ORIGINAL ARTICLE

Defeat and proficiencies

Christoph Kelp

COGITO Epistemology Research Centre, University of Glasgow, Glasgow, UK

Correspondence
COGITO Epistemology Research Centre, University of Glasgow, Glasgow, UK
Email: christoph.kelp@glasgow.ac.uk

Abstract

Virtue epistemology is the view that beliefs are attempts at truth (or perhaps knowledge) and, as a result, can be assessed as successful, competent, and apt. Moreover, virtue epistemology identifies central epistemic properties with normative properties of beliefs as attempts. In particular, knowledge is apt belief and justified belief is competent belief. This paper develops a systematic virtue epistemological account of defeat (of justification/competence). I provide reason to think that defeat occurs not only for beliefs but for attempts more general. The key constructive idea is that defeaters are evidence that attempting (in a certain way) isn't successful and that defeaters defeat the competence of an attempt when one stands in a certain normative relation to the defeater. I argue that while this account handles paradigm cases of defeat both within epistemology and beyond nicely, cases of external (sometimes also 'normative' or 'propositional') defeat continue to cause trouble. To handle these cases, I develop a distinctively functionalist version of virtue epistemology. This functionalist version of virtue epistemology allows me to countenance proficiencies, that is, roughly, abilities that have the function to produce successes under certain conditions. It is the normative import of proficiencies that delivers the normative relation that serves to explain defeat in cases

Philosophical Issues. 2022;32:82–103.

of external defeat. In this way, the functionalist version of virtue epistemology ushers the way towards a satisfactory account even of external defeat.

1 | INTRODUCTION

Virtue epistemology is the view that knowing and believing justifiably have to do with believing from epistemic ability. It is one of the most popular accounts of knowledge and justified belief in the literature. At the same time, one of the most important properties of justified belief is that it is defeasible. Very roughly, what this means is that we can lose justification for our beliefs in light of new information. Given just how popular virtue epistemology is and given how important the phenomenon of defeat is in the theory of justification, it is surprising just how little discussion of defeat there is in the literature on virtue epistemology.[1]

The central aim of this paper is to supply this lack and to develop a systematic virtue epistemological account of defeat. To achieve this aim, I will first say a few words about defeat, to get the phenomenon into clear view, and I will distinguish between two kinds of defeat (internal and external) that any adequate account of defeat must accommodate (Section 2). Section 3 introduces virtue epistemology and argues that standard versions of virtue epistemology threaten to be unable to accommodate the phenomenon of defeat altogether. Section 4 develops a solution to this problem. In Section 5, I develop the backbone of a more substantive virtue epistemological account of defeat and show how it can handle cases of internal defeat.

The remainder of this paper aims to show how we can make room for external defeat. To this end, I first argue that cases of external defeat mean trouble for standard versions of virtue epistemology (Section 6). To make room for external defeat, Section 7–8 offer a distinctively functionalist version of virtue epistemology. Section 7 develops a functionalist account of abilities and draws out its normative implications. Section 8 introduces the notion of a proficiency, that is, roughly, an ability that in addition has the function of producing successful attempts under certain trigger conditions and argues that some epistemic abilities are also proficiencies. Finally, Section 9 argues that that the functionalist version of virtue epistemology does allow us to make room for external defeat in our epistemology.

2 | DEFEAT: INTERNAL AND EXTERNAL

I characterised defeat in terms of the loss of justification in light of new information. To get a better handle on defeat, let's look at a couple of paradigm cases:

You are preparing to go and see a show by your favourite band. You have bought tickets months ago and, finally, the big day has arrived. You have double and triple checked the dates and times. You believe that the show is happening today. Clearly, your belief is justified. Just as you are about to leave, you are notified that the show is cancelled because one of the band members has fallen ill.

You are looking at a surface in front of you, which looks red. Based on this, you form a belief that the surface is indeed red. This is yet another case in which your belief is justified. Shortly

after, I tell you that the surface is illuminated by red light and, as a result, would look red to you even if it were white.

These paradigm cases of defeat. In both cases, justification is lost in light of new information. In the cancellation case, initially, you justifiably believe that the show is happening today. However, your justification is lost in light of the new information is that the show was cancelled. Likewise, in the red-light case, initially, you justifiably believe that the surface is red. However, your justification is lost in light of the new information that the surface is illuminated by red light.

I am hoping that the characterisation of defeat as loss of justification in terms of new information is not only useful but also theoretically lightweight in that it is acceptable to most contributors to the debate on defeat. In fact, I think that the lightweight characterisation of defeat can be fleshed out a little further. When defeat occurs, what happens is that the new information leads to the loss of justification *by providing reason against holding certain beliefs*.[2] In the cancellation case, the information that the show has been cancelled due to illness provides a reason against believing that it is happening tonight. And in the red-light case, the information that the surface is illuminated by red light is a reason against believing that the surface is red.[3]

Now, in both of our toy cases, you come to know some fact that constitutes the defeater for your belief. In the cancellation case, you come to know that the gig is cancelled. In the red-light case, you come to know that you ingested a drug that messes with your colour vision. It is widely acknowledged that facts that we know can defeat our justification. In what follows, I will refer to cases in which an agent's justification is undermined as a result of a defeater that they have psychologically registered[4] as cases of 'internal defeat'. However, there is also a question over whether facts that we do not know (or even believe) can defeat our justification as well. In what follows, I will refer to cases in which an agent's justification is undermined as a result of a defeater that they have not psychologically registered as cases of 'external defeat'.[5]

While the question of whether there are cases of external defeat has historically been more controversial[6], recent research has unearthed strong evidence that we do need to countenance the existence of cases of external defeat. Most prominently, there is excellent reason for thinking that the phenomenon of testimonial injustice requires us to allow for the existence of external defeat. Testimonial injustice is a distinctively epistemic kind of injustice. More specifically, in one important kind of case of testimonial injustice the hearer doesn't give the speaker the credibility that they deserve as a result of systematic prejudice (Fricker, 2007). To take a famous example, consider the case of a scientist, A, who doesn't believe anything his female colleagues say, because he is a sexist (Lackey, 2018). Now suppose that A carries out some experiments that strongly support the hypothesis that p and that he comes to believe that p on this basis. Suppose, next, that a female colleague of his, B, discovers a serious flaw with the experiment, which she points out to A. Due to sexist bias, A discounts B's word and continues in his belief that p.

This is a paradigm case of testimonial injustice. In discounting B's word due to sexist bias, A fails to give B the credibility she is due, as a result of systematic identity prejudice. At the same time, it is also a paradigm case of defeat. After B's testimony about the flaw in the experiment is a reason against believing that p.

Now, it is worth mentioning that this case does not quite require us to countenance the existence of external defeat. After all, while A neither knows nor believes what B told him, it is plausible that A does know that (q =) B told him that there was serious flaw with the experiment. Crucially, the fact that q is itself widely acknowledged to be a reason against believing p. As a result, the fact that q itself constitutes a defeater for A's belief that p. Since A knows that q, A knows a fact that constitutes a defeater for his belief that p. So, we can explain why their justification for believing p

is defeated without invoking external defeat. We don't yet have compelling reason to countenance the existence of external defeat.

That said, the route to external defeat is quite short from here. To begin with, let's consider a variation of the case. More specifically, consider a variation in which A is radicalised in the following way. Not only does he have the disposition to dismiss testimony by women as a result of systematic identity prejudice, but he also thinks that what women have to say is simply not worth listening to in the first place. He is disposed to simply tune out whenever a female team member tells him something with the result that there isn't even uptake of what he was told in the first place. Otherwise put, in the original version of the case, A continues to keep track of what his female colleagues say (but then goes on to discount it). In the new version, he doesn't even keep track of what they say any longer.

Now, here is one crucial point. It cannot be that now that A has become radicalised (in that his bad epistemic disposition has gotten worse because now he is not only disposed to discount the word of their female colleagues, but he is also disposed to not even listen to them), he has successfully insulated himself from defeat he would have had had he been less radicalised (in that he is only disposed to discount the word of his female colleagues, but still disposed to listen to them). As a result, if the original version of the case is a case of defeat, then so is the new version. But since we have already seen that the original version is indeed a case of defeat, we have excellent reason to think that the new version is as well.

