

# Corrupted

## *An Essay on Intellectual Character and Epistemic Vice*

Taylor Roger Charlie Matthews

Thesis submitted in partial fulfilment of the requirements for the degree of  
Doctor of Philosophy

September 2023

Department of Philosophy



**University of  
Nottingham**  
UK | CHINA | MALAYSIA

# Acknowledgements

Of all the words contained in this thesis, the easiest to write were the acknowledgements. This was because of the sheer number of people I needed to thank, not just for helping me write this piece of work, but for supporting me throughout my PhD over the past four years. Without this help and support, none of this would have ever been possible. I can only hope that the chapters to come are something we can all be proud of.

The first point of call is my supervisors. If you were to ask anyone on my MA course to describe me, they would have undoubtedly mentioned the *Routledge Handbook of Epistemic Injustice* glued to my hands. It was upon reading the chapters contained within this volume that I was first introduced to virtue epistemology, and eventually vice epistemology. So, when I decided to apply for PhD programmes, it was a no-brainer that I went to Nottingham to work with one of its editors, Ian James Kidd. As somebody with no formal philosophy background, Ian's supervision on this project not only taught me a great deal about philosophy; he helped teach me *how* to do philosophy. Whether it was in Christine's Delights, Bendigo, or Greenhood, I will always be thankful for the time he took to read over drafts of these chapters, as well as the rigorous feedback that made each one significantly better than before. I cannot thank him enough for helping me, both academically and personally, over the past four years. That my thesis builds on his work is a testament to his influence on me.

No less an influence is Michael Hannon. Michael joined my supervisory team at the outset of the pandemic, and ever since, has been an inspiration to me. Although our supervision meetings were more conventional than Ian's, they were just as important in shaping my philosophical work. As somebody who doesn't work primarily in virtue epistemology, he was a blessing for my thesis: whenever I sent him drafts of chapters, he was able to offer a fresh and critical pair of eyes to my many assumptions; whenever I found myself stumped by a problem, he was able to help me see the issue differently and overcome it; and whenever I needed help or just general advice about something, he could provide it. This was made all the better by our shared sense of humour. Thank you for all your help with this thesis.

In September 2020, Quassim Cassam joined my supervisory team. Being able to work under somebody who, arguably, put vice epistemology on the philosophical map has been both a privilege and an honour. From our Zoom meetings, to sitting in his office in Warwick, I would like to thank Quassim for reading over drafts of chapters 4 and 5, for providing incisive comments and suggestions to my arguments, and for taking the time to discuss vice

epistemology together. I will always remember his advice: “one cannot do philosophy without sugar”, before proceeding to take a pack of custard creams from his desk drawer during one of our meetings. I can assure you that I consumed much sugar in the process of writing this thesis.

It was only through receiving an Arts and Humanities Research Council Midlands4Cities (M4C) Doctorial Scholarship that I was able to get Quassim on-board. I am very grateful to M4C for funding my project over the past three years. Without this funding, I would also not have been able to travel to the Universities of Glasgow, Cardiff, Sheffield, Birkbeck, and Humboldt Universität Berlin, where on each occasion my work benefitted from the rich comments, suggestions, and feedback I received from participants.

Behind the scenes, there were a number of people at Nottingham, past and present, who supported and encouraged me throughout my PhD. You know who you are. For the avoidance of doubt, I would like to thank Janset Cetinkaya, Giorgio Mazzullo, Alice Monypenny, Andrea Raimondi, and the PGR community at Nottingham Philosophy. As an Aristotle scholar, Janset knew the importance of friendship from day one, helping me settle into the department as well as providing invaluable comments and feedback on my work. This was especially true when it came to digesting the *Nicomachean Ethics*. Giorgio, our shared sense of humour meant that coming into the department was never dull. Not only could I rely on your sharp wit to make me laugh, but your philosophical insights were even sharper, helping me think through a number of issues across this thesis. Alice, as a fellow virtue epistemologist, it was a pleasure to work with you. Being able to exchange ideas on a range of subjects was invaluable to the development of this thesis, particularly the ideas found in chapters 1 and 4. Andrea, having sat next to you for the past several years, I am still unsure which of us types the loudest. Either way, your willingness to read my work and offer penetrating suggestions was highly valuable, thank you.

