

Defeaters as Ignorance Indicators

This chapter puts forth and defends a novel view of defeat, and it shows that it is superior to its competition in that it can account for the epistemic impermissibility of defeat resistance cases and normative defeat cases, as well as for the effect ignored defeat has on doxastic justification. On this account, defeaters are ignorance indicators¹: facts that one is in a position to know and that reduce one's evidential probability that p . Furthermore, I also put forth a novel account of the normativity at work in cases of normative defeat and negligent inquiry and evidence gathering.

8.1 The Nature and Theoretical Importance of Defeat

The notion of defeat is central to epistemology, practical reasoning, and ethics. Within epistemology, it is standardly assumed that a subject who knows that p or justifiably believes that p can lose this knowledge or justified belief by acquiring a so-called defeater, whether this is evidence

¹ I have developed, defended, and presented my account of epistemic reasons, evidence, and defeat as knowledge/ignorance indicators starting back in 2018. Dutant and Littlejohn (2021) also call defeaters 'ignorance indicators', but their account is spelled out in very different terms, so any affinity is merely terminological. On their account, defeaters consist in evidence that one is not in a position to know. Gibbons (2013) and Kelp (2023) also develop and defend accounts along these lines. There are two main problems for accounts like these. The first is structural: these accounts are epistemically second order, in that defeaters are evidence that some epistemic status is missing. But for an agent to have this evidence, they need to be able to process the relevant content; many agents that can undergo defeat are not sophisticated enough to have the relevant contents, however. The second problem parallels a wrong-kind-of-reasons problem: matters that are intuitively irrelevant to justification and defeat can be evidence for or against one being in a position to know. Here is a case from Jonathan Jenkins-Ichikawa (personal conversation): say that your grandfather is not feeling well and you are searching for the thermometer to check whether he has a fever. Now, finding the thermometer is evidence that you are in a position to know that your grandfather has a fever, and, indeed, it is evidence that you will come to know that he does. However, clearly, it does not affect the justification of your corresponding belief. Conversely, not finding the thermometer is evidence that you are not in a position to know that he does have a fever, but it surely does not defeat whatever justification you might have had to believe that he does have a fever.

that not-*p*, evidence that the process which produced their belief is unreliable, or evidence that they have probably misevaluated their evidence. Within ethics and practical reasoning, it is widely accepted that a subject may initially have a reason to do something, although this reason is later defeated by their acquisition of further information.

Investigations into the nature and normativity of defeat come with high theoretical stakes. The notion of defeat has been central to a wide range of different philosophical debates, including, but not limited to:

- (1) The nature of justification and knowledge: since knowledge and justification are taken by many to be defeasible, the extent to which one account or another of the nature of knowledge/justification is able to account for/accommodate defeat constitutes an important ground for assessing its theoretical credentials.
- (2) Internalism versus externalism: several epistemologists worry that epistemic externalism has a hard time accommodating psychological defeat; at the same time, conversely, if justification supervenes on mental states alone, as per internalism, it seems mysterious that it could be defeated by normative defeaters lying outside of the cogniser's ken.
- (3) Epistemic norms and reasons: for accommodating the phenomenon of defeat, debates on epistemic norms and epistemic reasons owe us, at a minimum, an account of epistemic normative overriding, as well as an account of reasons against belief.
- (4) Evidence and higher-order evidence: since evidence is widely taken to be defeasible, a plausible account of evidence should come with a corresponding plausible account of defeat. For instance, one important desideratum on any such account is that it explains the defeating power of higher-order evidence, namely of evidence that one's first-order beliefs are the output of a flawed process.²
- (5) Closure and transmission: one popular solution to alleged failures of closure principles for knowledge and transmission principles for warrant is known as 'the defeat solution'. Roughly, the thought goes, closure and transmission hold *prima facie*, and the intuitions of failure are to be explained in terms of psychological defeat. This solution, of course, hangs on the assumption that there is such a

² Some philosophers believe that knowledge is not defeasible by higher-order evidence (e.g. Lasonen-Aarnio 2017, Williamson (forthcoming)): the picture I put forth does not rest on the denial of this claim. Even sceptics of knowledge defeat by higher-order evidence are not full-on sceptics about the defeating power of higher-order evidence.

thing to begin with (i.e. that psychological defeat is a genuine epistemic category).

