

*Resistance to Evidence, Epistemic Responsibility,
and Epistemic Vice*

We have seen that available accounts of possessed evidence, evidence one should have possessed, and permissible suspension of judgement struggle to accommodate the phenomenon of evidence resistance. Along the way, we have, in particular, seen that virtue reliabilist accounts of reasons to believe, permissible suspension, and propositional warrant don't do the needed work. At that point, some readers would have already thought that one straightforward explanation of the resistance data is afforded by the competing, virtue responsibilist camp: roughly, on this view, evidence resistance could be conceptualised as a failure to manifest epistemic responsibility in inquiry and/or as a manifestation/indication of epistemic vice. This chapter looks into the credentials of this move. I argue that once we distinguish epistemic virtues and vices proper from mere moral virtues and vices with epistemic content, it transpires that accounting for resistance cases, as well as accounting for epistemic virtue and vice, requires epistemic value-first unpacking.

5.1 Responsible Inquiry and Epistemic Character Traits

A prominent project in responsibilist virtue epistemology (henceforth also VR for short) is to develop 'maps' or 'perspicuous representations' of intellectual virtues, such as intellectual humility, intellectual courage, open-mindedness, and curiosity. The basic idea is to develop empirically grounded characterisations of epistemically admirable or praiseworthy character traits and to use these characterisations as guidance in a kind of regulative epistemology, sometimes with the aim of informing education theory. To take a few examples, Jason Baehr characterises the 'open-minded' person as someone who is 'willing [...] to transcend a default cognitive standpoint in order to [...] take seriously a distinct cognitive standpoint' (2011, 152). Heather Battaly characterises the 'epistemically humble' person as someone who is 'disposed to recognize her own fallibility, and to recognize and value the epistemic abilities of others [...]'

(2014, 194). James Montmarquet characterises the ‘epistemically courageous’ person as someone who has the disposition to ‘persevere in the face of opposition from others (until one is convinced that one is mistaken) [. . .]’ (1993, 23). And Roberts and Wood characterise the ‘epistemically autonomous’ person as someone who has the ‘proper ability to think for herself and not be [. . .] improperly dependent on or influenced by others’ (2007, 259).

In many ways, this work is an exciting new development in epistemology. For example, it highlights possible avenues for widening the traditional epistemological project. In particular, and importantly for my purposes, it might be thought that intellectual virtue and vice bear in crucial ways on the phenomenon of evidence resistance. Look back at our toy cases of resistance: in many of them, one might think, plausible epistemic vices – sexism, racism, laziness, partiality, closed-mindedness, wishful thinking, etc. – are explicitly present and manifest. The particular resistant episode traces back to said vice.

Let me, however, get a few things out of the way: it is crucial to note that, whatever the virtue responsibilist account of the impermissibility of evidence resistance will turn out to be, it cannot be too straightforward a view; that is, it cannot account for said impermissibility simply in terms of the absence of virtue/presence of vice in the resistant cogniser, nor can evidence resistance be simply accounted for in terms of the manifestation of some bad character trait. Agents with excellent epistemic characters can undergo isolated episodes of evidence resistance, and their virtuousness does not make the latter any less epistemically bad. Whatever the view is, then, it needs to be more subtle than this.

Jason Baehr notably champions an early (2009) stance on the issue that treats the data and VR’s theoretical resources with the necessary care. Baehr argues at length against traditional evidentialism by bringing forth two types of cases: cases of failure in inquiry – where S’s body of evidence is problematically affected by lazy or biased inquiry – and cases of bad responses to available evidence due, again, to some variety of bias or wishful thinking. His diagnosis is that traditional evidentialism needs to be supplemented with a good epistemic character condition. In what follows, I will mostly focus my analysis on Baehr’s account, but, since his account features a fairly minimal responsibilist condition, everything I say should apply, *mutatis mutandis*, to all accounts that want to explain cases of evidence resistance as, in one way or another, having to do with character failure. Here is Baehr’s proposal for what it means for a subject S to be justified:

Responsible Evidentialism (RE): *S* is justified in believing *p* at *t* if and only if *S*'s evidence at *t* appears to *S* to support *p*, provided that if *S*'s agency makes a salient contribution to *S*'s evidential situation with respect to *p*, *S* functions (qua agent) in a manner consistent with intellectual virtue. (2009, 16, emphases in original)

Note a few advantages of RE. First, while it does feature a good character requirement on top of traditional evidentialist necessary conditions for justification, is not a very strong character requirement: the justified subject need not manifest intellectual virtues, nor be a virtuous agent to begin with – the epistemically vicious among us can also form justified beliefs, insofar as either their agency is not much involved in arriving at a particular body of evidence at work in a particular belief-formation episode (and here Baehr has in mind as paradigmatic garden-variety automatic perceptual belief formation) or else, should their agency be thus involved, their evidential situation will have been arrived at in a manner consistent with intellectual virtue. In a nutshell, the vicious can form justified beliefs either automatically or via virtuous-like agency. The fact that the account does not require virtue to be present or manifest is a great feat: it does not implausibly exclude the vicious from having a chance at justified beliefs, and, even more crucially, it can, at least at first glance, accommodate one-off cases of evidence resistance, in which the subject's bad epistemic character is not to blame – since the subject is a virtuous yet fallible believer.

I think Baehr's account, in not imposing a particularly demanding epistemic character condition on justification, is probably as good an account as possible for trying to explain the impermissibility of instances of evidence resistance within a virtue responsibilist framework: after all, again, clearly, episodes of resistance do not require vice, nor vice manifestation, but rather can be mere one-off failures of otherwise perfectly virtuous believers.

Nevertheless, in what follows I will argue that Baehr's account does not work for two main reasons, having to do with (1) difficulties in precisifying the character condition and (2) epistemic virtue individuation. If I am right, these two problems together will paint a pessimistic picture for the prospects of accounting for the impermissibility of evidence resistance for the champion of virtue responsibilism: since Baehr's character condition is a very non-demanding one, *mutatis mutandis*, all of Baehr's difficulties will translate to most virtue responsibilist attempts to accommodate these data. In what follows, I will take these worries in turn.

The crucial responsibility requirement in Baehr's account of justification, and the one that is designed to take care of resistance cases, requires

there to be the case that, if *S*'s agency makes a salient contribution to *S*'s evidential situation with respect to *p*, *S* functions (qua agent) in a manner consistent with intellectual virtue. Two things require more spelling out: first, what it is for a virtue to be a genuine intellectual (rather than, for example, moral or prudential) virtue – this is the individuation issue that we will look at in the next section; and second, what it is for a believer to function qua agent in a manner that is consistent with such intellectual virtue.