Doing a PhD in philosophy can be a pretty isolating experience, and mine was made much better with the support and love of my friends, partner, and family. You always checked in on me when times were hard, and when I needed an escape from philosophy you were the best people to relax and unwind with. In particular, I would like to thank Fabrice, Iona, Reece, and Beth for your support over the past decade and for always being there for me through thick and thin. Erinn, it goes without saying that your love and encouragement over the past four years has enriched my life. Whether I came to you in despair about my PhD or subjected you to you mini lectures on vice epistemology, you were always able to pick me up, keep me focused, or tell me when to take a break. Thank you.

Finally, to my mum and dad: you have always been there for me, supported me in everything I have done, and been a constant source of love in my life. Whenever I doubted myself, you were always there to remind me of the positives. While you might not immediately recognise your influence on my work, the truth is that I wouldn't have been able to do philosophy, let alone go to university, without your help, advice, and unfailing support. For that alone, thank you.

## Abstract

This thesis examines the relationship between character and intellectual or *epistemic* vices. The philosophical study of epistemic vices is called vice epistemology. To date, much of the work in this emerging field has focused on the nature and epistemological significance of particular intellectual vices such as close-mindedness or dogmatism. Far less has been said about how it is that people come to acquire and develop these intellectual vices. My aim is to fill this lacuna by articulating how this phenomenon occurs. Specifically, this thesis develops an account of *epistemic corruption*.

I start in chapter 1 with a critical discussion of character. Despite character forming a foundational pillar of responsibilist virtue epistemology, there is surprisingly little discussion as to what intellectual character might look like. The discussion in this chapter attempts to get clearer on this matter by offering an *ontology of character* that is both conceptually and empirically robust going forward. In chapter 2, I introduce the phenomenon of *epistemic corruption* in detail, distinguishing a variety of ways it might occur. Chapter 3 examines how these variants operate in practice by considering the finer *mechanics of epistemic corruption*. There, I raise difficulties for an existing account and offer solutions to offset these difficulties. In chapter 4, I turn to consider whether *collective agents* might be susceptible to epistemic corruption. Drawing on Miranda Fricker's account of institutional character, which she terms 'institutional ethos', I distinguish two ways in which this might occur. Ultimately, though, I urge caution about Fricker's account, and with it, the prospects of applying epistemic corruption to collective agents. Finally, in chapter 5, I raise concerns about a different analogy in virtue theory: the *skill-analogy*. I argue that this analogy has problematically influenced how vice epistemologists treat skill in their theorising of epistemic vice, resulting in an assumption that the latter will be characterised by an absence or deficit of the former. I refer to this as the *deficiency thesis* and reject it. In doing so, I note how its rejection provides a novel insight into the nature of epistemic vices going forward.

## Table of Contents

Introduction.....	8
1. Virtue Epistemology: A Brief History.....	10
2. Vice Epistemology and Epistemic Corruption.....	13
3. Thesis Overview.....	17
Chapter 1: An Ontology of Character.....	20
1. Persons and Characters.....	22
2. Situationism and Empirically-Informed Character.....	29
3. The Ontological Problem.....	37
4. Conclusion.....	50
Chapter 2: Varieties of Epistemic Corruption.....	52
1. The Aretaic Conception.....	53
2. The Phronetic Conception.....	56
3. The Attitudinal Conception.....	61
4. Conclusion.....	62
Chapter 3: The Mechanics of Corruption.....	65
1. Corrupting Conditions.....	66
2. The Axes of Corruption.....	69
3. Passive Epistemic Corruption.....	76
4. Active Epistemic Corruption.....	82
5. Conclusion.....	87
Chapter 4: Epistemic Corruption and Collective Agents.....	89
1. Institutional Ethos.....	90
2. Corrupted Epistemic Ethos.....	98
3. The Limits of an Analogy.....	107
4. Conclusion.....	119
Chapter 5: Competence in Vice Epistemology.....	121
1. The Deficiency Thesis.....	123
2. The Skilfully Vicious I: Weak Deficiency.....	133
3. The Skilfully Vicious II: Strong Deficiency.....	142
4. Conclusion.....	153
Concluding Remarks.....	155
Bibliography.....	159