- (6) Disagreement: one way to characterise the debate between conciliatory and steadfast views of disagreement is as centred around the question: can the testimony of one's peer carry defeating power? Steadfastism answers 'no', conciliationism answers 'yes'. The correct account of the nature of defeat can help settle the issue.
- (7) Reductionism versus anti-reductionism about testimony: say that a suspect for murder *S* tells you that she did not do it. According to both of the main views in the epistemology of testimony, you are not justified to believe *S*. According to reductionism, that's because you always need positive, non-testimonial reasons to believe what you are being told. In contrast, according to anti-reductionism, you are *prima facie* justified to believe *S*, but your justification is defeated. The correct account of the nature of defeat will likely go a long way towards settling this issue.

Given these high theoretical stakes, investigations into the nature and normativity of defeat carry significant philosophical weight. Unfortunately, not many systematic, full accounts of defeat have been put forth on the market (although see Brown and Simion (2021) for the first full volume on defeat and Kelp (2023) for the first book-length treatment). In what follows, I will look at the classical accounts of defeat on the market – one internalist evidentialist, coming from John Pollock, and one externalist reliabilist, championed by Alvin Goldman – to help situate my account in the extant landscape.³

8.1.1 *Traditional Evidentialism About Defeat*

The first and what is now considered the classic view on the nature of defeat in epistemology is due to Pollock (1986). According to this view, *d* is a defeater of *e*'s support for *p* for *S* if and only if (1) *e* is a reason to believe *p* for *S* and (2) *e&d* is not a reason to believe *p* for *S* (henceforth 'Pollock's view').

The account has a lot going for it: it nicely promises to cut across normative domains in virtue of being framed in terms of reasons; after all, epistemologists hardly enjoy an exclusivity on reasons. With Pollock's view

³ For a comprehensive overview of extant accounts of defeat in the literature, see Simion (forthcoming).

in play, it is easy to see how we could generalise it to cover different targets (e.g. actions) and types (e.g. moral, prudential) of normativity. Second, Pollock's view makes good on the intuitive thought that defeaters are actualisers of the possibility of a positive normative status to be overridden or undercut; what the view says, in a nutshell, is that defeaters are the kind of things that render a permissible belief impermissible.

For our purposes here, there are two important limitations to Pollock's view: first, it does not account for partial defeat – nor is it trivial to see how it could be extended to do so. Since many cases of defeat resistance will be cases of partial defeat, the view will not deliver the needed theoretical resources for the data I am trying to explain. Second, the account remains silent on the nature of reasons and, most importantly, on what it is for something to be a reason *for S* to believe: however, in order to understand the impermissibility of defeat resistance, it is crucial to understand both of these things. We need to know what reasons are and which reasons are reasons *for S* to believe, since this is essential to understanding the impermissibility of *S*'s resistance to defeat.

One way to spell out Pollock's view that suggests itself, given his evidentialist leanings, is a traditional, seemings-based recipe: on this account, reasons for *S* to believe are *S*'s relevant seemings. Of course, as we have seen in Chapter 2, a view like this will get us into trouble with resistance cases rather rapidly: on the necessity direction, recall only the very sexist George, who zones out whenever a woman speaks to him. This guy doesn't host any relevant seemings – intuitively, however, his beliefs are defeated by women's testimony. Against the sufficiency direction, notably, cases of cognitive penetration will create trouble for a seemings-based defeat account (e.g. see Lyons 2011): the fact that it seems to me – due to sexist bias – that women don't know what they're talking about is not enough to defeat my justification to believe their testimony.