Here is another crucial point. The case in which A is radicalised is a case of external defeat. This is because, thanks to his radicalisation, none of the defeaters, i.e. that the experiment is flawed or that the female colleagues' asserted that is flawed, is psychologically registered by A. In this way, the phenomenon of testimonial injustice does provide a compelling reason to think that there is external defeat.[7]

Before moving on, it may be worth noting that we do not need to appeal to distinctively testimonial cases to make this point. In fact, once the testimonial cases are in clear view, it is easy to see that we can construct similar cases that don't involve testimony as a source of belief. For instance, consider a case in which a teacher is asking a question to the class. The only student who raises their hand is black. Even though the entire class, including the black student, is in plain view, as a result of racist bias, the teacher simply doesn't register that the black student raised their hand. They form the belief that no one in the class is willing to answer the question and proceed to explaining the answer themselves. It is hard to deny that if the case of the radicalised sexist scientists is a case of external defeat, then so is the case of the racist teacher.[8]

3 | DEFEAT AND VIRTUE EPISTEMOLOGY

Virtue epistemology is associated with a normative framework that allows us to assess attempts. (Following Sosa (2021) I will henceforth also refer to this framework as 'telic normativity'.) Attempts here have constitutive aims. As a result, we can ask whether or not a given attempt is *successful*, i.e. whether it attains its aim. Most importantly for present purposes, we can also ask whether a given attempt is *competent*, i.e. produced by an ability to attain its aim. Finally, we can ask whether a given attempt is *apt*, i.e. whether it is successful because competent.

Virtue epistemologists standardly take beliefs to be attempts that have truth as their constitutive aim. That said, my own preferred view is that belief constitutively aims at knowledge rather than truth (Kelp, 2021a, 2021b). While I will leave the question as to who is right on this issue open for now, I will return to it in due course. Given that belief is a kind of attempt, telic normativity

applies to belief. We can ask whether beliefs are successful, i.e. whether they are true/knowledge. In addition, we can also ask whether they are competent, i.e. whether they are produced by an ability to believe truly/know. And, finally, we can ask whether they are apt, i.e. successful because competent.

Virtue epistemology identifies epistemic properties such as justified belief and knowledge with normative properties of the framework for assessing attempts. In particular, knowledge is identified with apt belief[9] and, most importantly for present purposes, justified belief is identified with competent belief.[10]

It may be worth noting that that the resulting view is attractive on a number of counts. It offers attractive solutions to a range of central epistemological problems, perhaps most notably the Gettier problem and the value problem[11]. What's more, the view offers an attractive account of the epistemic normativity of belief as an instance of telic normativity, and promises to make sense of the centrality of knowledge in epistemology throughout history.

There is reason for optimism about a virtue epistemological account of defeat. This is because there is excellent reason to think that defeat can undermine competence of attempts more generally. Suppose that you are about to take a shot in archery. As you aim for the bullseye, I tell you that there is a strong wind blowing from the right. In this case, you need to adjust your aim to shoot competently. If you aim for the bullseye in exactly the same way you would were there to be no wind, your shot will not be competent. It is plausible enough to think that what is going on when you take a shot that is aimed right at the bullseye is that your shot is not competent because it is subject to defeat. In particular, it is plausible enough that my testimony that there is a wind blowing from the right constitutes a defeater for your shot, at least if it is aimed right at the bullseye. But, of course, if there is reason to think that defeat can undermine competence of attempts in general, there is reason to think that it can do so in the particular epistemic case we are interested in. This means that it should be possible to give a virtue epistemological account of defeat.

So much for the good news. At the same time, incorporating a viable account of defeat into virtue epistemology is by no means a trivial task. To see why not, note that beliefs for which we have defeaters may be produced by epistemic abilities. Suppose you tell me that the pen I am standing in front of is populated predominantly by cleverly disguised mules. Suppose I nonetheless form a perceptual belief that the animal I am looking at is a zebra. My belief will be the product of an exercise of an epistemic ability to tell a zebra from the way it looks. If this isn't obvious, note that had there not been cleverly disguised mules around, my belief would have qualified as knowledge.[12] As a result, it is hard to deny that my belief was produced by the exercise of an epistemic ability. After all, if it hadn't been produced by such an exercise, it is hard to see why it should be that it would have qualified as knowledge had there not been cleverly disguised mules around. In this way, there is reason to believe that beliefs for which we have defeaters may be produced by epistemic abilities.

What's more, the problem generalises to attempts for which we have defeaters more generally. To see this, let's return to the archery case. Suppose I am about to take a shot and you tell me that the there is a wind blowing from the right. Suppose I nonetheless take a shot that is aimed right at the bullseye. My shot will be the product of an exercise of my ability to hit the target. If this isn't obvious, note that had there not been a wind blowing, my shot would have found the bullseye. As a result, it is hard to deny that my shot was produced by an exercise of my ability to hit the target. After all, if it hadn't been the product of such an exercise, it is hard to see why it should be that it would have found the bullseye had there been no wind. In this way, there is reason to believe that

you can also exercise your ability to hit the target no matter whether you have also learned that there is wind blowing from the right.

We saw earlier on that there is reason to think that defeat can undermine competence of an attempt. At the same time, it transpires that attempts for which we have defeaters can be produced by relevant abilities, both in epistemic and non-epistemic cases. As a result, standard virtue epistemology, according to which an attempt is competent if and only if it is produced by a relevant ability, runs into trouble. It doesn't have the resources to accommodate the idea that defeat can undermine competence of attempts.[13]

4 | THE STRUCTURE OF DEFEAT: DEFEATERS AS RANGE LIMITERS

It transpires that thinking about defeat quickly leads to difficulties for virtue epistemology. Fortunately, there is a solution to this problem, which I will develop in this section.

First, note that there is independent reason to think that competence of attempts requires more than merely being produced by an exercise of an ability to attain the relevant success. To see this, consider the following case. You are a basketball player who has the ability to make layups. The game you are currently playing is about to end. In fact, you only have two seconds to score a basket to win the game from your midcourt position. Suppose that you produce a shot via an exercise of your ability to make layups, which, of course, doesn't even come close to the basket (Kelp, 2018).

In this case your shot is not competent. At the same time, it is produced via the exercise of an ability to score baskets. This means that a competent attempt requires more than being produced by the exercise of an ability to succeed. What more? One very plausible thought is that a competent attempt needs to be produced by an exercise of the *right kind* of ability. In the above case, your ability to produce layups isn't the right kind of ability for the kind of thing that you are attempting to do, i.e. score a basket from midcourt.

This raises the question as to what it takes for an attempt to be produced by the right kind of ability. To answer it, note that it is independently plausible that abilities are relative to ranges of attempt types. For instance, your ability to score layups in basketball is relative to a range of types of attempts you may make. It extends to attempts to score baskets from some distances not others, it may extend to attempts to score baskets from some angles and not others, and so on. Crucially, note that shots from midcourt aren't in the range of your ability to score layups. Accordingly, here is how the point that abilities are relative to ranges of attempt types can give us an attractive way of unpacking what it takes for an attempt to be produced by the right kind of ability: the attempt must be in the range of the ability that produced it. The view of competent attempts that we get then is one on which a competent attempt requires not only that it is produced via an exercise of a relevant ability but also that the attempt is in the range of the ability exercised. On this view, then, the reason why your shot from midcourt is not competent when it is produced via an exercise of an ability to score layups is that midcourt shots aren't in the range of the ability to score layups (Kelp, 2018).

With these points in play, let's return to the case of defeat of competence. Recall that here I tell you that there is a wind blowing from the right. When you go on to take a shot that is aimed right at the bullseye, your shot is not competent. In particular, your knowledge of what I told you constitutes a defeater here. Now, the crucial point is that once you know that there is a wind blowing from the right, the ability that involves aiming straight at the bullseye isn't the right kind

of ability for what you are attempting, i.e. hitting the target while a wind is blowing from the right. The problem here is exactly the same kind of problem as in the above basketball case.

We are now in a position to see how a view that takes abilities to be relative to ranges of attempt types can accommodate the idea of defeat undermining competence of attempts in general and of belief in particular. It does so by limiting the range of the abilities at issue. For instance, coming to know what I tell you limits the range of your ability to hit the target by aiming right at the bullseye with the result that if you take a shot via the exercise of this ability, your shot will not be competent. Similarly, coming to know that most of the animals in the pen before you are mules cleverly disguised to look like zebras limits the range of your ability to acquire true beliefs/knowledge about the presence of zebras. In particular, if you now come to believe that the animals in the pen you are looking at is a zebra, your belief will not be competent. In fact, in both cases, your range-limiting knowledge constitutes a defeater which prevents your attempt from being competent.

5 | THE SUBSTANCE OF DEFEAT: DEFEATERS AS EVIDENCE FOR UNSUCCESSFUL ATTEMPTS

The account of defeaters as range limiters allows us fit defeaters into the virtue epistemological framework. While this allows us to understand what defeaters are structurally, as it were, within a virtue epistemological framework, the question of a more substantive account of defeat remains open. If this isn't immediately obvious, consider once more the case in which you come to know that the pen before you is mostly populated by cleverly disguised mules. This constitutes a defeater for your belief that the animal you are looking at is a zebra. What we can say now is that your knowledge about the cleverly disguised mules limits the range of your ability to come to know that the animal is a zebra by looking. However, the question as to why exactly it is that this knowledge constitutes a defeater—i.e. why exactly it is that the range of your ability gets limited—remains open. In what follows, I will begin to answer it by developing a more substantive account of defeat.