# Introduction

At 6:20am on the 29<sup>th</sup> October 2018, Lion Air Flight 610 took off on a domestic journey from Jakarta to Pangkal Pinang, Indonesia. The flight was scheduled to land an hour later, but it never did. Just thirteen minutes after take-off, it crashed into the Java Sea. All one-hundred and eighty-nine passengers were sadly killed. At 8:38am on the 10<sup>th</sup> March 2019, Ethiopian Airlines Flight 302 set off on an international flight from Addis Ababa, Ethiopia towards Nairobi, Kenya. Six minutes after take-off, the plane plunged towards the ground and crashed, killing all one-hundred and fifty-seven passengers on board.

Within the span of five months, how could two different planes have crashed so devastatingly? On closer inspection, there was a common denominator between Lion Air Flight 610 and Ethiopian Airlines Flight 302: both planes were manufactured by Boeing. Until very recently, Boeing was the most profitable aerospace company in the world, responsible for designing and manufacturing the famous 737 jets we regularly board to go on holiday. In 2017, Boeing introduced its newest fleet of jets, the 737 Max. The engines on the 737 Max were built forward onto the wings, which introduced a risk that too much thrust could cause the plane's nose to lift, increasing what those in aviation call the *angle of attack* (AOA) – the angle between the wings and the flow of air. Above an optimum angle, a plane's nose will sharply decrease due to turbulence and subsequently stall. To prevent stalling, Boeing relies on its Manoeuvring Characteristics Augmentation System (MCAS).

The problem is that Boeing was aware of a fatal flaw built into this system. In testing, they found the MCAS received flawed information from an external AOA sensor on the plane, which caused the 737 Max's nose to repeatedly dive. Unfortunately, there was nothing pilots could do about this. Boeing, to ensure they remained in competition with Airbus, concealed this information from their own pilots and the Federal Aviation Administration (FAA). Both planes involved in the crashes above were 737 Max jets.<sup>1</sup> The 737 Max crashes were the culmination of a series of failings within Boeing. Prominent amongst these were its intellectual or *epistemic* failings. In the report that followed, the House Committee on Transportation and Infrastructure found Boeing to have created a 'culture of concealment'

---

<sup>1</sup> For an investigation on the 737 Max crashes, see the United States House of Representatives' Committee on Transportation and Infrastructure findings (2020), available at: [https://www.govinfo.gov/content/pkg/GOVPUBY4\\_T68\\_2fb0f3812fefe3515ebcf3f4170fce64b/pdf/GOVPUB-Y4\\_T68\\_2-fb0f3812fefe3515ebcf3f4170fce64b.pdf](https://www.govinfo.gov/content/pkg/GOVPUBY4_T68_2fb0f3812fefe3515ebcf3f4170fce64b/pdf/GOVPUB-Y4_T68_2-fb0f3812fefe3515ebcf3f4170fce64b.pdf).

(2020: 3), a part of which involved withholding crucial technical information from employees, customers, pilots, and the FAA. However, Boeing did not just *withhold* this information, it actively *suppressed* it. When engineers raised concerns about the MCAS, Boeing either removed them from projects or engaged in ‘efforts to obfuscate information about the operation of the aircraft’ (*ibid.*: 12). Those who retained their jobs told Congressional hearings how they were expected to deliver technical designs at ‘double the normal pace’ and how ‘rushed engineers’ created ‘sloppy blueprints’ (Gelles et. al., 2019, paras. 8, 29). Instead of fostering carefulness, creativity, and innovativeness, Boeing created a working environment that encouraged deceit, intellectual carelessness, and narrow-mindedness in its employees and executives.