I will start with the latter: my worry is that it will be difficult to spell out the virtue consistency relation in a manner that is permissive enough for allowing for ubiquity of epistemic justification but also efficient in dealing with resistance cases. To see why, note, once more, that we are fallible creatures: whatever virtues we may have, it will be consistent with them that we fail to manifest them on occasion, and also that we fail to do the right thing on occasion, in spite of manifesting them. No matter how courageous I might be, I may, on occasion, fail to be the first to reach the enemy's lines. I need not do the perfectly courageous thing all of the time in order for my actions to be consistent with having the relevant character virtue. Furthermore, even in manifesting courage, I may end up doing things that look cowardly due to misreading the facts on the ground. Again, we are fallible creatures: in us, virtues can be manifested in instances of failure. If so, however, one-off cases of evidence resistance will be consistent with both having and manifesting intellectual virtue. In what sense of consistency, then, are we to interpret Baehr's character constraint?

One option would be to follow Williamson (forthcominga) and go modal: maybe what it is for a believer to function qua agent in a manner that is consistent with intellectual virtue has to do with doing what the virtuous person would do in the situation at stake. Unfortunately, this will not work either, once more, in virtue of virtue fallibility: after all, the impeccably virtuous person involved in a one-off case of evidence resistance is doing what a virtuous person – namely, themselves – would do in that situation.

As far as I can tell, what the Baehr account requires to solve this problem is character virtue infallibilism: on a view like this, resistance cases are not consistent with epistemic virtue because the virtuous person just would not fail in this way. Unfortunately, though, as soon as we go down this route, we lose both general plausibility for our account of epistemic character traits (we are fallible creatures!) and, even more crucially for our purposes, Baehr's responsibilist evidentialism renders justification a rare commodity, to be held only by the flawless among us.

I started off by saying that I had two worries for Baehr's view, having to do with spelling out consistency with epistemic virtue and, respectively, with individuating epistemic virtue. The remainder of this chapter concerns itself with the latter.

5.2 Content-Individuating Responsibilist Virtues and Vices

Virtue responsibilist epistemologies reveal some of the complex ways in which the disciplines of epistemology and ethics seem to overlap: character matters for normativity, the thought goes – be it moral or epistemic. However, this raises an important question: that of whether, or to what extent, a given intellectual virtue counts as distinctively *epistemic* as opposed to moral or prudential; after all, as we have already seen in Chapter 3, we have strong reason to be sceptical of the possibility of explaining away resistance cases in terms of mere moral failures.

It is here that I think the virtue responsibilist project requires careful handling. After all, the notion of virtue (in the responsibilist sense of a character trait) is familiar first and foremost from ethics. Moreover, it's not clear what relation the virtues at stake stand in to more traditional moral virtues, for example. Indeed, it seems plausible that some of them *just are* paradigm cases of garden-variety moral virtues as opposed to some other kind of virtue. A fully worked out account of the intellectual virtues should be able to clearly address this issue. It seems central to our understanding of the relevance of an investigation into some virtue or another for epistemology. It also seems central to our understanding of the scope of epistemology itself.

In what follows, I do three things: first, I argue that two popular individuation recipes for epistemic virtues and vices don't work. Second, I defend a value-centric way of individuating virtues, including intellectual ones. Third, I argue that this way of individuating virtues gives us reason to be somewhat more cautious in our claims to make progress in the epistemology of evidence resistance (and, more generally, in the epistemology of justification) by investigating intellectual virtues.

How can we individuate distinctively epistemic character traits? In order to answer this question, I'd like to start by looking into the way in which philosophers in the epistemic norms literature individuate epistemic norms and take it from there. After all, virtues and vices are normative; if so, individuating epistemic virtues and vices will benefit from whatever the correct individuation recipe is for epistemic norms. Let's first consider the following proposal:

Content Individuation (CI)

If a norm N concerns epistemic features required for permissible phi-ing, then N is an epistemic norm.

It is safe to say that, if there is such a thing as a received view in the epistemic norms literature concerning what that literature is theorising about in the first place, it is CI.¹ Many philosophers, for instance, when they ask what the epistemic norm for phi-ing is, take themselves to be asking, roughly, how much epistemic warrant one needs for proper phi-ing. For example, here is Jennifer Lackey (arguing that cases of isolated second-hand knowledge show that knowledge is not the epistemic norm of assertion):

It should be emphasized that it is clear that the problem with the agents in the above cases is that it is not *epistemically* appropriate for them to flat-out assert that *p* [...]. One reason this is clear is that the criticism of the agents concerns the *grounds* for their assertions [...]. (2014, 38, emphases in original)

Thus, according to Lackey, insofar as norms concern epistemic grounds, they will be genuine epistemic norms. On a similar note, here is Ishani Maitra (on a view about the nature of assertion that she takes to be widely endorsed):

Assertions are governed by an alethic or an epistemic norm – that is, a norm that specifies that it is appropriate to assert something only if what is asserted is true, or justifiably believed, or certain or known. (2011, 277)

There is very good methodological reason to endorse CI: it is both simple and user-friendly. CI provides a neat and straightforward way to individuate epistemic norms, ensuring that the debate can be framed on common terminological ground.

Most importantly for present purposes, CI might help us to individuate distinctively epistemic virtues. After all, if epistemic norms are individuated by content, as CI has it, since virtues are normative, it is independently plausible that the same goes for epistemic virtues.

Virtue Content Individuation (VCI)

If a virtue V concerns epistemic features required for virtuous phi-ing, then V is an epistemic virtue.

¹ For explicit endorsements, see, for example, Lackey (2011), Maitra (2011), Brown (2012), and Benton (2014). For implicit assumptions to this effect, see Hawthorne (2004), Gerken (2011), and Littlejohn and Turri (2015).

Note that the sorts of virtues that the VR literature has focused on, including open-mindedness, intellectual humility, curiosity, epistemic courageousness, and temperance, clearly concern epistemic features. By VCI, these virtues are epistemic virtues.

In fact, there is yet another way of making the very same point. Once more, virtues are widely taken to be normative. The following is an attractive way of capturing this thought:

Normative Charge of Virtues (NCV)

One's actions and states ought to manifest virtuous character traits.²

By NCV, virtues are associated with oughts. Note also that the oughts they are associated with are typed by the type of virtue in question. For instance, if the virtue to be manifested is a moral virtue, the ought at issue in NCV is a moral ought; if the virtue to be manifested is an aesthetic virtue, the ought at issue in NCV is an aesthetic ought; and so on.