8.1.2 Defeaters as Reliable Processes

Reliabilist theories of justification have been extremely popular in the last three decades and come in a variety of forms, but the gist of the view is that a belief is justified if and only if it is formed via a (normally) reliable procedure or ability. Reliabilism is a theory of *prima facie* justification. As such, in line with normative theories in general, it needs a theory of defeat in order to hold water. The standard reliabilist account of defeat comes from Alvin Goldman:

The alternative reliable process account (ARP): S's belief is defeated iff there are reliable (or conditionally reliable) belief-forming processes available to S such that, if S had used those processes in addition to the process actually used, S wouldn't have held the belief in question (Goldman 1979).

One can see how ARP is an elegant reliabilist translation of the Pollockian thought that defeat is the kind of normative entity that, when taken in conjunction with the extant epistemic support for the relevant belief, fails to render it justified.

Bob Beddor (2015) is the *locus classicus* for criticism of ARP; if Beddor is right, ARP is both too weak and too strong. Against ARP's sufficiency direction, Beddor offers the following case:

Thinking About Unger: Harry sees a tree in front of him at *t*. Consequently, he comes to believe the proposition TREE: ⟨There is a tree in front of me⟩ at *t*. Now, Harry happens to be very good at forming beliefs about what Peter Unger's 1975 time-slice would advise one to believe in any situation. Call this cognitive process his Unger Predictor [...]. What's more, [...] whenever it occurs to Harry that Unger would advise him (Harry) to suspend judgement about *p*, this causes Harry to [...] suspend judgement about *p*. So if Harry had used his Unger Predictor, he would have come to [...] suspend judgement regarding TREE. (Beddor 2015, 152)

What this case shows is that ARP is too weak normatively: contra ARP, for my belief that *p* to be defeated, it is not enough that I would change my mind about *p* in a counterfactual world due to employing some reliable process. Just because I would change my mind in world *W*, it does not follow that I *should* change my mind in world *W*: defeat is a normative notion.

More importantly for our purposes here, however, Beddor's case against ARP's necessity direction is, indeed, a paradigmatic case of evidence resistance. Here it is:

Job Opening: Masha tells Clarence that her department will have a job opening in the fall. Clarence believes Masha; assuming that Masha is usually reliable, Clarence's belief counts as *prima facie* justified. Sometime later, Clarence speaks with the head of Masha's department, Victor, who informs him that the job search was cancelled due to budget constraints. Now suppose that Clarence harbours a deep-seated hatred of Victor that causes him to disbelieve everything that Victor says; what's more, no amount of rational reflection would rid Clarence of this inveterate distrust. Consequently, he continues to believe that there will be a job opening in the autumn.

This case shows that ARP is also too strong: just because, in all counterfactual worlds, I would irrationally and stubbornly hold on to my belief, it does not follow that I *should* do so. Once again, ARP is not normative enough to do the job it is supposed to do. Going back to our purposes here, Job Opening is a paradigmatic case of testimonial evidence resistance: it shows that ARP's normative weakness results in difficulties for accounting for resistance to defeat.

Along similar lines, recall the case of Professor Racist:

Case #6. Misdirected Attention: Professor Racist is teaching college-level maths. He believes people of colour are less intelligent than white people. As a result, whenever he asks a question, his attention automatically goes to the white students, such that he doesn't even notice the Black students who raise their hands. As a result, he believes Black students are not very active in class.

To see the problem that this case poses for ARP, let us again ramp up the epistemically problematic aspects of the case and stipulate that Professor Racist is not only racist but also dogmatic (indeed, this assumption should not be very hard to make, since these two epistemic vices tend to be encountered together): even if he had seen the Black students raising their hands, he would have still strongly believed that they're not very active in class. If so, ARP predicts that there is no defeat at stake in this case: after all, there is now no alternative reliable process that is such that, had Professor Racist used it, he would have abandoned his belief that Black students are not very active in class. Not only is ARP then not predicting defeat in this case, but it predicts that Professor Racist is epistemically better off in this version of the case than in the original version. Since being racist and dogmatic is epistemically worse than being only racist, ARP remains unsatisfactory.⁴

Note, finally, that we can make the same stipulation of dogmatism in all of the other resistance cases as well, with the same unsatisfactory result.