Deceit, intellectual carelessness, and narrow-mindedness are all examples of qualities that tend to make us *bad thinkers*. Being intellectually creative, by contrast, is something that makes us a *good thinker*. The former set of qualities are what epistemologists have come to call intellectual or *epistemic* vices, whereas the latter quality is an intellectual virtue.<sup>2</sup> Within contemporary epistemology, it is the job of virtue and vice epistemologists to study the nature and significance of intellectual virtues and vices. While there has been increasing work on this front, especially on the virtue-epistemological side of things, far less attention has been paid to *how* it is people come to lose their intellectual virtues or develop intellectual vices.<sup>3</sup> In this thesis, I am concerned precisely with the ways in which our social environments can erode our intellectual virtues or encourage us to develop intellectual vices. By encouraging deceit, intellectual carelessness, and narrow-mindedness, Boeing created an *epistemically corrupting* environment. (Kidd, 2018a, 2019, 2020). My aim is to develop a full account of this phenomenon.

In this Introduction, I offer a brief trajectory of contemporary virtue epistemology, before homing in on the development of vice epistemology and situate my project within this emerging discipline. I end by providing an overview of the chapters to come.

---

<sup>2</sup> In this thesis, I tend to use the terms ‘intellectual’ and ‘epistemic’ largely interchangeably, though I somewhat prefer ‘intellectual’ over ‘epistemic’ because of the latter’s narrower connotations with truth as opposed to knowledge, understanding, and wisdom more broadly. Not much hangs on this preference, however.

<sup>3</sup> For the most part, theorists have been interested in the intellectual character *development* of agents, usually in the form of intellectual character education. See, for example, Battaly (2016), Baehr (2013, 2016), and Kotzee (2014).

# 1. Virtue Epistemology: A Brief History

A concern for peoples' intellectual conduct has a long and venerable history in philosophy. In the *Analects*, Confucius stresses the importance of personal and inter-personal reflection not just as a means of acquiring knowledge, but as a way of developing one's character (2003, 1.4, 5.20).<sup>4</sup> Several of Plato's dialogues, including the *Republic*, *Symposium*, and *Phaedrus*, offer guidance on how we can improve our intellectual affairs. Aristotle, in addition to his exploration of the character virtues, explicitly identifies several distinct intellectual virtues (*NE*, bk. VI). In his *Summa Theologiae*, moreover, Aquinas builds on Aristotle's intellectual virtues (2014, 2a.2ae).<sup>5</sup> The early-modern period was also characterised by similar concerns, with John Locke (particularly his *Some Thoughts Concerning Education* [1693] (1996), Francis Bacon (*Novum Organon*, [1620] (1859), David Hume (*An Enquiry Concerning Human Understanding*, 5.1, [1748] (1975), and René Descartes (*Rules for the Direction of Mind*), all taking an interest in what made for good thinking.<sup>6</sup> More recently, these concerns were taken up by classical pragmatist such as Dewey (*How we Think*) and Pierce ('The Fixation of Belief').<sup>7</sup> As Alasdair MacIntyre (2013) might have put it, a concern with the virtues that make us good thinkers is one that goes all the way back through our history.

Until very recently, however, contemporary epistemology had failed to uphold these concerns. Rather than focus on the intellectual conduct of agents, epistemologists have focused instead on their belief-states, paying particular attention to what makes one's beliefs justified, warranted, or non-accidentally true. Part of this was the result of epistemology's incessant attempts to fend off external-world scepticism and provide solutions to an ever-growing list of Gettier problems. Fortunately, this state of affairs began to shift in the early 1980s with the publication of two influential articles: Ernest Sosa's (1980) 'The Raft and the Pyramid' and Lorraine Code's (1984) 'Toward a "Responsibilist" Epistemology'. Taking up Elizabeth Anscombe's (1958) earlier call for a return to virtue theory, both argued that epistemology would do well to reorient its focus on the epistemic agent. Just as Anscombe called for rule- and act-based ethical theories to be abandoned, so Sosa and Code pushed back against the perceived failure of belief-based epistemology. In its place, they made the case for

---

<sup>4</sup> See Mi (2015, 2017) for an illuminating comparison of Confucius' insights on knowledge and Ernest Sosa's virtue epistemology. See, also, Kidd (2018b) and Cheng-Hung (2014).