NCV gives us the result that virtues of a certain type are associated with oughts of the same type. Now we may expect to use our recipe for individuating epistemic norms to home in on distinctively epistemic virtues. After all, given that types of virtues are by NCV associated with the corresponding types of norms, if we have a recipe for individuating types of norms, we might be able to use it to individuate the associated virtues as well. In particular, should CI be the right way to individuate epistemic norms, we get the result that, if the virtuous character trait at issue in NCV concerns epistemic features, then the norm is epistemic. And since the sorts of virtues that the VR literature has focused on and that are at issue in the relevant instances of NCV do concern epistemic features, we once more get the result that these virtues are indeed epistemic virtues.

And while it might be thought that this is all entirely as it should be, on reflection, there is reason to think otherwise. To see why, note that both of the above arguments rely on CI for their motivation. The first uses CI to motivate VCI, which is key to that argument, and the second uses CI expressly in the relevant derivation. The trouble is that there is excellent reason to think that CI is false. More specifically, as I will argue

² NCV requires some noteworthy qualifications. First, it's quite plausible that it needs to be restricted to *relevant* actions (i.e. those that can be performed in a way that manifests the relevant virtue). For instance, while it is plausible that I ought to be generous in my dealings with my friends, it does not seem to be the case that I ought to be generous in sitting down. Second, the oughts at issue here are defeasible. For instance, if I have overriding reason not to be generous in my dealings with my friends (perhaps because someone threatens to kill me if I am), it is not the case that I ought to be thus generous, at least not all things considered. Since these qualifications are of little consequence for the purposes of this chapter, in what follows we will take them as read.

momentarily, first, CI does not generalise in the right way, and, second, it is extensionally inadequate.

To see this why CI doesn't generalise as it should, let's get the generalisation on the table. Here is what it looks like:

Generalised Content Individuation (GCI)

If a norm N concerns features of type T required for permissible phi-ing, then N is a norm of type T.³

Unfortunately, there is excellent reason to think that GCI is false. Consider, for instance, traffic norms: driving one's car within city bounds will surely be subject to whatever the local traffic regulations have to say about it. Say that the relevant traffic norm forbids one from driving faster than 30 miles per hour. However, imagine that a terrorist group has placed a bomb in the centre of town and you are the only one able to diffuse it. In order to get there in time, you must break the traffic norm and drive at 40 miles per hour. Clearly, the latter (moral, prudential) requirement overrides the traffic norm and renders driving at 40 miles per hour the all-things-considered proper thing to do. In this case, the bomb threat drives the all-things-considered proper speed up to 40 miles per hour. The moral requirement, in this case, has traffic-related content: it regulates the morally (and all-things-considered) appropriate speed. If that is the case, however, it looks as though, just because a norm has traffic-related content – just because it regulates the appropriate speed – it need not follow that it is a traffic norm. GCI fails for traffic norms.

What's more, the traffic case is hardly isolated. Similar examples can be construed for many types of normativity. It can be prudentially or morally appropriate to drive faster or slower, to have a better or a worse grade point average, to wear a lighter or a darker dress at a funeral, or to speak louder or more quietly. Just because a norm regulates the appropriate degree of a traffic-related feature (i.e. speed), it need not follow that it is a traffic norm. Just because a norm regulates a fashion-related feature (i.e. colour of dress), it need not follow that it is a fashion norm. And so on. GCI is false.

What about CI? Given that GCI does not hold, there is, of course, every reason to believe that the same goes for epistemic norms: just because a

³ There are many problems with CI that I will not discuss here. One such problem, for instance, is: how is the typing supposed to work for biconditional norms, such as 'hit the emergency brake if and only if you see someone stuck on the escalator'? Is this a norm for permissible brake-hitting? If so, it has features of accident-witnessing as permissibility requirement. Or is it a norm for permissible accident-witnessing? If so, it has features of brake-hitting as permissibility requirements. Is it both, and does it have both as permissibility requirements and would thus be classed as a norm of two types?

norm regulates the appropriate degree of an epistemic feature, it need not follow that it is an epistemic norm. It may indeed happen to be an epistemic norm; but it may also be the case that it is a norm of a different nature – say, a prudential or moral norm – with epistemic content. It may be a norm of some other kind that simply happens to regulate the (morally, prudentially, etc.) proper degree of an epistemic feature. To see this, consider the following examples:

SING. One must sing only songs one knows.

JUMP. One must not jump in lakes unless one knows how to swim.

ASSERT. One must assert only what one knows.

BELIEVE. One must believe only what one has sufficient evidence for.

All four norms have epistemic content. Accordingly, CI will predict that they are all distinctively epistemic norms. However, that is intuitively implausible; while ASSERT and BELIEVE are plausibly epistemic norms, SING and JUMP are not. Instead, SING is plausibly an aesthetic and/or prudential norm, and JUMP is plausibly a prudential norm, although both have epistemic content.

What transpires, then, is that just because a norm has epistemic content, it does not follow that it is an epistemic norm. Just because it concerns epistemic features (i.e. what epistemic position one needs to be in in order to permissibly phi), it does not follow that it concerns distinctively epistemic permissibility. It can also regulate an epistemic feature required for prudentially, morally, aesthetically, etc., permissible phi-ing. What CI allows us to home in on are norms with epistemic content rather than genuinely epistemic norms.

If all of this is right, however – that is, if CI fails as a criterion for individuating genuinely epistemic norms – we should also be suspicious of the ways of individuating epistemic virtues that it serves to motivate, since virtues are normative. For instance, since CI fails to distinguish between epistemic norms and mere (moral, prudential, aesthetic, etc.) norms with epistemic content, we may legitimately wonder whether VCI correspondingly fails to distinguish between genuinely epistemic virtues and virtues of different stripes (e.g. moral, prudential, etc.) with epistemic content. And, of course, this worry is only more pressing for the individuating recipe that relies on CI directly.