⁴ See Beddor (2015) for an excellent and comprehensive critique of ARP, Graham and Lyons (2021) for a rejoinder on ARP's behalf and Simion (forthcoming) for criticism of both. Beddor (2021) proposes a reasons-first reliabilism that unpacks reasons as inputs to reliable processes. Beddor's view will also struggle with resistance cases. After all, reliable processes are not infallible processes: they can fail on occasion, either by not taking up the right kind of stuff or by taking up the wrong kind of stuff. If so, the account does not have the resources to explain why testimony from a woman is a reason, while sexist beliefs are not.

As such, process reliabilists still owe us an explanation of what is going wrong in cases of defeat resistance.

The account I will develop next, somewhat unsurprisingly, complements the account of evidence as knowledge indicators developed in the previous chapter. On the view I will defend, conversely, defeaters are ignorance indicators. It is interesting, I think, to consider the view against its historical evidentialist and reliabilist predecessors: like the evidentialist, my account takes evidence and defeaters to have normative strength independently of whether they are being taken up via particular types of processes or abilities. Like the reliabilist, however, the account bottoms out in processes and abilities: a fact only constitutes a defeater if one is *in a position* to take it up via one's cognitive capacities.

8.2 Defeaters as Ignorance Indicators

In my view, 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:

Defeaters as ignorance indicators: A fact d is a defeater for S 's evidence e for p iff S is in an position to know d and S 's evidential probability that p conditional on $e \& d$ is lower than S 's evidential probability that p conditional on e .

Or, slightly more formally:

Defeaters as ignorance indicators: A fact d is a defeater for S 's evidence e for p iff S is in a position to know d , and $P(p/e \& d) < P(p/e)$.

Recall also that, on my view, 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. It is easy to see that the view of defeat defended here nicely predicts that the justification of some occurrent beliefs hosted by the characters in several of resistance cases is defeated by the presence of ignorance indicators. Take Bill, Dump's supporter: since the information coming from several sources around him is such that he is in a position to know it, and it lowers the evidential probability that Dump is a good president, Bill should lower his confidence that Dump is a good president. Similarly, Mary should lower her confidence that her husband is just making friends, and Professor Racist should not believe, as he does, that Black students are not very active in class.

In my view, rebutting and undercutting defeaters share one and the same central epistemic normative property: they are evidential probability decreaseers. What differs is the mechanism by which they achieve this effect: rebutters lower one's evidential probability for p by raising one's evidential probability for not- p . In contrast, undercutters reduce the degree of confirmation that a particular piece of evidence e confers on p (see also Kotzen (2019) for a detailed formal treatment along these lines). This comes in stark contrast to literature that gives different treatment to first- and higher-order evidence or rebutting and undercutting defeat. I think mine is the right result, and we should, all else equal, prefer this unified treatment on grounds having to do with theoretical adequacy. Here are some quick reasons why scepticism about the defeating power of higher-order evidence does not work: one's evidence comes with a having relation and an evidential support relation (also known as degree of confirmation: how much a piece of evidence probabilifies p). Plausibly, one's confidence in p should match the degree of confirmation that one's evidence offers to p . Higher-order evidence/undercutting defeat works via raising/lowering the degree of confirmation that first-order evidence provides to p . Now, here is a case of higher-order evidence that increases the degree of confirmation of the first-order evidence: I believe p based on my neighbour George's testimony (alternatively, I have, e.g., 0.8 credence that p). Mary tells me that George is the top expert in the world on the matter. Her testimony is evidence that q : 'George's testimony gives very high support to p '. In probabilistic terms: if George's original testimony probabilifies p to x , Mary's testimony translates roughly as 'George's testimony probabilifies p to y & $y > x$ ' (how high y is will depend on Mary's epistemic credentials). So now that Mary has spoken, I am in a position to know something along the lines of r : 'I have a 0.8 credence that p based on George's testimony and the probability that p conditional on George's testimony is 0.9'. Intuitively, I should revise to 0.9.