<sup>5</sup> See DeYoung (2009) for discussion of virtues and vices in the early Christian Tradition.

<sup>6</sup> Corneanu (2011) offers an insightful discussion of intellectual character within the early-modern period; Vitz (2009) draws out the virtue-epistemological roots of Hume's work; and Kidd (2022) provides a rich discussion of the need to develop an historically-sensitive vice epistemology.

<sup>7</sup> Alfano (2013b) offers a compelling case for thinking of Nietzsche as a virtue epistemologists.

agent-based epistemologies, at the heart of which was the concept of an intellectual virtue. What followed was the beginning of contemporary virtue epistemology, or, perhaps more accurately, virtue *epistemologies*.

I say this because there is no one single “virtue epistemology”. From the outset, Sosa and Code were interested in two different projects. These projects are nicely captured by Nicholas Wolterstorff’s (1996) distinction between ‘analytic’ and ‘regulative’ epistemology. The former produces theories of knowledge, justification, and rationality while the latter attempts to generate guidance for our epistemic practices and conduct (Roberts and Woods, 2007: 21-22). On the one hand, Sosa sought to apply the concept of an intellectual virtue to analytic epistemology. The motivation for this was the impasse reached by foundationalists and coherentists about epistemic justification. Rather than approach this problem through the lens of process reliabilism – i.e., the view that a belief is justified if it is produced by a reliable belief-forming process (Goldman, 1979) – Sosa argued that we should instead understand these processes as intellectual *virtues*, conceived as reliable truth-conducive cognitive competences of agents (1991, 2007, 2015). The inspiration for this was an insight dating back to Plato, according to which ‘anything with a function – natural or artificial – does have virtues’ (Sosa, 1991: 271). Just as the virtue of a knife is its sharpness, so the virtue of our cognitive faculties is their reliability. Hence, the kind of intellectual virtues of interest to Sosa included things like good eyesight, memory, hearing, etc. Nowadays, prominent virtue *reliabilists* such as John Greco (2010), Duncan Pritchard (2012, 2017, 2020) and others not only employ the concept of an intellectual virtue (or cognitive competence/ability) to analyse epistemic justification, but to also explore the nature, normativity, and value of knowledge (Carter, 2016; Kelp and Greco, 2020; Turri, 2015).

Code, on the other hand, set out to apply the concept of an intellectual virtue to regulative epistemology. Her aim was to shift attention away from debates between foundationalists and coherentists and instead provide a robust account of epistemic responsibility. Whereas Sosa took inspiration from Plato’s conception of an intellectual virtue, Code relied heavily on Aristotle’s moral virtues to develop her account. Accordingly, she understood epistemic responsibility largely as a matter of one’s orientation towards the world and one’s ‘knowledge-seeking self’, with the virtues of intellectual *character* at the forefront of this (1984: 41). Whilst these virtues included intellectual courage, rigour, and curiosity, she argued that epistemic responsibility was the ‘central virtue from which others radiated’ (*ibid.*: 34). In the ensuing forty years, virtue *responsibilists* have continued to conceive of intellectual

virtues as praiseworthy features of intellectual character.<sup>8</sup>

Of course, this is a rather simplified trajectory of the virtue-epistemological story. In reality, things are more complicated. For one, early responsibilists did try to apply the concept of an intellectual character virtue to traditional epistemological problems. James Montmarquet (1993), for example, employed the concept of an intellectual character virtue to develop an account of doxastic justification (see, also, Fairweather, 2001). Linda Zagzebski went even further by providing a virtue-responsibilist account of knowledge, according to which knowledge is the result of true belief arising out of ‘acts of intellectual virtue’ (1996, 2003).<sup>9</sup> Such responsibilist ambitions have recently fallen on hard times, however, with many theorists now conceding that the concept of an intellectual character virtue is poorly equipped to address traditional epistemological questions, at least on its own.<sup>10</sup> Nevertheless, the responsibilist character virtues have still proved useful in supporting evidentialist accounts of epistemic justification (Baehr, 2011). Moreover, many virtue reliabilists acknowledge an important, auxiliary role for intellectual character virtues in their accounts of knowledge (Greco, 2010; Sosa, 2015, 2021).

