceptic should favour it over the warranted conclusion of the Moorean argument.

In what follows, I will develop a new neo-Mooreanism. My view falls squarely within the radical neo-Moorean camp, in that it takes transmission to hold in Moorean inferences and finds no flaw – epistemic or dialectical – with Moore's argument. However, as opposed to extant radical neo-Mooreanism, it does predict that there is something epistemically good about the sceptic's doxastic response that sets them apart from believers merely displaying full-on cognitive malfunctions, such as wishful thinking.

Recall that, on the account developed here, evidence consists of facts that are knowledge indicators, in that they enhance closeness to knowledge: it consists of 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 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: 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.

As such, 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. 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. For the same reason,

mere beliefs, as opposed to justified and knowledgeable beliefs, will not be evidence material; they lack the relevant indicator property.

Conversely, defeaters are indicators of ignorance: they are facts that one is in a position to know and that lower one's evidential probability that  $p$  is the case.

Going back to our sceptic: just like the wishful thinker, on this view of evidence and defeat, the sceptic has no epistemic reason to believe their preferred sceptical hypothesis. There are no knowledge indicators available to them to this effect. There are no facts that raise the evidential probability of the sceptical hypotheses within their reach. Furthermore, Moore's assertion that HANDS provides the sceptic with evidence that there are hands, as Moore's testimony to this effect is a knowledge indicator. Also, as the sceptic's sceptical belief is not an ignorance indicator (i.e. it does not lower the relevant evidential probability), it does not qualify as a defeater for HANDS. In this, the sceptic is in double breach of justification-conferring epistemic norms: they have unjustified sceptical beliefs, and they resist knowledge indicators on offer because of them. The sceptic does not have defeaters for HANDS; rather, they have mere motivating reasons to this effect: evidentially irrelevant facts (i.e. the fact that they believe non-WORLD/doubt WORLD) that lead her to unjustifiably reject HANDS.

What is it, then, that explains our intuition of reasonableness in the sceptic case and the lack thereof in the case of the wishful thinker? Recall that, according to the view developed here, the sceptic ought not to hold sceptical beliefs to begin with, ought to come to believe that WORLD based on Moore's argument, and thereby ought to draw the inference to WORLD with Moore and abandon their antecedently held sceptical beliefs. If they fail to do all that, they are in breach of the justification-conferring epistemic norm: their resistance to Moore's argument is epistemically impermissible.

Now, here is, however, a well-known fact about norms, generally speaking: sometimes, when we engage in impermissible actions, this gives rise to contrary-to-duty obligations. Consider the following normative claims:

- (1) It ought to be that John does not break the neighbour's window.
- (2) If John breaks the neighbour's window, it ought to be that he apologises.

(1) is a *primary obligation*, saying what John ought to do unconditionally. In contrast, (2) is a *contrary-to-duty obligation* about (in the context of (1)) what John ought to do conditional on his violating his primary obligation.

(1) is a norm of many sorts: social, prudential, moral, and one of politeness. Should John break the neighbour's window, there would be nothing good about it. That being said, John would be even worse off if, should he break the neighbour's window, he would also fail to go and apologise to the neighbour.

Our functionalist normative schema has the resources needed to explain this datum: input-independent proper functioning – of the type that governs hearts – remains a dimension of functional evaluation in its own right, independently of whether the general proper functioning of the trait in question is input dependent or not. Just like we can ask whether a heart is doing what it's supposed to do with the stuff that it takes up – be it blood or orange juice – we can also ask whether the lungs are doing the stuff that they're supposed to do with the stuff that they have taken up – be it oxygen or carbon dioxide. There's going to be an evaluative difference, then, between two pairs of lungs that are both improperly functioning simpliciter (i.e. in the input-dependent sense, in that they take up the wrong kind of stuff from the environment) in terms of how they process their input gas: are they carrying the input gas through the respiratory system, and subsequently through the lining of the air sacs, to the blood cells? The pair of lungs that do are better than the pair of lungs that don't in that, even though strictly speaking both are malfunctioning overall, the former are at least displaying input-independent proper functioning.