And, indeed, on closer inspection, one encounters such cases in the literature. Consider, for instance, Roberts and Wood's example of the intellectual virtue of epistemic temperance:

In Toni Morrison's novel *Beloved*, the slave-narrator comments on a slave-owner who was a cut about the average that he looked away when the slave

women were nursing their infants. This intentional foregoing of acquaintance expresses respect for the women's privacy and a sense of the limits that human proprieties set to appropriate knowledge [. . .]. To be an indiscriminate ogler is a trait of bad intellectual character, a failure of discipline of the will to know. (2007, 175)

Although this virtue is associated with an epistemic feature, it is not at all clear that it qualifies as a genuinely epistemic virtue rather than a virtue of some other denomination that has epistemic content. After all, by the lights of the authors themselves, this particular virtue expresses '*respect for the women's privacy* and a sense of the limits that *human proprieties* set to appropriate knowledge' (Roberts and Wood 2007, 175, emphases added). As such, it is more natural to read it as a kind of moral virtue, albeit one that has something to do with the epistemic domain – a moral virtue with epistemic content. Indeed, epistemic temperance through the foregoing of acquaintance looks like a paradigmatic example of a moral virtue, the exercise of which *limits* the gathering of information and knowledge. And from a purely epistemic perspective, at any rate, this seems like precisely the opposite of what the exercise of an epistemic virtue would be apt to limit.

Let us take stock. The challenge we aim to meet on behalf of VR is to find a criterion for normative typing that will allow VR to remain within the boundaries of epistemology proper. This will help us determine whether and to what extent the intellectual virtues being mapped by VR theorists are genuinely epistemic virtues and thus to be usefully employed for an account of epistemic justification and, conversely, of the epistemic impermissibility of evidence resistance. On the first view we looked at, epistemic norms and, by extension, epistemic virtues are concerned with distinctively epistemic features; that is to say, epistemic norms and virtues, as such, are norms and virtues with epistemic content. We have seen, however, that CI for epistemic norms and virtues runs into trouble on at least two counts. First, it doesn't generalise to other normative domains in the way it ought to. And second, on reflection, there is reason to think that it even makes counterintuitive predictions in the epistemic domain.

In what follows, I will look at a different proposal for individuating epistemic virtues and vices: by the psychological reality that they describe.

5.3 Psychological Vice Individuation

In her recent book, Alessandra Tanesini (2021) offers a systematic, comprehensive, thoroughly empirically informed picture of the nature and normativity of epistemic vice. The book is also aiming to carve out new

methodological space in the epistemology of vice beyond the reliabilist/responsibilist divide in virtue epistemology. Tanesini sees her project to be one of ‘autonomous’ epistemology. The ground of epistemic vice on this view is neither responsibility nor reliability: it lies with psychological reality. If successful, the view offers an alternative individuation recipe for epistemic virtues and vices and thereby an alternative, psychology-of-vice-based explanation of what is going on in resistance cases.

On Tanesini’s account, epistemic vices are taken to be essentially sourced in attitudes towards the self: fatalism, self-satisfaction, narcissistic infatuation, and self-abasement. The thought, roughly, is that some people have a self-infatuated stance towards their intellectual qualities, which they therefore assess as superlative without pausing to consider their true epistemic worth. Others, in contrast, adopt a self-abasing and negative stance towards their intellectual abilities. Consequently, they become ashamed of their intellectual qualities, which they perceive to be extremely limited.

Fatalism, self-satisfaction, narcissistic infatuation, and self-abasement are attitudes towards the self that ground epistemic vices of self-assessment. These are exemplified by those who do not have the measure of their intellectual abilities because they assess their epistemic worth using the wrong unit of measurement:

[F]or instance, those who are motivated to self-enhance tend to compare themselves for how they differ from less capable individuals so as to find further confirmation of their excellence. I use the metaphor of measuring oneself by the wrong unit to describe this phenomenon of biased selection of the yardstick (as represented by the relative ability of another person or group) by which to evaluate one’s own performance. (Tanesini 2021, 15)

Since these evaluations are crucial in the setting of realistic epistemic goals, in the choices of methods and strategies to adopt in inquiry, and in the process of epistemic self-improvement, those whose self-assessments are thus misguided are unlikely in ordinary circumstances to excel in their epistemic pursuits.

That being said, on Tanesini’s account, the presence of vice is essentially connected to its being sourced in self-assessments that employ the wrong unit of measurement and not to the falsity of its constitutive doxastic attitudes, nor its (likely) unfortunate epistemic consequences.

In what follows, I take issue with Tanesini’s vice internalism: epistemic vices are vices, I argue, only if externalistically individuated. I’ll focus on what Tanesini calls ‘vices of self-satisfaction’, like narcissism and superbia, because they are the ones that are most paradigmatically relevant to

episodes of evidence resistance – although evidence resistance in virtue of unwarranted lack of trust in one’s epistemic abilities is also perfectly possible, and likely often encountered in historically marginalised groups. Nothing hinges on this though – the worries I outline generalise neatly to the entire framework.

The presence of vice, on Tanesini’s account, is independent of the accuracy of the vice-constitutive beliefs about oneself – I might be right that I am the smartest person in the world, but if this belief is sourced in bad self-assessment processes, it has the disposition to be vice-constitutive nevertheless. Furthermore, it may be that my narcissism is, *de facto*, extremely reliable, in that it mostly outputs true beliefs: it remains an epistemic vice on Tanesini’s view nevertheless. For this, Tanesini takes vices to supervene on subjects’ psychologies (i.e. on particular attitudes towards the self). Here is what Tanesini thinks about individuals who are in the grips of vices of self-satisfaction:

[These] individuals adopt a self-satisfied stance towards what they regard as their intellectual strengths. They believe that a great number of their intellectual features are impressive. These individuals are also often averse to working towards improvement. They adopt this stance because they believe that they are already great and thus have no need to improve. Hence, their mindset is [...] fixed since they judge themselves to be naturally talented and thus capable of effortless success. (Tanesini 2021, 15)

Vices of self-satisfaction will be individuated by the corresponding attitude. Accuracy doesn’t matter: my self-satisfied beliefs about myself may well be (luckily) true. What’s crucial to vice presence is that they are not sourced in/based on evidence, but rather are sourced in/based on mistaken self-directed attitudes (self-admiration, self-defence, etc). Here is Tanesini:

We should expect narcissistic and self-satisfied self-evaluations to be off the mark by underestimating shortcomings and overestimating strengths. However, in unusual circumstances, it is possible that such individuals may have impressive intellectual strengths and through sheer luck their self-assessments may prove to be largely accurate. [...] The person with narcissistic tendencies, for example, is disposed to bullshit even though he holds true beliefs about his capacities. What makes his claims about the self, among other things, bullshit is that he does not care whether they are true. (Tanesini 2021, 15)

One can distinguish between two truth-independence claims in Tanesini’s view: first, the presence of the vice is compatible with it being (mostly) constituted by true beliefs about oneself. Second, the presence of the vice is also independent of its epistemic consequences: it might be that, in virtue

of holding this attitude towards myself, I am highly successful in inquiry – in the sense that I am more likely to discover truths and avoid falsehoods. Here they are, just for simplicity of use:

Constitutive Truth Independence (CTI): epistemic vices are compatible with the truth of their constitutive beliefs.