Moving on to a corresponding case of undercutting defeat: I believe p (indeed, I know p) based on my neighbour George's testimony (alternatively, I host, e.g., 0.8 credence that p). Mary tells me that George is a well-known liar on p -issues. Her testimony is evidence that q : 'George's testimony gives lower support to p '. George's original testimony probabilifies p to x ; Mary's testimony, just like before, amounts to: 'George's testimony probabilifies p to y & $y < x$ ' (it will depend on Mary's epistemic credentials how low y is). Just like before, I am in a position to know that r : 'I have a 0.8 credence that p based on George's testimony and the probability that p conditional on George's testimony is 0.4'. Intuitively,

I should revise to 0.4. Scepticism about the defeating power of higher-order evidence is wrong, and defeat affords unified treatment.

Going back to our central cases of evidence resistance – again, crucially, real-world, high-stakes cases of climate change denial and vaccine scepticism will sometimes be diagnosed by this account of evidence and defeat 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.

As previously shown, however, 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 defeat: if reliable testifiers in one's community testify against p : 'climate change is 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' defeaters (i.e. evidential probability decreaseers). Note, however, that these cases – cases of justified evidence rejection in virtue of misleading defeat – will be fairly specific cases epistemically: for example, cases in which the cogniser has more reliable (although misleading) testimony that not- p than evidence that p , or cases in which the cogniser has overwhelming undercutting defeaters (e.g. based on reliable although misleading testimony) for the source of p . While this may happen in fairly isolated communities, the more one has access to evidence for p , the less justified their evidence rejection will be.

Here is a question that arises for the account put forth: why think that there is such a thing as an ought governing our belief formation to take up defeat? After all, on the account proposed, defeaters are ignorance indicators (i.e. facts that lower evidential probability for one). Why should a system with the function of generating knowledge be under an obligation⁵ to take up ignorance indicators, which, by stipulation, decrease closeness to knowledge?

Two things about this: first, note that defeaters are not merely ignorance indicators for p , but also either knowledge indicators for not- p (rebutting defeaters) or knowledge indicators for 'piece of evidence e does not confirm/offers weaker confirmation for p ' (undercutters). This, in turn,

⁵ Recall that I follow classic deontic logic in using 'ought' and 'obligation' interchangeably. Philosophers who think of obligations as being thicker than mere oughts, or as governing only human beings but not systems, artifacts, etc., should read 'obligation' as 'ought' throughout.

affords them the same normative explanation as that offered for garden variety evidence: defeaters exercise normative pressure on our cognitive systems in virtue of them being knowledge indicators.

Second, and furthermore, I think that there is more to the normative pressure of defeat than this merely general evidential normative pressure; in particular, I think that the normative pressure of defeaters for my belief that p can also be explained in a p -centric fashion. Here it goes: consider, for starters, a case in which I know that p at t_1 , and (misleading) evidence that not- p becomes available to me at t_2 . In a situation like this, crucially, whether I take up the defeating evidence or not, my knowledge that p is defeated (short of my knowledge constituting defeat defeat): I am now left, at best (depending on the relevant evidential weights), with a somewhat – but not fully – justified true belief that p . My full belief has thereby been rendered impermissible: I now hold a doxastic attitude that is stronger than what the evidence affords. Note, also, that this is so even if what I started with is not knowledge, but rather, for example, a justified credence of 0.6. Now that defeat is available in my environment, whether I take it up or not, my 0.6 credence is no longer permissible, since my level of justification has been lowered by defeat.⁶ What is it that my cognitive system ought to do now that it's hosting an impermissible doxastic attitude? The answer that suggests itself is: abandon it, either altogether (if defeat is full defeat) or in favour of the weaker attitude that remains supported by the evidence.