Recall the lightweight characterisation of defeaters as reasons against. According to virtue epistemology, a justified belief is a species of competent attempt. What's more, we saw that there is reason to think that defeaters can undermine competence of attempts more generally. Given that this is so, it's plausible to think that defeaters are reasons against attempting; or, at the very least, they are reasons against attempting in a certain way.

It is worth noting that, on the virtue epistemological picture that emerges, not every reason against attempting (in a certain way) is competence-undermining. Suppose that I offer you a considerable sum for not aiming your next archery shot right at the bullseye. Now, you have a reason against attempting in a certain way. However, this doesn't mean that you have a competence-undermining reason for not aiming your next shot at the bullseye. If this isn't obvious, note that while taking a shot aimed right at the bullseye may well be practically irrational, it doesn't mean that a shot thus aimed is not competent.

Again, something similar may happen in the epistemic case. If you offer me a considerable sum for not believing that there is a laptop before me, I may have a reason against believing it. However, this doesn't mean that I have a justification-undermining reason for my perceptual belief that there is indeed a laptop before more. Even if forming this belief is now practically irrational, it doesn't mean that my perceptual belief isn't justified.[14]

What comes to light is that not all reasons against are competence-undermining. This may be because not all reasons against are defeaters or because not all defeaters undermine competence, at least by the lights of virtue epistemology. I do not mean to settle this question here. Instead, I

will rest content with noting that since, on virtue epistemology, it is competence-undermining reasons against that are of distinctive epistemological interest, it is competence-undermining reasons against that I will focus on in what follows and that I will refer tby 'defeaters'.

The obvious question that arises is what distinguishes defeaters from reasons against attempting that aren't competence-undermining. Here is my suggestion. *Defeaters are evidence that attempting, or at least attempting in a certain way, isn't successful.* While a handsome sum of money or the promise of a trove of knowledge may constitute a reason against attempting, it does not constitute evidence that attempting (in a certain way) isn't successful.[15] On the present view, we do not have defeaters for shooting or believing in a certain way here. At the same time, that a wind is blowing from the right is evidence that taking a shot aimed at the bullseye isn't successful. Likewise, that the pen you are looking at is populated predominantly by cleverly disguised mules is evidence that your perceptual belief that the animal before you is a zebra doesn't qualify as knowledge. These pieces of evidence are reasons that are competence-undermining and so defeaters for the attempts in question.

Does a defeater for an attempt always undermine competence of the attempt (at least unless it is itself defeated)? Or does it do so only sometimes? To get this issue into clearer view, consider a case in which it is true that most of the animals in the pen before you are cleverly disguised mules, but you have no inkling of this. Does this fact constitute a defeater for your belief that the animal is a zebra? Now, one might think that the answer is yes. After all, that most of the animals are cleverly disguised mules is a competence-undermining reason against believing that you are looking at a zebra. But if the answer is indeed yes, then defeaters don't always undermine competence of attempt. In the present case, for instance, your belief that the animal is a zebra is justified despite the existence of a defeater. By the same token, if defeaters always undermine competence, then more is needed for something to count as a defeater.[16]

I don't mean to settle the question whether defeaters always undermine competence here. What I would like to do instead is look at what else is needed for defeaters to undermine competence, leaving the question as to whether what else is needed is essential to defeaters to one side. For ease of exposition, I will assume that what else is needed isn't essential to defeaters. On this view, it is plausible that what is needed for defeaters to undermine competence is that one is related to defeaters in some way.

The question about this relaton is one of the key divides in the epistemological debate on defeat. First, one might think that the relation is psychological (henceforth also 'the psychological view'). On this view, a defeater undermines one's justification if and only if one stands in some member of a designated set of psychological relations to the defeater. It is easy to see that the psychological view cannot make room for external defeat. Since we have seen that there is reason to think that we must make room for external defeat, there is also reason to think that the psychological view cannot be correct.

The second prominent view takes the relation to be normative (henceforth also 'the normative view' (e.g. Alston, 2002; Graham & Lyons, 2021; Simion in press)). One way in which champions of the normative view take their view to differ from the psychological view is that psychological relations aren't what matters when it comes to defeaters undermining justification. Rather, what matters is positive epistemic standing: for a defeater to undermine justification it must be epistemically proper for one to have it.[17] Crucially, the normative view differs from the psychological view is that it has the resources to deny the psychological view's claim that a psychological relation is necessary for a defeater to undermine justification. The normative view can allow that a defeater can have positive epistemic standing, that it can be epistemically proper for one to have it, even when one does not bear any psychological relation to it. On the normative view that emerges,

then, a defeater, d, undermines one's justification if and only if it is epistemically proper for one to have d. And, more generally, a defeater, d, undermines competence of attempt if and only if it is epistemically proper for one to have d.[18]

With these points in play, let's take a look at how the virtue epistemological view of defeat handles paradigm cases of defeat. In the archery case, you acquire testimonial knowledge that there is a wind blowing from the right. Since you epistemically properly have defeaters that you know to be true and since it is epistemically proper for you to have defeaters that you epistemically properly have, it follows that it is epistemically proper for you to have defeaters that you know to be true, including that there is a wind blowing from the right. Moreover, since the fact that there is a wind blowing from the right is a defeater for your attempt to hit the target by aiming straight at the bullseye, if you do aim your shot right at the bullseye, it will not be competent. Likewise, in the zebra case, when the zookeeper tells you that the pen before you is mostly populated by cleverly disguised mules, you come to know what you are told. In addition, that the zookeeper tells you about the presence of cleverly disguised mules is a defeater for your belief that the animal you are looking at is a zebra. As a result, if you form this belief anyway, it will not be competent/justified.

6 | THE PROBLEM OF EXTERNAL DEFEAT

That the substantive virtue epistemological account of defeat can deal with paradigm cases of defeat is certainly good news. That said, note that all the paradigm cases of defeat we looked at are cases of internal defeat. At the same time, one of the central aims of this paper is to make room for external defeat. Accordingly, the question I want to ask next is whether our virtue epistemological account of defeat can do so and if so how.

Unfortunately, there is reason for pessimism. To see why, I'd like to begin by looking at three key facts about external defeat. First, external defeat affects the justificatory status of beliefs. After all, cases of external defeat are cases of *defeat*.

Second, external defeat turns on normative facts. After all, recall that, on the normative view, for defeat to obtain, i.e. for defeaters to undermine competence of attempts, it must be epistemically proper for one to have them. In this way, defeat turns in general turns on normative facts.[19]

Third, in cases of external defeat, these normative facts aren't normative facts about beliefs one has formed. Otherwise, they wouldn't be cases of *external* defeat. In fact, more generally, the normative facts on which external turns aren't (or at least needn't be) facts about attempts one has made. If this isn't immediately obvious, just consider our sexist scientist who doesn't bother to tune into what his female colleague tells him. That we have a case of external defeat doesn't turn on him attempting to do anything. For instance, to generate a case of external defeat we don't have to suppose that our sexist scientist attempted to listen to his female colleague but failed. Rather, to generate a case of external defeat, it will be enough if our sexist scientist does not attempt anything at all. As a result, the normative facts on which external turns aren't (or at least needn't be) facts about attempts one has made.

With these features of external defeat in play, let's turn to some features of virtue epistemology. First, according to virtue epistemology, knowledge and justified belief are identified with categories of telic normativity, to wit, competent and apt belief. That's one of the key claims of the view.

Second, according to virtue epistemology, telic normativity is autonomous. That is to say, it is not encroached upon by other kinds of normativity (Sosa, 2021, p. 40).[20]

Third, telic normativity presupposes that the agent has made an attempt (Sosa, 2021, p. 66). After all, whether an attempt is successful, competent, or apt presupposes that an attempt was made. In the epistemic case, this means that telic normativity presupposes that one has formed a belief. After all, whether one's belief is successful, competent, or apt, presupposes that one has indeed formed a belief.

We are now in a position to see exactly why the phenomenon of external defeat means trouble for virtue epistemology. By the first feature of virtue epistemology, justification of belief is identified with a normative category of telic normativity, i.e. competence. By the first feature of external defeat, external defeat affects the justificatory status of beliefs. This means that external defeat affects the status of beliefs as competent. By the second feature of external defeat, external defeat turns on normative facts. Since, additionally, by the second feature of virtue epistemology, telic normativity is autonomous, we get the result that the normative facts on which external defeat turns must be facts of telic normativity. After all, if this weren't the case, the status of beliefs as competent would turn on normative facts outside of telic normativity. However, this would mean that telic normativity is encroached upon and would thus be incompatible with the autonomy of telic normativity.