Consequence Valence Independence (CVI): vices are compatible with a positive epistemic valence of their epistemic consequences.

Furthermore, Tanesini proposes that epistemic vice fully supervenes on one's psychology – both metaphysically and normatively. Let's formulate this claim for ease of use as well:

Tanesini's Vice Internalism (TVI): epistemic vices supervene on the subject's psychological attitudes.

In what follows, I argue for three claims: first, that CTI and CVI do not suffice to support TVI, nor any other internalism about vice. Rather, CTI and CVI merely reinforce the already popular view that a simple, *de facto* reliabilist view of epistemic normativity is wrong. Compatibly, I argue, epistemic vices might still require externalistic individuation of a different flavour.

Second, I argue that, indeed, epistemic vice will need a hook in the world outside of one's skull if it is to be plausibly epistemically normatively problematic.

Finally, I consider a comeback on behalf of vice internalism: even though, if I'm right, CTI and CVI do not do the analytic work that the vice internalist needs them to do, they might still be useful for doing the social psychological work – that is, while vice internalism need not follow from CTI and CVI, it might still be the case that, in the world we inhabit, and given the kinds of creatures that we are, it is paradigmatically the case that vices will survive truth and reliability. I will put forth some worries for this claim.

Here it goes: I think Tanesini is right about CTI and CVI. Plausibly, epistemic vice can survive accuracy of constitutive beliefs and *de facto* reliability. That's hardly surprising, one would think: we already know from research on externalist theories of justification and the norm of belief that (1) plausibly, there's more to attributively good belief than truth, and (2) blunt, *de facto* reliabilism just won't do as a theory of epistemic justification (e.g. Norman the Clairvoyant⁴ has taught us as much). If so (i.e. if positive normative properties of beliefs don't supervene on either

⁴ Here is the famous case by Laurence Bonjour: "Norman, under certain conditions, which usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter.

truth or de facto reliability), we should also expect that negative epistemic dispositional properties need not imply the lack thereof.

Luckily true beliefs based on, for example, coin tosses don't make for (attributively) good beliefs: they don't make for good tokens of their type. If so, the fact that a particular attitude towards the self is grounded in true beliefs need not suggest it's not an (attributively) bad attitude. Vices can be grounded in true beliefs.

Wishful thinking is not a proper way to form beliefs, nor does it lead to good beliefs, even if it's reliable. If so, just because a way to form beliefs reliably results in true beliefs, it does not follow that it is not a bad way to form beliefs. Vices can be reliable.

That being said, CTI and CVI do not imply vice internalism, more than, for example, the knowledge norm of belief is an internalist norm or normal worlds reliabilism is an internalist view of justification. Champions of both of these views agree, respectively, that true beliefs are not good tokens of their type, and that de facto reliability need not imply that a method of belief formation is a good way to form beliefs. Indeed, any externalism about epistemic normativity in general that denies these two claims will be perfectly compatible with CTI and CVI, while, at the same time, denying that epistemic normative categories are internalistically individuated. Knowledge normers, for instance, are free to claim that non-knowledgeable (albeit true) beliefs can constitute vice, and that dispositions or attitudes that reliably lead to truths – but not knowledge – can be epistemic vices. Non-de-facto reliabilists will agree that vices can be de facto reliable: they will just hold that they are incompatible with normal worlds reliability, or proper function, and so on.

CTI and CVI do not imply vice internalism and thereby fail to offer support to TVI. But is TVI independently plausibly true? If yes, maybe CTI and CVI are mere symptoms of this reality.

I don't think so: vice internalism is false. To see this, let's ask the following question: what is it that makes, for example, narcissism and superbia into vices? On Tanesini's view, recall, it is the biased selection of the measuring unit used to measure oneself that explains the problematic nature of vices of self-assessment. But what is wrong with biased selection? What explains its negative epistemic valence, in virtue of which it grounds

He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact the belief is true and results from his clairvoyant power under circumstances in which it is completely reliable" (BonJour 1985, 41).

vice? Here are a few answers that are not available to internalism: self-measurements involving biased unit selection are epistemically problematic because they cannot lead to knowledge, because they don't have a tendency to get it right in normal conditions, or in normal worlds, because they were selected for biological rather than epistemic success, etc. All of these normative grounds are not available to the vice internalist because they lie outside of the skull's limits. In a nutshell, then, when Tanesini talks of vice being grounded in self-measurements that employ the wrong measuring unit, what is it that explains the relevant wrongness? More precisely, what is it, within the subject's skull, that explains it?

I conjecture that vice internalism will have just as hard a time answering this question as general internalism about epistemic normativity has historically had: what is wrong with beliefs based on wishful thinking? Well, they are formed via a bad belief-forming process. Why is wishful thinking bad? Because it's not the right kind of process for forming beliefs. What is the right kind of process? Short of offering an ad hoc list, notoriously the answer will have to appeal to something outside the believer's skull.⁵

Maybe TVI was never intended as a dismantling analysis of epistemic vice (i.e. as offering necessary and sufficient conditions for its instantiation) but rather as a paradigm case analysis thereof: maybe, that is, the claim is rather that, paradigmatically, vices are independent of the truth of their constitutive beliefs, as well as of the valence of their epistemic consequences.

I worry about the plausibility of this take on the view: first, it seems to me as though it is both psychologically and epistemologically implausible that the exercise of epistemic vice will often and easily co-exist with accuracy. Consider: I believe I'm very good at maths due to self-admiration alone. Next, I do some maths. As it turns out, I'm getting things right all of the time. I find it implausible, at this juncture, to think that the inductive evidence is not (at least part) of the basis of my belief. Implausibility is not impossibility, of course: it may be that I totally ignore this inductive evidence. I submit, however, that it is psychologically implausible, so it will not serve as a paradigm case analysis of the phenomenon we are looking at.

Second, the problem generalises: if, whenever I believe (from self-mismeasuring) that I am good at phi-ing, I phi and get inductive evidence

⁵ In recent work in reply to my worries, Tanesini concedes that the way forward for vice individuation will have to appeal to truth – in particular, that what makes epistemic vices into epistemic vices regards, roughly, their being sourced in a non-truth-orientated psychological tendency.

that I am good at phi-ing, it will be hard to see how it is that I am still instantiating narcissism or superbia rather than a merely justified belief that I'm great, sourced in a solid inductive basis.