Note, though, that abandoning the (now) impermissible doxastic attitude in conjunction with not taking up the relevant defeaters seems irrational. If I, at the same time, (mistakenly) continue to take my evidence to support a 0.6 credence and adjust to a 0.5 credence, something has gone amiss, rationality-wise. In sum, it would seem as though, if I don't take up the relevant defeat from the environment, I am faced with a normative dilemma: either I hold on to an impermissible doxastic attitude that is no longer supported by my evidence or I hold a novel doxastic attitude that enjoys propositional justification but not doxastic justification, for lack of proper basing. Grasping neither of these horns is knowledge conducive. If so, it would seem, the only path left for knowledge conduciveness is

⁶ For readers who believe that knowledge is not defeasible by higher-order evidence, while other epistemic states may be – in line with, for example, Lasonen-Aarnio (2014) and Williamson (forthcomingb): the picture I put forth does not rest on the denial of this claim. That is because it is compatible with my picture that, for example, knowledge always constitutes defeat defeat.

taking up the relevant defeaters and adjusting my doxastic attitude in light of them.⁷

8.3 External Defeat and Norms of Evidence Gathering

Consider a doctor, X, who believes that p but missed a recent development in the field that ($q =$) the data that have been taken to support p don't really do so. X does not justifiably believe that p : q is a defeater for this belief, which undermines X's justification for believing p . Now contrast X with a layperson, Y, who had been told by their doctor that p and still believes that p . Despite the fact that there has been a recent development in the field of medicine, Y's belief that p continues to be justified. In particular, q does not undermine Y's justification for their belief that p . Since the central difference between X and Y is that X occupies a certain social role (i.e. they are a doctor), there is reason to think that social roles can be sources of defeat via giving subjects reasons to inquire (Goldberg 2018).

Note, first, that so far we have looked only at epistemic functions in individual agents. And while epistemic functions may arise in individual agents, they also arise in broader social systems. It is precisely this idea that will be of central importance in what follows.⁸

To begin with, I take social systems to be systems that feature multiple agents who are connected to one another in at least some ways. The social roles we are interested in are properties of agents in social systems. Being a doctor, teacher, parent, etc., is a property of an agent in a social system.⁹

One interesting feature of social roles is that many of them have constitutive functional properties in that what it is to be an X (doctor, teacher, baker, fireperson, etc.) is to have the particular function in question (to treat ill people, to teach people stuff, to make baked goods, to put out fires, etc.). To see this, note that we cannot even fully understand the roles in question without understanding their functions: to fully understand what a doctor is, you have to understand what the function of a doctor is (i.e. to

⁷ For a discussion of the epistemic value of coherence, see Chapter 11.

⁸ See Kelp and Simion (2021b) for a full treatment of normative defeat.

⁹ Note that the existence of agents who occupy social roles need not imply the existence of other agents. We might want to allow that a doctor continues to be a doctor even if they are the only human being left alive. However, this does imply the existence of other agents at least at some point in the past. If there had only ever been one human being around, this person could not have been a doctor. At best, they may have engaged in healing practices, but they couldn't have been a doctor. We take this to be independently plausible. After all, social roles are just that: *social* roles.

treat ill people). It is easy enough to see that the same is true of a whole host of other social roles, including teachers, bakers, firepeople, among many others.

Given that many social roles have constitutive functional properties, the prospects for an analysis of the epistemic norms constituting these social roles already look bright. The question that we need to consider is: are the functions constitutive of these social roles generating any constitutive epistemic norms of proper functioning? If the answer to this question is yes, then the route to an account of epistemic norms constitutive of social roles is a going to be a short one. By the same token, an account of how social roles may lead to defeat may come into view.

To get an idea of how this might be, let's take another look at the case of the doctor. Note that having an up-to-date understanding of their field is part of the normal functioning of doctors in the social system that we occupy. More specifically, it is part of such normal functioning that doctors engage in inquiries into recent developments in the field, as a result of which they maintain an up-to-date understanding of the field and thereby know how to treat people. In fact, that doctors maintain an up-to-date understanding is a key element in the feedback loop that explains the continued existence of this important social role in the social system that we occupy¹⁰: doctors' understanding informs their treatment practices, and the fact that it is kept up to date enhances their success rate of these treatments, which in turn explains why the social role of doctor continues to exist in our social system. But since maintaining an up-to-date understanding of the field is part of the normal functioning of doctors in our social system, we get the by now familiar normative import. It is thereby part of such proper functioning. This, in turn, means that we get a norm that doctors violate if they fail to maintain an up-to-date understanding of their field and, by the same token, if they fail to engage in the inquiries needed to do so. And, of course, it is easy enough to see that the same holds, *mutatis mutandis*, for many other social roles, including teachers,