By the third feature of external defeat, the normative facts on which external defeat turns aren't normative facts about beliefs one has formed, and, more generally, they aren't (or needn't be) normative facts about attempts one has made. But all of this is incompatible with the third feature of virtue epistemology, according to which telic normativity presupposes that an attempt was made. After all, if the normative facts on which external defeat turns aren't (or needn't be) normative facts about some attempt one has made, then they cannot be facts of telic normativity. And since telic normativity presupposes that an attempt was made, there are no normative facts of telic normativity about attempts one hasn't made.

What comes to light is that virtue epistemology is incompatible with our three key facts about external defeat. By the same token, there is reason to think that virtue epistemology cannot make room for external defeat.[21]

7 | THE FUNCTIONALIST ACCOUNT OF ABILITIES

The argument that virtue epistemology cannot make room for external defeat rests on three key features of virtue epistemology. To the extent that virtue epistemologists want to make room in their epistemology for external defeat, they might think again whether they really want to hold on to all three of these claims. Perhaps one of them can be rejected after all, and perhaps this will usher the way towards a viable virtue epistemological account of external defeat.

While I do not mean to deny that this is one way to tackle the issue, it is not the approach I will pursue here. Rather than fixing a defect, I want to turn over a new leaf and develop a virtue epistemology from a slightly different starting point. More specifically, I want to start by thinking about the nature and normativity of abilities. The central property of this approach is the property of a function. Accordingly, to see how it works, it will be useful to start by taking a quick look at what a function is.

The kind of function I am interested in here turns on the existence of a feedback loop involving the functional item and a good functional effect, which it produces in a system[22]. ('Function' will henceforth refer to this kind of function unless otherwise noted). The heart is an example of an item with a function (Graham, 2012). It is widely recognised that one (if not the) key function of the heart is to pump blood. Why is that? According to the present account, the answer is that pumping

blood is good in that it contributes to the proliferation of genes that are responsible for its existence, by keeping us alive long enough to procreate. The fact that the heart pumps blood contributes to explaining why hearts exist which, in turn, contributes to explaining why they continue to pump blood. In this way, hearts exemplify exactly the kind of feedback loop characteristic of functions.

Now, here is the first key idea of the view I am trying to develop: abilities are ways of producing attempts that have the function of producing the relevant kind of success for the agent (henceforth also 'the functionalist account of abilities'[23]). For instance, to have the ability to score free-throws in basketball you must have a way of shooting free-throws that has the function of scoring free-throws for you. Of course, functions are unpacked as expected in terms of a feedback loop, here involving a way of producing attempts and a relevant success, which it produces for the agent: the way of producing attempts explains why the successes are produced and the production of successes explains why the way of producing attempts is in place. In the case of your ability to score free-throws, that you shoot in the way that you do must explain why you score free throws and the fact that you score free-throws must explain why you shoot in the way that you do.[24]

Now, crucially, it is widely agreed that functions have normative import. In particular, their normative import can be read off the answers to the following two questions:

1. Does the functional item fulfil its function?
2. Is the functional item functioning properly?

Regarding function fulfilment, the item meets the first normative standard if and only if the item does produce the functional effect in question. Function fulfilment is rather straightforward.

Things are a little more complex when it comes to proper functioning. To understand this normative standard, we need a bit of conceptual machinery. We have already seen one crucial concept: function fulfilment. The other two are normal functioning and normal conditions. Roughly, normal conditions are the conditions that obtain in the feedback loop in which the functional item produces the functional effect. And, again roughly, normal functioning is the way of functioning that produces the functional effect in the feedback loop, under normal conditions. To make these ideas a little more concrete, consider the heart once more. Here normal conditions include being hooked up to the arteries and veins of a certain kind of organism in a certain way and normal functioning is beating at a certain rate. Functioning normally (beating at a certain rate) under normal conditions (whilst hooked up to arteries and veins), the heart produces its functional effect (pumping blood).

With these points in play, we can now see what proper functioning amounts to for items with functions. In a nutshell, the idea is that proper functioning is normal functioning. When the heart is beating at a certain rate it is functioning properly. In order to figure out whether a functional item is functioning properly, then, we need to ask whether it is functioning normally, i.e. whether it is functioning in the way it does when producing the functional effect in the feedback loop, under normal conditions. And, again, the item meets the second normative standard if and only if it is.[25]

Of course, if the functionalist account of abilities holds, then so does the normative import of functions. In particular, we can ask whether, on a given occasion, an ability fulfilled its function and whether it was functioning properly. Since function fulfilment is unpacked in terms of production of the functional effect and since the functional effects of abilities are successful attempts, we get the result that an ability fulfils its function if and only if the attempt it produces is successful. Moreover, proper functioning is analysed in terms normal functioning, i.e. in terms of functioning in the way that produces the functional effect in the feedback loop, under normal conditions.

For abilities, normal functioning involves the production of attempts that are in the range of the ability. After all, that's the kind of functioning that produces the functional effect, i.e. successful attempts, under normal conditions. But now note that what we end up with is an account of the normativity of abilities, that features normative categories that attempts produced by them satisfy if and only if they are, respectively, successful and competent. In this way, the functionalist account of ability, in conjunction with the normative import of functions, will effectively allow us to recover two out of the three central normative categories of telic normativity.[26]

8 | PROFICIENCIES

The question remains, however, how this functionalist version of virtue epistemology can make room for external defeat. To answer it, recall first that telic normativity is a normative framework for attempts which invokes properties such as success and competence. Note that the normative framework that telic normativity provides is limited in scope. What telic normativity provides is a normative framework for is attempts in particular. It is far from clear whether the framework can be expanded in any substantive manner beyond attempts and, if so, how this might be done. Crucially, functionalist normativity is not limited in scope in this way.[27] After all, we can use functionalist normativity to assess lots of things besides attempts. In this way, functionalist normativity expands beyond attempts. We'll have functions when we have the relevant feedback loop between the functional item and the functional effect, and we have normative import when we have functions. When we saw that we can recover two normative categories of telic normativity in functionalist normativity, what we effectively saw was that (at least a certain part of) telic normativity can be embedded in the broader framework of functionalist normativity. While that is in itself an attractive result, what is most important for present purposes is that the normative framework we are employing is much broader in scope. As a result, its resources are not exhausted by the normativity of attempts, nor by the normativity of abilities. And it is precisely these additional resources that will allow us to make headway towards a better account of external defeat.

How so? To answer this question, I'd first like to distinguish abilities from what I will call proficiencies. To get a better handle on this distinction, consider the basketball case again. Whether you have an ability to make shots with your right, say, in a certain range turns on whether you have a way of shooting that has the function to make shots in that range. Note that while you may have this ability, it may well be that you rarely if ever exercise it. Perhaps this is because you are cautious, perhaps it is because you have taken a vow never to shoot with your right again, or perhaps it is for some other reason entirely. Now contrast this case with a case in which you don't have an ability, but you produce many shots. You are prolific at producing shots and have what we may call a prolificacy. Perhaps the number of successful shots you produce in a day is exactly the same in both cases.

Now, the key suggestion is that proficiencies combine abilities and prolificacies in a certain way. First, any genuine proficiency is also an ability. A mere way of attempting that is not an ability is not a proficiency, not even if it produces successful attempts prolifically. Suppose my way of attempting free-throws is by throwing basketball right up in the air. In this case, I don't have a genuine ability to make free throws, not even if I make many free throws, say because there is an army of clandestine helpers with wind machines that see to it that my attempts find the basket (Kelp, 2018). Likewise, I don't have a genuine proficiency to make free throws, again no matter whether I happen to make a lot of free-throws.

If a genuine proficiency is also an ability, the question that arises is what more is required for an ability to qualify as a proficiency. One might think that the answer is simply that one exercises one's ability a lot. But, again, this can't be quite right. To see this, consider two agents, A and B, who both have a certain ability. A exercises their ability rarely, but when they do, their attempts are virtually always successful. In contrast, B exercises it frequently but indiscriminately with the result that their attempts are virtually never successful. Let's suppose that A produces successful attempts as often as (or perhaps even more often than) B. While B is more prolific than A, it is not the case that B is more proficient than A. As a result, whether an ability qualifies as a proficiency cannot just be a matter of how often one exercises one's ability. Note also that what it takes for an ability to qualify as a proficiency cannot only turn on whether one exercises one's ability a lot and one frequently succeeds upon exercising one's ability. Again, the fact that one succeeds frequently might be accidental to exercise of the ability in which case the ability isn't a proficiency.