In sum, Tanesini's rich account teaches us a lot about the psychology and epistemology of vice: de facto lack of reliability does not matter, and false constitutive beliefs are not needed for vice. Compatibly, though, I have argued, Tanesini's self-mismeasuring attitudes need externalist normative grounding: the 'mis' in the 'mismeasure of the self' can't be restricted to the limits of the skull. Epistemic normativity – be it of virtue or of vice – is externalist normativity: it has to do with epistemic values that lie outside our skulls, such as truth and knowledge.

In what follows, I will look at a competing proposal for individuating genuinely epistemic normative notions: by the values they are associated with. This option, I argue, is not only theoretically superior, but it is also extensionally adequate. However, alas, with this improved method to hand, many of the virtues discussed in the VR literature will fail to qualify as genuinely epistemic virtues. Rather, as ever, they are moral virtues with epistemic content. If so, the prospects of employing them in unpacking the epistemic impermissibility of resistance to evidence are dim.

5.4 Value-Based Vice Individuation

This section looks at a theory-neutral individuation recipe for epistemic virtues and vices based on a widely accepted claim concerning the relation between the axiological and the deontic.

The theory of normativity has a theory-neutral answer to the question of normative individuation ready to hand: norms can be typed by the type of *good* they are associated with.

Value Individuation (VI)

A norm N is of type T if and only if N is associated with goods of type T.

According to VI, prudential norms are associated with prudential goods, moral norms are associated with moral goods, and so on. All normative domains have goods (values) that are central to them, in virtue of the kind of normative domains they are: survival is a prudential good; promise-keeping is a moral good; politeness is a social good; beauty is an aesthetic good; money is a financial good. Similarly, if etymology is any guide to what normative domain 'the epistemic' is supposed to refer to, knowledge is an epistemic good. But given that this is so, philosophers cannot just stipulate that, starting tomorrow, they will use 'moral' to refer to a type of

normative domain that does not care about promise-keeping but does care about money, or ‘financial’ to refer to a domain of which the chief good is safe driving. Similarly, it would be odd to count wealth and having short nails amongst epistemic goods. We have some independent hold on the relevant types of goods. Given that this is so, VI holds out the hope of offering a helpful way of typing norms by the goods associated with them.

This is, of course, still rather vague, and in particular the association relation at issue requires spelling out. After all, one way in which a norm can be associated with a particular good is by requiring more or less of that good. This, however, will put us back in the same trouble we faced with CI: just because a norm requires me to know how to swim before jumping into lakes and is thereby associated with an epistemic good (i.e. knowledge), it does not follow that it is an epistemic norm.

On the view I favour, the association relation stands for one or another direction of explanation: either the goods explain the norm or the other way around. To see how this goes, it may be worth noting that VI is value-theoretically neutral in the sense that it does not come with any substantive commitments about the relation between the axiological and the deontic. That is because the association claim between norms and goals of the same type does not *imply* any particular direction of explanation. As a result, it is compatible with both of the two leading views about the relationship between the axiological and the deontic. Teleologists (e.g. Sidgwick 1907, Slote 1989, Moore 1993) explain the ‘ought’ in terms of the ‘good’; they claim that the norm of type X is there to guide us in reaching the good of type X. In contrast, deontologists (e.g. Ewing 1947, Scanlon 1998, Rabinowicz and Rønnow-Rasmussen 2004) reverse the order of explanation: according to ‘fitting attitude’ accounts of value, for example, the goods of type X are only valuable to begin with because the norm of type X gives us reasons to favour them. Crucially, in either case, the mere *association* claim at issue in VI holds.⁶

Since we are interested specifically in the epistemic domain, I want to take a moment to take a quick look at how the association claim may be unpacked for distinctively epistemic norms within both a teleological and a deontological framework. Here is Peter Graham for an explicit statement of the teleological direction of explanation:

Epistemic norms in this sense govern what we ought to say, do or think from an epistemic point of view, from the point of view of promoting true belief and avoiding error. (Graham 2015, 247)

⁶ For a good general overview of the relevant literature in value theory, see, for instance, Schroeder (2012).

Here is Kurt Sylvan for a statement of the deontological direction of VI:

[C]entral epistemic properties like justification, coherence, and substantive rationality derive non-instrumental epistemic value in virtue of the fact that they manifest different *epistemically fitting ways of valuing accuracy*. (Sylvan 2018, 389, emphasis in original)

Note also that VI makes plausible predictions about epistemic norms and norms with epistemic content. For instance, BELIEVE is correctly characterised as an epistemic norm. After all, it is a norm that is associated with the epistemic good of true belief. The easiest way to see this is by looking at the teleological direction of the association claim. Believing what one has good evidence for believing is a good means to true belief; it promotes true belief. Consider, by way of contrast, SING. Rather than being associated with distinctively epistemic goods (e.g. promoting true belief), it is associated with an aesthetic good. Again, in a teleological framework, singing only songs one knows is a way of promoting beauty. As a result, VI classifies it, again correctly, as an aesthetic norm, albeit one with epistemic content.

Now, just as we took CI to motivate the parallel content-based recipe for individuating epistemic virtues (due to the normativity of virtues), we may take VI to motivate the obvious value-based recipe for so doing. Here it goes:

Virtue Value Individuation (VVI)

A virtue V is a distinctively epistemic virtue if and only if V is associated with epistemic goods.

One might initially worry whether VVI preserves VI's value-theoretic neutrality. After all, given that virtues are associated with epistemic goods, one may get the impression that VVI commits me to a distinctively teleological value theory. Fortunately, the answer to this question is 'no'. After all, it may be that, in accordance with VI, the axiological is analysed in terms of the deontic. If so, it may still be that norms explain both goods and virtues. By the same token, VVI does preserve the value-theoretic neutrality after all.⁷

Importantly, there is reason to believe that several virtue responsibilists are attracted to VVI. In several places, leading VR theorists hint at something along VVI lines. Let's start with what Roberts and Wood themselves say about this issue:

⁷ Again, we can also individuate epistemic virtues by VI and NCV. Since it's easy enough to see that the results will be the same, I will not go into detail about this here.