What explains our intuition of reasonableness in the sceptic's case, I claim, is not an epistemic norm simpliciter but rather an epistemic contrary-to-duty imperative: now that the sceptic is in breach of the justification-conferring epistemic norm, short of abandoning their unjustified beliefs, the next best thing for them to do is to embrace the commitments following from their unjustified beliefs and reject the commitments that follow from their negation. The next best thing for the sceptic, now that they believe/have a 0.5 credence that non-WORLD and reject HANDS, both impermissibly, short of abandoning their impermissible beliefs, is to reject whatever follows from HANDS. The sceptic's cognitive system, just like the wishful thinker's and just like lungs taking up carbon dioxide from the environment, is overall malfunctioning on several counts: it takes up improper inputs (the sceptic's sceptical beliefs) and rejects excellent inputs (Moore's testimony that HANDS). That being so, though, the sceptic's cognitive system does something right in terms of input-independent functioning: it processes the (bad) stuff that it has taken up in the right way. The sceptic's cognitive system would be even

worse were they, now that they believe/have a 0.5 credence that non-HANDS, to go ahead and infer that WORLD.

Before I close, I would like to consider a possible objection to my view. So far, I have been assuming, with Pryor and Williamson, that the sceptic's sceptical beliefs/doubts are unjustified. One could worry, though, that my view of evidence might allow for the (reasonable) sceptic to have induction-based evidence for their 0.5 credence that non-WORLD. After all, the sceptic could reason as follows: (1) when I can't tell the difference between pears and apples, I can't come to know that there's an apple in front of me. (2) When I can't tell the difference between John and his twin brother, Tim, I can't come to know that John is in front of me. (3) Therefore, when I can't tell the difference between  $x$  and  $y$ , I can't come to know that  $x$  is the case. (4) I can't tell the difference between WORLD and non-WORLD. (5) Therefore, I don't know that WORLD. In turn, if the sceptic believes that (5), on pain of Moorean paradoxicality, they can't believe that WORLD.

There are two points to consider about this. First, crucially, the envisaged sceptic is wrong, as (1) is notably too strong: I can come to know that there's a pear in front of me in a world where there are no apples, or where apples are extremely rare, even if I can't tell the difference between pears and apples. That being said, of course, (1) may well be justified inductively, which would lead to (5) being justified inductively. Second, though, note that Moorean paradoxicality, just like incoherence, tells us nothing about which of the two beliefs should be abandoned: it merely predicts that one needs to go. Why think WORLD needs to go rather than (5)? Furthermore, notice that in everyday testimonial cases it's the previously held ignorance belief that should be abandoned: I believe I don't know whether you are thirty-two years old, you tell me that you are thirty-two years old, and I thereby come to know that you are thirty-two years old and abandon my belief that I don't know that you are thirty-two years old. That's how it normally goes.

Here is a last attempt: maybe the sceptic's inductively justified belief that they can't tell the difference between WORLD and non-WORLD acts as an undercutting defeater for Moore's testimony that HANDS? This could work. The problem, though, is that undercutting defeaters need to exhibit particular strength properties in order to successfully undercut. For instance, my three-year-old's testimony that Dretske is wrong about closure failure because he took a hallucinogenic drug before writing his paper 'Epistemic Operators' will not successfully undercut my evidence that closure fails sourced in Dretske's paper. Why not? My three-year-old

is just not a very reliable testifier on the issue – not reliable enough to undercut Dretske's written testimony at any rate. The testimony from my three-year-old does not lower my evidential probability conditional on Dretske's testimony that closure fails. If so, what would need to happen in the case of the sceptic for their induction-based sceptical belief to undercut Moore's testimony would be that the former is weighty enough, epistemically. Why, though, should we think that the sceptic's induction has such devastating epistemic effects against Moore's testimony? Also, recall that the inductive argument only warrants the reasonable sceptical belief 'I don't know that WORLD', not the anxious sceptical belief that 'non-WORLD'. Of course, though, the former is much weaker than the latter and thus has much less defeat power.<sup>5</sup>

#### 11.4 Conclusion

This chapter has developed a novel, functionalist variety of radical neo-Mooreanism. I have argued with Williamson that, just like the wishful thinker, the sceptic is displaying epistemic malfunction in rejecting Moore's testimony. On my account, that is because their cognitive processes fail to pick up knowledge indicators. I have also shown, however, that the intuition that there's something reasonable about the sceptic who resists going through Moore's inference is right: the sceptic is in compliance with a contrary-to-duty obligation akin to input-independent well-functioning.

To be clear: this account does not make any concessions to the sceptic in terms of justification-conferring epistemic norms (i.e. primary epistemic obligations): no justification for sceptical beliefs, nor any defeat against Moore's testimony, is instantiated in the context. The account merely explains why we find the sceptic reasonable (albeit wrong) to resist Moore's inference from HANDS to WORLD: they are in compliance with their contrary-to-duty epistemic obligations. Now that they have broken the window, as it were, the sceptic might as well go ahead and apologise.

<sup>5</sup> Thanks to Chris Kelp for pressing me on this.