Unsurprisingly, I want to suggest unpacking the difference between an ability and a proficiency in functionalist terms. More specifically, at a minimum, a proficiency is an ability that has the function of producing successes across a range of trigger conditions, i.e. conditions that trigger the exercise of the ability.[28] Since proficiencies are functional entities, they have normative import. We can ask whether a proficiency fulfilled its function and whether it was functioning properly. Of course, proper functioning here is understood in the expected way in terms of normal functioning, i.e. the way of functioning that produces the functional effect in the feedback loop, under normal conditions. For proficiencies, normal functioning involves the obtaining of some trigger conditions, which trigger the exercise of the ability, which produces an attempt, which is successful, at least under normal conditions. The key point here is that normal functioning of the proficiency involves functioning that proceeds from the obtaining of some trigger conditions to the production of an attempt.

Let's take a look at how these rather abstract points about proficiencies play out in the sort of epistemic case that is of central interest here. First, note that many epistemic abilities are also proficiencies. Suppose, for instance, that I am looking at a blue dot on a white background. I have the perceptual ability to recognise blue things: I have a way of forming perceptual beliefs that has the function of producing perceptual knowledge about blue things. However, my perceptual ability is also a proficiency. It has the function of producing perceptual knowledge about blue things across a range of conditions that trigger its exercise. For instance, looking at a blue dot on a white background is among these trigger conditions. My ability to recognise blue things has the function of generating knowledge of the presence of something blue under those conditions.

Suppose, next, that I am told that you like Degas. I have the ability to learn from testimony: I have a way of forming testimonial beliefs that has the function of producing testimonial knowledge. At the same time, my testimonial ability is also a proficiency. It has the function of producing testimonial knowledge across a range of conditions. For instance, being told that you like Degas is among these conditions. My ability to learn from testimony has the function of generating knowledge that you like Degas under those conditions.

In all of these cases, normal functioning of proficiencies involves functioning from the obtaining of some trigger conditions to the production of a belief, which, under normal conditions, qualifies as knowledge. But since normal functioning is proper functioning, the result that we get is that an epistemic proficiency will function properly only if the obtaining of the trigger conditions leads to the formation of a belief. In this way, then, the result that we get is that if the trigger conditions obtain, it is proper for the proficiency to output a belief.

Before moving on to external defeat, I'd like to mention one important point about epistemic proficiencies. Note that these break down into at least two components. One component involves the uptake of information, the other the formation of beliefs. What I'd like to focus on here is

the component involving information-uptake. Note that this is an epistemic proficiency in its own right which is a constituent of the broader epistemic proficiency that has the function of producing knowledge. Here function fulfilment consists in the uptake of information and normal functioning consists in the functioning from the obtaining of some trigger conditions to the uptake of propositional content (henceforth 'content' for short) that, under normal conditions, qualifies as information.[29] In the blue dot case, my ability to recognise blue is a proficiency with the function of generating knowledge of the presence of blue things via a contained proficiency to take up information about the presence of blue things. Likewise, in the Degas case, my ability to learn from testimony is a proficiency with the function of generating testimonial knowledge via a contained proficiency to uptake testimonial information—here that you like Degas. It is these information-uptake proficiencies that are key to making room for external defeat, as I will explain in what follows.

9 | EXTERNAL DEFEAT

With these points on proficiencies and their normative import in play, let's return to external defeat. In what follows, I will outline how proficiencies—and, in particular, information-uptake proficiencies—can pave the way towards a better account of external defeat.

To begin with, note that the norm that corresponds to the proper functioning of information-uptake proficiencies is a genuinely epistemic norm. This is because it derives from a genuinely epistemic function of epistemic proficiencies (e.g. Simion, 2018), i.e., ultimately, the production of knowledge that is the function of the proficiencies of which information-uptake proficiencies are constituents.

Second, this epistemic norm is a norm one may fail to live up to without forming any belief or taking up any contents at all. For instance, in the case in which I am looking at a blue dot on a white background, my epistemic ability to perceptually recognise blue things contains a proficiency which, when functioning epistemically properly, will take up the content (and, under normal conditions, the information) that the dot is blue upon being presented with a blue dot on a white background. If I fail to take up the content that the dot is blue, I am in violation of this norm. Similarly, in the testimony case, my ability to learn from testimony involves a proficiency to take up information. When it functions epistemically properly, it will take up the content (and, under normal conditions, the information) that you like Degas upon me being told that you do. If I fail to take up the content that you like Degas, I violate this norm.

What comes to light is that the epistemic proficiencies give us norms that are genuinely epistemic and that one may violate without forming beliefs or even taking up any contents at all. It is not hard to see that this is takes us a long way toward making room for external defeat in our epistemology.

The final step is to connect these norms with the account of defeat. My suggestion is that epistemic proper functioning of information-uptake proficiencies corresponds to a way in which it is epistemically proper for one to have the relevant contents (and, under normal conditions, the information). For instance, in the blue dot case, when epistemic proper functioning of my proficiency involved with the uptake of information about the presence of blue things will lead me to take up the content (and, under normal conditions, the information) that the dot on the white background I am looking at is blue, it is epistemically proper for me to have the content (and, under normal conditions, the information) that the dot is blue. And, in the Degas case, when epistemic proper functioning of my testimonial proficiency will lead me to take up the content

(and, under normal conditions, the information) that you like Degas, it is epistemically proper for me to have the content (and, under normal conditions, the information) that you like Degas.

We are now in a position to see how we can make room for cases of external defeat in our epistemology. These cases are cases in which it is epistemically proper for one to take up and thus have a certain contents (and, under normal conditions, certain information) that qualify as a defeaters, but one does not take up said contents (information).

By way of illustration, let's return to the case of the racist teacher from Section 2. Recall that, in this case, a black student raises their hand in class. As a result of racism, the teacher doesn't even register that the black student raised their hand and forms the belief that no one is willing to answer the question that was asked. In this case, the teacher's belief that no one is willing to answer the question is clearly not justified. The present account can explain this. In this case, it is epistemically proper for them to take up the information that the black student who raised their hand did indeed raise their hand. This is because their perceptual epistemic abilities are also proficiencies which, when functioning properly, take up the content (and, under normal conditions, the information) that the black student in question raised their hand. But, of course, that the black student in question did raise their hand is a defeater for their belief that no one in the class is willing to answer the question. Since it is epistemically proper for the racist teacher to take up a content (and, under normal conditions, the information) that is a defeater for their belief that no one in the class is willing to answer the question, their belief that no one in the class is willing to answer the question suffers from defeat. Since, at the same time, they do not take up the defeater, what we are looking at is a case of external defeat.

In the above case, it is a perceptual proficiency that generates a defeater. It is easy enough to see that other proficiencies can also do so. To see this, let's return to the case of the radicalised sexist scientist, A, who doesn't even tune into what his female colleague tells him because he is a sexist and his female colleague, B, who found a flaw with one of his experiments. In this case, the A's epistemic ability to learn from testimony is a proficiency, which contains an information-uptake proficiency. When functioning epistemically properly, this information-uptake proficiency will take up the content (and, under normal conditions, the information) that there is a flaw with the experiment based on which he believes that p. As a result, it is epistemically proper for A to take up the content (and, under normal conditions, the information) that there is a flaw with the experiment based on which he believes that p. But, as we have already seen, this means that it is epistemically proper for him to have a defeater for his belief that p. This means that A's beliefs that p suffers from defeat. Since, in this case, A is radicalised and doesn't even psychologically register the defeater that it is epistemically proper for him to have, the case is another case of external defeat.

What comes to light is that proficiencies support epistemic norms that one can violate without forming beliefs or taking up information. When we fail to have information that, thanks to the existence of a proficiency, it is epistemically proper for us have, we may have cases of external defeat. This happens when the information that it is epistemically proper for us to have but that we don't have are defeaters. In this way, proficiencies allow us to make room for external defeat in our epistemology.

10 | CONCLUSION

This paper has developed a novel virtue epistemological account of defeat. I have developed an account of defeaters as evidence that attempting, or at least attempting in a certain way, is unsuc-

cessful. In addition, I argued for an account of defeat that is normative at least in that a defeater defeats some justification one has when it is epistemically proper for one to have this defeater, no matter whether one also stands in some psychological relation to it. I showed that this account can accommodate plausible cases in which defeat undermines competence of attempts both in epistemology and beyond.

Another central idea I developed in this paper was that cases of testimonial injustice provide us with excellent reason to think that we must make room for external defeat in our epistemology. This turned out not to be an easy task. I argued that standard virtue epistemology cannot successfully do so. To overcome this difficulty, I developed a distinctively functionalist version of virtue epistemology, which embraces a functionalist account of abilities, and uses the normative import of functions to recover two of the three central normative categories of standard virtue epistemology. One important consequence of the move to functionalist virtue epistemology was that it gave us a broader normative framework, one that isn't limited to the normativity of attempts. It is this broadening up that made room for proficiencies, which are abilities with the further function of producing successful attempts under certain trigger conditions. The functionalist normativity of proficiencies supports epistemic norms that one can violate without holding beliefs or having information. More specifically, the idea was that in cases of external defeat one doesn't have a defeater that it is epistemically proper for one to have. In this way, proficiencies allowed us to make room for external defeat in our epistemology.