The difference between our study and a study in virtue ethics is simply that we are interested in the relations between the virtues and the intellectual goods. (Roberts and Wood 2007, 60)

Here is Montmarquet on epistemic virtues:

What I want to suggest, then, as a first approximation, is that the epistemic virtues are those personal qualities (or qualities of character) that are conducive to the discovery of truth and the avoidance of error. (Montmarquet 1993, 20)

Also, Jason Baehr's view about what distinguishes epistemic virtues from other types of virtues seems very similar to VVI;

While structurally similar to moral virtues, they are also distinct from what we ordinarily think of as moral virtues on account of aiming at distinctively epistemic goods like truth, knowledge, and understanding. (Baehr 2017, 96)

Last but not least, here is Linda Zagzebski:

I will argue that truth conduciveness is an essential component of intellectual virtues and I will attempt to ground these virtues in the motivation for knowledge. (Zagzebski 1996, 13)

Unfortunately, if VVI is correct, many of the purportedly epistemic virtues discussed in the literature turn out to be moral virtues with epistemic content rather than genuinely epistemic virtues. By the same token, these virtues will not be relevant to epistemic justification, nor will they be helpful in accounting for the *epistemic* impermissibility of resistance cases.

Let us start with an easy case. Consider, again, the virtue of epistemic temperance discussed by Roberts and Wood. Intuitively, Roberts and Wood are right: the slave-owner who looked away when the slave women were nursing their infants was, to this extent, manifesting virtue.⁸ As Roberts and Wood well put it, 'To be an indiscriminate ogler is a trait of bad intellectual character, a failure of discipline of the will to know' (2007, 175). However, since we have seen that typing norms and virtues by content will not do, the question is: when Roberts and Wood talk about 'bad intellectual character', what type of badness is that? Is it genuinely epistemic (i.e. associated in the relevant way with epistemic goods) or rather badness of a different sort (i.e. associated with different types of goods), albeit badness that has epistemic content? It is hard to

⁸ To be sure, the claim here is that the slave-owner manifests virtue *in this specific regard*, and not, of course, that the slave-owner in any way manifests virtue by being a slave-owner.

deny that it must be the latter. After all, what seems to be going on here is that a properly epistemic virtue, ‘the love of knowledge’ (which is, in the relevant way – either teleologically or deontologically – associated with an epistemic good, in this case knowledge), is being overridden (‘disciplined’) by another virtue, namely the virtue of ‘epistemic temperance’. Plausibly, though, the latter is precisely not associated with epistemic goods in the relevant sense. After all, on neither direction of explanation does it plausibly stand in an association relation with epistemic goods. Neither is it conducive to their acquisition, nor does it give one reasons to favour them. Rather, what seems to be the case is that epistemic temperance stands in the relevant relation with moral goods, such as respect for privacy or discretion. Indeed, depending on the details, one might conceive of the Roberts and Wood case as a case of all-things-considered permissible evidence resistance due to moral considerations overriding epistemic considerations.

If all this is so, then, the Roberts and Wood account is guilty of normative ambiguation: it is not the case that the slave-owner is an epistemically virtuous person in virtue of tempering his will to know in the relevant case; rather, the slave-owner manifests a moral virtue with epistemic content by tempering his will to know.

On a similar note, consider Heather Battaly’s discussion of epistemic temperance as well as the corresponding vice that she calls ‘epistemic self-indulgence’:

[T]he passions and actions associated with these traits are epistemic rather than physical, and include wanting, consuming, and enjoying beliefs, knowledge, and belief-forming practices. I argue that the epistemically temperate person desires, consumes, and enjoys only appropriate epistemic objects, only at appropriate times, and only in appropriate amounts. The epistemically self-indulgent person, however, [...] desires, consumes, and enjoys epistemic objects at inappropriate times (e.g., while having sex with his partner); or desires, consumes, and enjoys epistemic objects too much (thus preventing him from pursuing other things of value). (Battaly 2010, 232)

Note that the goods secured by ‘epistemic temperance’ are again social or prudential, precisely at the *expense* of epistemic goods. According to Battaly herself, the epistemically temperate person sacrifices the consumption of epistemic goods for the sake of ‘pursuing other things of value’. Conversely, it is not clear why ‘epistemic self-indulgence’ should be considered an epistemic vice. After all, on both directions of explanation it is strongly associated with epistemic goods rather than bads. Thus, by the lights of VVI, it seems that Battaly’s discussion also manifests the familiar mistake.

Importantly, this is not to say that a case cannot be made for the claim that epistemic temperance, in a particular context, can be an epistemic virtue proper, or that epistemic self-indulgence could be an epistemic vice – again, depending on context (i.e. depending on whether the relevant context associates these virtues with epistemic goods). Here is, very roughly, how an argument to this effect might go: consider someone who spends all of their epistemic resources on acquisitioning disparate items of knowledge about completely unrelated topics. Perhaps this person should temper their will for a high quantity of epistemic goods in favour of the quality of the epistemic goods. That is, perhaps a project of diving more deeply into a particular domain, in order to achieve understanding thereof, would be more *epistemically* worthwhile than one of simply acquisitioning disparate items of knowledge. At the very least, this question is well worth investigating.

Nenad Miscevic's (2016) discussion of curiosity (which he takes to be the basic epistemic virtue) raises similar concerns. Miscevic carefully draws our attention to the fact that what he means by the 'virtue of curiosity' is something that excludes *nosiness*. Indeed, he stipulates the term 'curiosity+' to pick out a kind of curiosity that is not too strong. Curiosity+ is just strong enough to secure some epistemic goods without thereby giving rise to bad moral side effects. Here is Miscevic:

Curiosity, when a virtue, call it curiosity+ includes knowledge of appropriateness, and motivation for appropriate exercise. Curiosity–, the vicious inquisitiveness, is not really curiosity. [...] We would then in general have two sub-species of cognitive intrinsic desire to know, intrinsic curiosity+, and curiosity–, the bad intrinsic curiosity. The first is truly a virtue, the second is not [...]. Typical [...] cases of low-level object curiosity [are] aiming at private and intimate matters of others (nosiness), or [are] connected to morally problematic goals or consequences. (Miscevic 2016, 150–152)

Conduciveness to moral goods, though, does not bear on whether or to what extent a particular character trait is an epistemic virtue or not. Indeed, it is unclear why nosiness is not legitimately understood as an epistemic virtue, precisely insofar as it is associated with epistemic goods, and despite the fact that it plausibly fails to qualify as a moral virtue. Consider a teleological direction of explanation: nosiness, as morally indecent as it might be, will definitely be one good character trait to have for securing epistemic goods, such as knowledge or true beliefs. It is, then, unclear why we should endorse Miscevic's distinction between good curiosity (curiosity+) and bad curiosity (curiosity–) rather than allowing that curiosity simpliciter is an epistemic virtue, albeit one that may or may

not also count as a moral virtue depending on whether or not it brings about bad moral consequences.