¹⁰ These are, of course, not essential features of doctors, but merely constitutive features: doctors remain doctors even when they are in breach of some of the norms constitutive of their social role. This is a general feature of functionalist normativity: hearts remain hearts even when they malfunction. It is also a general feature of constitutive normativity more generally: one can break several constitutive norms of a game or language and still count as playing the game/speaking the language. Note, though, that too widespread a breach of constitutive norms with maximal systematicity will lead to discontinuing the constituted activity: if I only utter 'kakakakaka', I don't count as speaking English. Similarly, someone who doesn't have a medical degree, lacks any understanding of medicine, etc., will not count as being a doctor due to their being in too widespread a breach of the constitutive norms of the role.

lawyers, academics, and so on. In turn, since the norms in question are generated by the constitutive functions of these social roles, they will be constitutive norms.

It turns out, then, that we can explain normative defeat as a breach of a constitutive norm, sourced in the constitutive function of these social roles. Since our doctor *X* is a practicing doctor, they occupy the social role of doctor. As a result, they violate an epistemic norm associated with proper functioning for this role when they fail to maintain an up-to-date understanding of their field (e.g. by missing the research that indicates that *q*). In this way, it is epistemically proper for them to believe that *q*. And since *q* is a reason against believing that *p*, we get the desired result that *X*'s justification for believing *p* is defeated.

How does this account of normative defeat map onto my general account of defeaters? Recall the proposed view: a fact *d* is a defeater for *S*'s evidence *e* for *p* iff *S* is in an position to know *d* and *S*'s evidential probability that *p* conditional on *e* & *d* is lower than *S*'s evidential probability that *p* conditional on *e*. Recall also that, on my account, being in a position to know is spelled out in terms of availability, as restricted by the type of cogniser instantiated. What I want to suggest is that, in cases like that above, the social role individuates the type of cogniser at stake and the availability conditions follow the corresponding constitutive epistemic oughts: since cognisers like our doctor should be aware of recent developments in their field, the account will predict that these are available to them in the relevant sense – of course, with reasonable qualitative, quantitative, and environmental restrictions.

But won't this account suffer from a problem parallel to Goldberg's view? Recall that Goldberg wanted to explain the evidence one should have had in terms of social expectations. Recall, also, that we said that epistemically illegitimate (albeit reliable) social expectations make problems for Goldberg's account: if epistemic normativity is encroached by social normativity, reliable but epistemically problematic social expectations cannot be further explained in epistemic normative terms.

Won't my account have the same problem, in virtue of appealing to social roles? Can't there be social roles that are functionally constituted by norms that are bad, epistemically? Consider, for instance, the social role 'judge' in a judicial system where discrimination based on race is written into the laws of the land: isn't my account going to deliver the result that judges shouldn't update based on the testimony of, for example, Black testifiers?

It will not. To see this, note that one important advantage that my account has over Goldberg's is that epistemic normativity is not

encroached upon by social normativity: the epistemic remains an independent normative domain with its own independent evaluative structure. On my view, some genuine epistemic norms – associated with promoting epistemic values, such as knowledge – constitute social roles. Compatibly, norms constituting social roles that are bad epistemically, in that they conflict with norms sourced in the proper functioning of our cognitive system – such as ‘don’t believe Black testifiers!’ – are not epistemic norms: they are mere (bad) social norms with epistemic content.

8.4 Conclusion

This chapter developed an account of defeat that builds nicely on the account of evidence developed in the previous chapter. On this view, defeaters are ignorance indicators: they are facts that one is in a position to know and that decrease one’s evidential probability. What differs is the mechanism by which they achieve this effect: rebutters lower one’s evidential probability for p by raising one’s evidential probability for not- p . In contrast, undercutters reduce the degree of confirmation that a particular piece of evidence e confers on p . Finally, the chapter developed a novel, functionalist account of normative defeat and the impermissibility of negligence in evidence gathering.