ENDNOTES

[1] A notable exception is Greco who offers an account of defeat for his subjective justification condition on knowledge and justified belief (e.g. 2010). Since Greco's account runs into the in-principle problem that standard versions of virtue epistemology encounter for cases of external defeat (see Section 9), I will not discuss the view in detail here.

[2] One may wonder just how lightweight the lightweight characterisation of defeat really is. After all, aren't there accounts of defeat that explain defeat without invoking reasons at all? Most notably, consider the perhaps most influential account of defeat from the process reliabilist camp, the alternative reliable process view (Goldman, 1979; Lyons, 2009). Very roughly, according to process reliabilism, whether you believe justifiably turns on whether your belief is produced via reliable cognitive processes, i.e. processes that tend to produce beliefs with a favourable truth to falsity ratio. And whether your justification for believing something is defeated turns on whether you have an alternative reliable process available such that were it to be used, you would not hold your belief. Crucially, defeat is explained purely in terms of processes, reasons don't feature at all here. Note also that this is no accident. Process reliabilists take pride in the fact that their account of justification is naturalistically respectable. Part of what makes it so is that it features no normative properties such as reasons. But if there actually are live accounts of defeat that don't accommodate the lightweight characterisation, one may wonder just how lightweight the characterisation really is. Three comments on this. First, even though process reliabilists don't state their account of defeat in terms of reasons, it doesn't follow that their account of defeat isn't compatible with the lightweight characterisation of defeat. After all, it may be that reasons can be analysed in terms of available reliable processes. Note that in this case, the process reliabilist account does achieve a nice fit with the lightweight characterisation. After all, the lightweight characterisation isn't meant as a substantive account of defeat and leaves open the possibility that the key property of reasons against believing admits of further analysis, including along process reliabilist lines. Second, defeat is a general normative phenomenon. That is to say, it doesn't only occur in the epistemic domain. At the same time, it is far from clear that, in all normative domains in which defeat occurs, justification and defeat can be unpacked along process reliabilist lines. (For instance, note that process reliabilism is structurally a kind of rule consequentialism. While this may be plausible for the epistemic domain, it is not clear that it is equally plausible for e.g. the practical domain, which may require an act consequentialist treatment.) If it cannot, the prospects for a fully general account of defeat in terms of alternative reliable processes are dim, to say the least. What's more, we may just have to revert to a general account in terms of reasons, perhaps with a process reliabilist account of what it takes to have reasons in the epistemic

domain. Third, to the best of my knowledge, the alternative process account of defeat is the only account that doesn't naturally fit with the lightweight characterisation of defeat. At the same time, it has come under heavy criticism (e.g. Beddor 2014). And in order to develop a viable account of defeat, process reliabilists have started to try and make room for reasons in their epistemology, or at least some normative category in the vicinity (e.g. Beddor 2021; Graham and Lyons 2021). In light of this, no matter whether literally everyone is on board with the lightweight characterisation, I take it that it is lightweight enough to provide a useful starting point for theorising about defeat.

[3] Note that defeat works differently in these two cases. In the cancellation case, you get a reason against holding your belief by getting a reason for thinking that it is false. In the red-light case, in contrast, you get a reason against holding your belief by getting a reason for thinking that your source is inadequate. This distinction is known in the literature as the distinction between rebutting and undercutting defeat (Pollock, 1986).

[4] What does it take for a defeater, d, to be psychologically registered? For present purposes, I follow Lackey (e.g. 1999, 2003) in that one needs to have some doxastic attitude towards d (like beliefs or doubt).

[5] Note that the phenomenon sometimes goes under the labels 'propositional defeat' (e.g. Bergman, 2006) or 'normative defeat' (e.g. Lackey, 2008). I am not particularly happy with either term. The first suggests an analogy with propositional justification. Propositional justification is something that internalists can and indeed happily do acknowledge the existence of. The existence of what I call external defeat is incompatible with internalism. This is why I prefer not to go with 'propositional defeat' as a label. Normative defeat suggests that there is also non-normative defeat. Again, I take it that this is false, which is why I am not satisfied with this label either. That said, nothing hinges on this choice of terminology.

[6] In particular, note that epistemic internalists (e.g. BonJour, 1985; Chisholm, 1966; Conee & Feldman, 2004) will have to deny the existence of external defeat. While card-carrying externalists like myself may proceed to argue from the existence of external defeat to the falsity of internalism, internalists may venture to turn the argument on its head. Does this mean that we are in a deadlock? No. As I argue elsewhere, there is independent reason to think that the correct epistemology of defeat must be externalist (Kelp, 2020).

[7] One might wonder whether cases of external defeat featuring testimonial injustices aren't really cases of moral failures rather than cases of epistemic failures. I must confess that I don't find this move particularly plausible. Here is why. First, it has the unpalatable consequence that tuning up epistemically bad properties can lead to an improvement of an agent's epistemic position. In the above case, making the sexist scientist more sexist such that he not only discounts his female colleagues' words but doesn't even tune in to what they say will amount to an improvement in his overall epistemic state. Second, consider yet another variation of the case in which A systematically mishears what he is told by female colleagues about his work. Whenever he actually encounters disagreement, he hears agreement. It is perhaps even harder to believe that this trait should lead to an improvement of his epistemic position towards propositions about his work. Third, note that we can now even drop the injustice component of the case. We may suppose that A simply mistakes disagreement by anyone for agreement. Again, it's implausible that, as a result, A should be insulated from defeat. At the same time, the charge that the issue is really an ethical one becomes less credible once the injustice component is removed. It may be worth noting that there are further arguments for the existence of external defeat. One is inspired by Simion's (in press) argument that moral blameworthiness requires the absence of (epistemically) blameless ignorance. If so, in cases of moral blameworthiness, offenders are either aware that they are doing something wrong or else they (epistemically) should be. For instance, suppose our sexist scientist, A, is also the line manager of his female colleague, B, and that A promotes a male colleague, C, over B as a result of his sexism. Suppose further that A is not aware that he is doing something wrong. He thinks C deserves the promotion over B. But that's because, as a result of sexism, he pays close attention to all of C's work and none of B's. Now, we surely want to allow that A is morally blameworthy for his sexist promotion of C over B here. But if so, since A is ignorant that it is wrong to promote C over B, we must allow that A falls short epistemically in that he should be aware that it is wrong to promote C over B but doesn't. However, it is hard to see how this could be unless we allow for external defeat. Another argument is by Goldberg (2018) who uses cases of agents occupying certain social roles for this issue. While I agree with Goldberg, I believe that his cases require a slightly different treatment than the one I am discussing here. Due to limitations of space, I will have to leave discussion of them for another occasion. That said, once the existence of external defeat is duly acknowledged, I'd expect resistance to counting the above cases as cases of external defeat to wane accordingly. That's why Goldberg's cases are worth mentioning in any case.

[8] For more on defeat see e.g. (Bergman, 2006; Goldberg, 2018; Goldman, 1979; Lackey, 2008; and Pollock, 1986). For a recent volume on defeat see (Brown & Simion, 2021).

[9] Note that, on a knowledge-centric version of virtue epistemology, success and aptness coincide in the case of belief, i.e. a belief is apt if and only if it is successful. I have argued elsewhere that this is not a problematic consequence of the view (Kelp, 2017, 2018).

[10] For recent defences of virtue epistemology see e.g. (Broncano-Berrocal, 2017; Carter, 2016; Greco, 2010; Miracchi, 2015; Pritchard et al., 2010; Riggs, 2002; Sosa, 2015, 2021; Turri, 2016; Zagzebski, 1996). Recent collections on virtue epistemology include (Fernandez Vargas, 2016; Greco & Turri, 2012; Kelp & Greco, 2020). I have also defended a version of virtue epistemology e.g. in (Kelp, 2017, 2018).

[11] For more on the Gettier problem see e.g. (Gettier, 1963; Shope, 1983) For more on the value problem, see e.g. (Kvanvig, 2003; Pritchard et al., 2010).

[12] I am assuming (as I may) that had there not been cleverly disguised mules around, you wouldn't have told me that there are.