One question that the VR champion might rightly ask at this point is the following: is it not the case that, on the Aristotelian model, only *proper* exercise of a 'virtue' counts as an instance of genuine virtue, whereas *improper* exercise does not really count as an instance of virtue? Here is Miscevic on this issue, referencing a discussion by Philippa Foot:

The courage of a whistleblower is courage, the bravery of an SS-officer is not. [...] Similarly, nosiness is not really curiosity, at best it is pseudo-curiosity [...]. Curiosity, when a virtue, [...], includes knowledge of appropriateness, and motivation for appropriate exercise. Curiosity—, the vicious inquisitiveness, is not really curiosity. (Miscevic 2016, 152)

The problem with this way to go is that 'proper exercise' is a normative notion. As such, it requires typing itself. There is such a thing as the epistemically proper exercise of a virtue, the morally proper exercise of a virtue, and so on. If VVI is right, epistemically proper exercise of curiosity will represent an epistemically virtuous exercise. Once again, though, it is not clear why episodes of nosiness will not count as epistemically proper exercises of curiosity – after all, they are apt to secure epistemic goods.

Instances of normative ambiguation resulting from a lack of a clear individuation recipe for virtues are ubiquitous in the VR literature and beyond. It would take too much space to point them all out in this chapter. However, I need not do so for present purposes. Rather, the ambition here is to draw attention to a potential problem sourced in this lack for the project at hand: that of explaining the impermissibility of evidence resistance by reference to genuinely epistemic character traits.

Importantly for our project, by the lights of VVI, even virtues that have been at the very heart of the VR literature, and that might be taken to paradigmatically explain the impermissibility of resistance to evidence – such as open-mindedness or intellectual humility – deserve closer scrutiny.

To see why this is the case, let's close with a brief examination of open-mindedness. Of course, at least at first glance, open-mindedness comes across as a paradigmatic epistemic virtue, and one that is paradigmatically missing in cases of evidence resistance: being open-minded opens one towards properly appreciating the views of others and thus properly assessing available evidence. Note, however, that whether a given case really qualifies as *proper* appreciation may depend on (1) the direction of explanation for unpacking VVI and (2) contextual features. To see this, consider open-mindedness in a teleological framework according to which

virtues count as genuinely epistemic only insofar as they are conducive to epistemic goods. Now, plausibly, for most people of average epistemic endowment and living in an average epistemic environment, being open-minded is indeed conducive to epistemic goods. After all, when undertaking intellectual projects most of us are likely to encounter better (as in epistemically better) ideas/views, etc., than our own. Being receptive to these other ideas/views, etc., will thus be conducive to epistemic improvement.

In contrast, however, it is not clear that the same is the case when we move further up the scale of epistemic endowment. For a being that is well above average in cognitive ability, open-mindedness will often be conducive to epistemic loss. After all, it may be conducive to abandoning perfectly fine beliefs in the light of misleading evidence. An open-minded mathematical genius shouldn't abandon worthwhile beliefs in the light of less qualified testimony. Similarly, an expert in vaccines or climate change should not update on layman sceptical testimony. And so on.

If this is the case, whether open-mindedness is an epistemic virtue is a highly contextual matter.

Jeremy Fantl (2018) makes a similar point against the *tout court* epistemic goodness of open-mindedness, but in a more interesting way. He starts off with a fairly minimal account of open-mindedness, on which you are open-minded towards an argument if and only if (1) affective factors do not dispose you against being persuaded by the argument, (2) you are not disposed to unreasonably violate any procedural norms in your response to the argument, and (3) you are willing to be significantly persuaded conditional on spending significant time with the argument, finding the steps compelling, and being unable to locate a flaw. He then goes on to argue that open-mindedness is not always a good thing: in particular, he argues, there are many situations in which you know that a relevant counterargument is misleading whether or not you have spent significant time with the argument, found each step compelling, and been unable to expose a flaw. In cases like these, he argues, you should not let yourself be convinced by the argument. Such a dogmatism can be rational, according to Fantl, since often the best explanation of your situation is that your well-supported belief is correct and a clever individual has simply come up with a misleading counterargument (2018, 34). Furthermore, Fantl argues, such closed-minded dogmatism is a manifestation of intellectual humility – since you know of yourself to be fallible at identifying flaws in misleading arguments. This is why, according to Fantl, your knowledge can survive coming across an apparently flawless counterargument.

Of course, independently of its reliability in generating epistemic goods, open-mindedness is plausibly accurately characterised as a moral virtue

with epistemic content. After all, being an open-minded person quite plausibly means treating other human beings as worthwhile epistemic sources, independently of whether they actually are reliable sources. This, arguably, is an instance of acting in accordance with a more general moral law requiring us to respect humanity.

Similarly, and for similar reasons, intellectual humility is more properly thought of as a moral virtue with epistemic content that, contextually, may become an epistemic virtue proper in cases where the context is such that it is associated with epistemic goods. In other contexts, to the contrary, epistemic courage will be the genuinely epistemically virtuous character trait to have and manifest (see Ichikawa (forthcoming) for an excellent book-length treatment of epistemic courage in relation to evidence resistance and positive epistemology more broadly).

Virtue responsibilist character traits are epistemic virtues only insofar as they are associated with epistemic goods. If this is so, however, appealing to character traits cannot constitute the normative bedrock for any account that tries to explain the epistemic impermissibility of evidence resistance: epistemic values will do the work. Values will come before virtues in explaining resistance data (as well as in accounting for other interesting epistemic notions and phenomena). The view developed in this book will do just that: explain the impermissibility of evidence resistance by appeal to knowledge and its availability.

5.5 Conclusion

This chapter investigated the virtue responsibilist camp for resources to explain the epistemic impermissibility of evidence resistance. I have argued that, if plausible at all, a responsibilist account of these cases better not be too strong (i.e. had better not explain resistance as absence/lack of manifestation of virtue or presence/manifestation of vice) since, in slogan form, good people can also believe bad things (Levy 2021): the epistemically virtuous are fallible, as they can be involved in one-off resistance cases. I have then looked at what I take to be the most successful account on the market in keeping the character condition permissive: Jason Baehr's responsibilist evidentialism. I have argued that the notion of virtue consistency in the account does not afford plausible unpackings that explain evidence resistance. Furthermore, I have shown that individuating epistemic character traits requires an appeal to epistemic values. If this is so, intellectual character traits don't do the grounding normative work in explaining what goes wrong in resistance cases: epistemic values do.