[13] One might wonder whether this isn't too quick. After all, virtue epistemologists standardly take abilities to be relative to conditions, C. And couldn't they just hold that defeat undermines competence by precluding C? No. While abilities are relative to C, virtue epistemologists will do well not to take competent attempts to require that C be in place. Consider a case in which you take a shot that would have hit the target had it not be for a gust of wind that no one could have predicted. In this case, your shot is clearly competent. If we take competent attempts to require that C obtain, we cannot accommodate this datum. After all, your ability to hit the target is relative to sufficiently normal winds. This point is, if anything, even more important when we turn to the envisaged applications in epistemology. After all, there is a range of cases in which C are not satisfied and yet agents form justified beliefs. Gettier cases and sceptical cases are the most prominent examples here. To secure the correct verdict that agents in Gettier and sceptical cases have justified beliefs, it is imperative that virtue epistemologists allow that attempts can be competent even when C are not in place.

[14] Note that some think that there cannot be practical reasons against believing (e.g. Shah, 2006). It is possible to run a version of the case with what, on the face of it, is a kind of epistemic reason: you are offered a trove of knowledge in exchange for not believing. If you think that, in this case, there isn't a reason against believing either, what you think may entail that reasons against believing are defeaters in the epistemic case. Even so, since virtue epistemology's background normative framework extends beyond the epistemic case and since a similar move is not promising for attempts in general, the point remains worth bearing in mind.

[15] Standard virtue epistemologists who take belief to constitutively aim at truth may opt for an account of competence-undermining reason against attempting in terms of aptness rather than truth. Since on my view belief constitutively aims at knowledge, and since, in any case, aptness and success coincide in the case of belief, I will set this complication aside here.

[16] It is worth noting that we may in addition countenance a category of knowledge defeaters in addition to justification defeaters, where a knowledge defeater is a defeater that undermines knowledge without undermining justification. If we recognise knowledge defeaters, the fact that most of the animals in the pen are cleverly disguised mules will count as a knowledge defeater. That said, for present purposes I want to set the issue of knowledge defeaters aside and focus on justification defeaters only. Accordingly, throughout this paper, 'defeater' refers to justification defeaters.

[17] I take 'epistemically proper' to denote a generic normative property, signalling the existence of some corresponding epistemic norm. It is meant to remain neutral on what specific type of epistemic norm we are dealing with here. For instance, it is meant to remain neutral on whether the norm has the force of a 'should', a 'may', or something else entirely. As I will argue in due course, I take the relevant norms to be generated by functions (Section 6). Accordingly, for present purposes, I will rest content with observing that further questions about the nature of the norm under consideration will be settled by whatever the correct theory of the normative import of functions says about this.

[18] It may be worth noting that, for present purposes, I do not need to take a stance on the question as to whether some psychological relation is sufficient for a defeater to undermine justification. What matters is that no psychological relation is required for this. After all, that's what is needed to make room in our epistemology for external defeat. Accordingly, for present purposes, it may well be that the correct view is a hybrid one, according to which the psychological and the normative view each specifies a sufficient condition on what it takes for a defeater to undermine justification. That said, I do worry that the psychological view will overgenerate defeat. Sexists may

have defeaters for the word of women simply in virtue of the fact that they believe that women are not trustworthy. In my view, this can't be right. Defeat cannot be that easy to come by. In light of this, I will take it that the normative view is the right one in what follows.

[19] Recall that my account is compatible with a hybrid view, according to which the psychological and the normative view each specify a sufficient condition on what it takes for a defeater to undermine justification (fn.18). Crucially, in cases of external defeat, the defeater isn't psychologically registered. This means that, even on the hybrid view, external defeat turns on normative facts.

[20] It may be worth noting that virtue epistemologists hold that telic normativity is autonomous with good reason. An archer's shot that constitutes a heinous murder may nonetheless be a good shot (qua shot). In fact, it may be just as good a shot (qua shot) as a shot that heroically saves a life. Similarly, a heinous belief about how to best implement a fascist regime may still be a good belief (qua belief). In fact, it may be just as good a belief (qua belief) as a heroic belief about how to best implement a democracy. (It is easy enough to see that similar examples can be found for other potentially encroaching types of normativity such as practical or aesthetic normativity.) These considerations show that the quality of an attempt (qua attempt) is independent of whatever other normative qualities the attempt may have. Since it is hard to see how this could be unless telic normativity was autonomous, there is excellent reason to think that telic normativity is indeed autonomous.

[21] This is the kind of in principle problem I mentioned in fn.1, the one that Greco's account of defeat also encounters, simply in virtue of being a virtue epistemological account of defeat.

[22] The most popular account of functions is the etiological account of functions according to which functions turn on a history of successes. Prominent defences include (Godfrey-Smith, 1994; Millikan, 1984; Neander, 1991). For applications to epistemology see e.g. (Graham, 2012; Simion, 2019). That said, the etiological account of functions doesn't offer the only way of explaining feedback functions. A promising alternative is the organisational theory of functions which has been defended in e.g. (Christensen & Bickhard, 2002; McLaughlin, 2000).

[23] The perhaps most prominent defender of a functionalist account of abilities is Ruth Millikan (e.g. 2004).

[24] It may be worth noting that the above functionalist account of abilities differs from the standard virtue epistemological account of abilities in at least two important respects. First, according to the standard virtue epistemological account, abilities are analysed in terms of dispositions (e.g. Greco, 2010; Sosa, 2015), whereas on the functionalist account, they are analysed in terms of functions. Second, according to the standard virtue epistemological account, abilities are properties of agents rather than properties of ways of attempting (Sosa, 2015). I have defended both elements of the functionalist account of abilities elsewhere (Kelp, 2018, in press), and I will not rehearse these arguments here. Instead, I will rest content with exploring the prospects of the functionalist account of abilities for making room for external defeat in a virtue epistemology. If it can do so, this will in itself be a significant enough result, no matter whether, in addition, there is independent reason to think that it is preferable to the standard virtue epistemological account of abilities.

[25] Earlier on (fn.17) I mentioned that 'proper' signals the presence of some corresponding norm. I left the question as to the precise nature of this norm open (e.g. whether the norm has the force of a 'should', a 'may', or something else entirely). I did say that this question is to be settled by the correct theory of the normative import of functions. Now, in my view, proper functions give rise to norms that have the force of 'should'. For instance, to say that the heart is functioning properly by beating at a certain rate is to say that the heart should beat at a certain rate. If this is correct, then epistemic norms generated by proper epistemic functioning will also have the force of a 'should'. While I think that this is the right result, I recognise that the point is controversial. Perhaps it is less problematic once we acknowledge that these shoulds are generated by functions in a perfectly familiar way. Even so, this commitment is optional for present purposes in the sense that it may be that functions don't support norms with the force of a 'should'. This is also why I will continue to state the view in terms of 'propriety' and its cognates here.

[26] What about aptness, the third category of telic normativity? Perhaps functionalist normativity can be developed to make room for aptness. But note that, on my view, it's not really an issue if this cannot be done. To see why, recall first that while standard virtue epistemology takes the epistemic success of belief to be truth, on my view it is knowledge (e.g. Kelp, 2018, 2021a). But, of course, even on the above functionalist view, success is a central normative category, even if we cannot make room for aptness as a central normative category. After all, function fulfilment is a central normative category—in fact, it is the central normative category. And since for an ability to fulfil its function is for it to produce a successful attempt, success is a—and arguably, the—central normative category here too. On a knowledge-centric version of virtue epistemology, we don't really need the normative

category of aptness to explain the normativity of knowledge and justified belief. We can make do with the categories of success and competence. And since those categories can be recovered by the above functionalist view, the question of whether there is room for aptness as well is of comparatively little consequence, at least for champions of knowledge-centric versions of virtue epistemology. In fact, my own preferred view departs from virtue epistemology in that it abandons aptness as a central normative category and argues that the intuitions supporting the claim that aptness is a central normative category can be given a fully adequate alternative explanation (Kelp, 2021a).

[27] This is also clearly recognised by Simion (in press) who goes on to develop process functionalist account of external defeat. The central difference between Simion and myself is that Simion's functionalism is agent-neutral, whereas mine is agent-relative. I argue for the benefits of the agent-relative view in more detail elsewhere (Kelp, in press).

[28] Beyond that, proficiencies come in degrees. In particular, I want to suggest that a maximal proficiency is one that involves a maximal ability, i.e. a way of producing attempts that has a maximal success to failure ratio and does so no matter what conditions one may find oneself in. In addition, and more importantly for present purposes, a maximal proficiency is also maximally calibrated to the ability. By this I mean that the conditions under which the ability has the function of producing successes coincides with the conditions under which an attempt produced by the ability would be successful. Degrees of proficiency can then be measured in terms of approximations to maximal proficiencies.

[29] I am taking it that information is at least factive such that one can take up the information that p only if p is true. In cases in which it seems that p but p is false, one cannot take up the information that p. For instance, if the animal before you looks like a zebra but is a cleverly disguised mule, then you cannot take up the information that the animal is a zebra. In this case, I will say that what you take up is the propositional content that the animal is a zebra. A propositional content here, is, very roughly, what the good and the bad case have in common and what, in the good case, qualifies as information.

REFERENCES

Alston, W. (2002). Plantinga, naturalism, and defeat. In J. Beilby (Ed.) *Naturalism defeated? Essays on Plantinga's evolutionary argument against naturalism* (pp. 176–203). Cornell University Press.

Bergman, M. (2006). *Justification without awareness*. Oxford University Press. https://doi.org/10.1093/0199275742.001.0001

BonJour, L. (1985). *The structure of empirical knowledge*. Harvard University Press. http://doi.org/10.2307/2185273

Broncano-Berrocal, F. (2017). A robust enough virtue epistemology. *Synthese*, *194*(6), 2147–74. http://doi.org/10.1007/s11229-016-1043-y

Brown, J. & Simion, M. (Eds.) (2021). *Reasons, justification, and defeat*. Oxford University Press. http://doi.org/10.1093/oso/9780198847205.001.0001

Beddor, B. (2014). Process reliabilism's troubles with defeat. *The Philosophical Quarterly*, 65(259), 145–159. https://doi.org/10.1093/pq/pqu075

Beddor, B. (2021). Reasons for Reliabilism. In Brown, J. & Simion, M. (Eds.), *Reasons, Justification, and Defeat* (pp. 146–176). Oxford University Press. https://doi.org/10.1093/oso/9780198847205.003.0007

Carter, A. (2016). Robust virtue epistemology as anti-luck epistemology: A New Solution. *Pacific Philosophical Quarterly*, *97*(1), 140–155. http://doi.org/10.1111/papq.12040

Chisholm, R. (1966). *Theory of knowledge*. Prentice-Hall. http://doi.org/10.2307/2183834

Christensen, W., & Bickhard, M. (2002). The process dynamics of normative function. *The Monist*, *85*(1), 3–28. http://doi.org/10.5840/monist20028516

Conee, E., & Feldman, R. (2004). *Evidentialism: Essays in epistemology*. Oxford University Press. http://doi.org/10.1093/0199253722.001.0001

Fernandez Vargas, M. (Ed.) (2016). *Performance epistemology: Foundations and applications*. Oxford University Press. http://doi.org/10.1093/acprof:oso/9780198746942.001.0001

Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press. http://doi.org/10.1093/acprof:oso/9780198237907.001.0001

Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, *23*(6), 121–23. http://doi.org/10.1093/analys/23.6.121

Godfrey-Smith, P. (1994). A modern history theory of functions. *Noûs*, *28*(3), 344–362. http://doi.org/10.2307/2216063

Goldberg, S. (2018). *To the best of our knowledge: Social expectations and epistemic normativity*. Oxford University Press. http://doi.org/10.1093/oso/9780198793670.001.0001

Goldman, A. (1979). What is justified belief. In G. Pappas (Ed). *Justification and Knowledge* (pp. 1–25). D. Reidel.

Graham, P. (2012). Epistemic entitlement. *Noûs*, *46*(3), 449–82. http://doi.org/10.1111/j.1468-0068.2010.00815.x

Graham, P., & Lyons, J. (2021). The structure of defeat: Pollock's evidentialism, Lackey's framework, and prospects for reliabilism. In J. Brown & M. Simion (Eds.) *Reasons, justification, and defeat* (pp. 39–68). Oxford University Press. https://doi.org/10.1093/oso/9780198847205.003.0003

Greco, J. (2010). *Achieving knowledge*. Cambridge University Press. http://doi.org/10.1017/CBO9780511844645

Greco, J., & Turri, J. (Eds.) (2012). *Virtue epistemology: Contemporary readings*. MIT Press.

Graham, P. & Lyons, J. (2021). The Structure of Defeat: Pollock's Evidentialism, Lackey's Framework, and Prospects for Reliabilism. In Brown, J. & Simion, M. (Eds.), *Reasons, Justification, and Defeat* (pp.39–68). Oxford University Press. https://doi.org/10.1093/oso/9780198847205.003.0003

Kelp, C. (2017). Knowledge first virtue epistemology. In A. Carter, E. Gordon & B. Jarvis (Eds.) *Knowledge first: Approaches in epistemology and Mind* (pp. 223–245). Oxford University Press. http://doi.org/10.1093/oso/9780198716310.003.0011

Kelp, C. (2018). *Good thinking: A knowledge first virtue epistemology*. Routledge. http://doi.org/10.4324/9780429455063

Kelp, C. (2020). Internalism, phenomenal conservatism, and defeat. *Philosophical Issues*, *30*(1), 192–204. http://doi.org/10.1111/phis.12180

Kelp, C. (2021a). *Inquiry, knowledge, and understanding*. Oxford University Press. http://doi.org/10.1093/oso/9780192896094.001.0001

Kelp, C. (2021b). Theory of inquiry. *Philosophy and Phenomenological Research*, *103*(2), 359–84. http://doi.org/10.1111/phpr.12719

Kelp, C. (in press). Agent Functionalism. Sylvan, K. (Ed.), *The Blackwell Companion to Epistemology*. Wiley Blackwell.

Kelp, C. and Greco, J. (Eds.) (2020). *Virtue-theoretic epistemology: New methods and approaches*. Cambridge University Press. http://doi.org/10.1017/9781108666404

Kvanvig, J. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge University Press. http://doi.org/10.1017/CBO9780511498909

Lackey, J. (2003). Testimonial knowledge and transmission. *The Philosophical Quarterly*, *49*(197), 471–90. http://doi.org/10.1111/1467-9213.00154

Lackey, J. (2003). A minimal expression of non-reductionism in the epistemology of testimony. *Noûs*, *37*(4), 706–23. http://doi.org/10.1046/j.1468-0068.2003.00457.x

Lackey, J. (2008). *Learning from words. Testimony as a source of knowledge*. Oxford University Press. http://doi.org/10.1093/acprof:oso/9780199219162.001.0001

Lackey, J. (2018). Credibility and the distribution of epistemic goods. In K. McCain (Ed.) *Believing in accordance with the evidence* (pp.145–168). Springer. http://doi.org/10.1007/978-3-319-95993-1_10

Lyons, J. (2009). *Perception and basic beliefs. Zombies, modules, and the problem of the external world*. Oxford University Press. http://doi.org/10.1093/acprof:oso/9780195373578.001.0001

McLaughlin, P. (2000). *What functions explain: Functional explanation and self-reproducing systems*. Cambridge University Press. http://doi.org/10.1017/CBO9780511498510

Millikan, R. (1984). *Language, thought, and other biological categories*. MIT Press.

Millikan, R. (2004). *Varieties of meaning*. MIT Press.

Miracchi, L. (2015). Competence to know. *Philosophical Studies*, *172*(1), 29–56. http://doi.org/10.1007/s11098-014-0325-9

Neander, K. (1991). The teleological notion of 'function'. *Australasian Journal of Philosophy*, *69*(4), 454–468. http://doi.org/10.1080/00048409112344881

Pollock, J. (1986). *Contemporary theories of knowledge*. Hutchinson. http://doi.org/10.2307/2104442

Pritchard, D., A., Millar, & A., Haddock. (2010). *The nature and value of knowledge*. Oxford University Press. http://doi.org/10.1093/acprof:oso/9780199586264.001.0001

Riggs, W. (2002). Reliability and the value of knowledge. *Philosophy and Phenomenological Research*, *64*(1), 79–96. http://doi.org/10.1111/j.1933-1592.2002.tb00143.x

Shah, N. (2006). A new argument for evidentialism. *The Philosophical Quarterly*, *56*(225), 481–98. http://doi.org/10.1111/j.1467-9213.2006.454.x

Shope, R. (1983). *The analysis of knowing: A decade of research*. Princeton University Press. http://doi.org/10.2307/2107319

Simion, M. (2018). No epistemic norm for action. *American Philosophical Quarterly*, *55*(3), 231–238. https://doi.org/10.2307/45128617

Simion, M. (2019). Knowledge first functionalism. *Philosophical Issues*, *29*(1), 254–67. http://doi.org/10.1111/phis.12152

Simion, M. (in press). Resistance to Evidence and the Duty to Believe. *Philosophy and Phenomenological Research*.

Sosa, E. (2015). *Judgment and agency*. Oxford University Press. http://doi.org/10.1093/acprof:oso/9780198719694.001.0001

Sosa, E. (2021). *Epistemic explanations*. Oxford University Press. http://doi.org/10.1093/oso/9780198856467.001.0001

Turri, J. (2016). Knowledge as achievement, more or less. In M. Fernandez Vargas (Ed.) *Performance epistemology: Foundations and applications* (pp.124–134). Oxford University Press. http://doi.org/10.1093/acprof:oso/9780198746942.003.0008

Zagzebski, L. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge University Press. http://doi.org/10.1017/CBO9781139174763