It would be a mistake, though, to think that virtue reliabilism and responsibilism are in some sort of competition (Fleischer, 2017). As I noted above, both branches are motivated by different projects. For virtue reliabilists, it is analysing the nature and value of knowledge. For virtue responsibilists, it is analysing, evaluating, and guiding our intellectual conduct. While the two projects clearly overlap, my interests in this thesis are largely within the confines of responsibilist virtue epistemology.<sup>11</sup> More specifically, I am interested in something that Code picked up on long ago, but which has received only fleeting attention in the years since (e.g. Battaly, 2013; Kidd, 2019). This is that factors in a person’s environment and community ‘have crucial bearing upon the forms intellectual virtue can take’ (1984: 46). If one’s environment can determine one’s ability to cultivate intellectual virtues, that is, so

---

<sup>8</sup> Prominent responsibilists include James Montmarquet (1987, 1992), Linda Zagzebski (1996), Roberts and Woods (2007), Heather Battaly (2008), and Jason Baehr (2011), to name just a few.

<sup>9</sup> To be more accurate, knowledge on Zagzebski’s view is simply *belief* that arises out of acts of intellectual virtue because the beliefs in questions are already true as a result of being produced by acts of intellectual virtue (1996: 271).

<sup>10</sup> These ambitions fall under what Baehr (2011) calls ‘strong conservative responsibilism’. Baehr is far more optimistic about the prospect of ‘weak conservative responsibilism’, which see character virtues as playing an auxiliary role in solving traditional epistemological questions. Though, see Sylvan (2020) for a revised, ‘strong conservative responsibilist’ account of knowledge. Kvanvig (1992) and Roberts and Woods (2007) endorse ‘strong autonomous responsibilism’, which is the view that independent study of the intellectual virtues should replace traditional epistemology.

<sup>11</sup> As we will see in chapters 2 and 5, however, there is still an important place for reliabilist virtues and vices in an account of epistemic corruption.

too can it undermine their development. Worse still, those very environments can degrade our epistemic conduct by encouraging us to develop intellectual *vices*.

## 2. Vice Epistemology and Epistemic Corruption

In her *Virtues of the Mind*, Linda Zagzebski (1996: 171) remarks that ‘mistakes’ in belief-formation have historically been of more interest to philosophers than the remedies to those pitfalls. Such mistakes can take a myriad of forms. They might include poor vision, unreliable hearing, bad reasoning skills, and the like. However, a brief journey through intellectual history reveals an entirely different story.

In his *Some Thoughts Concerning Education*, Locke expressed how certain ‘errors in education’ were linked to what he saw as the ‘corruption of the youth’ [1663] (1996: 8). Later in his *Of the Conduct of Understanding*, this corruption is correlated with a range of bad intellectual habits, including despondency in inquiry and resigning oneself to the last beliefs or opinions encountered [1706] (1996: §§39, 27).<sup>12</sup> Hume, too, was well aware of the mental qualities that could hinder an agent’s belief formation, noting that ‘rash arrogance’ and ‘superstitious credulity’ could ‘detract from the character of the person, possessed of them’ (EHU, 5.1; SBN: 41, 312; EPM, 4.1). Similar concerns were expressed by early feminist philosophers. Mary Astell’s 1694 *A Serious Proposal to the Ladies* criticised societal deficiencies in upper class English women’s educational opportunities, lamenting the development of what she termed ‘feminine vices’ such as submissiveness and superficiality. These vices, she argued, caused women’s characters to become ‘degenerated and corrupted’ (2002: 62). A century later, Mary Wollstonecraft condemned the ‘negative virtues’ women were expected to develop in society, which, she argued, were ‘incompatible with any vigorous exertion of the intellect’ (1995: 133) [1792]. In her view, ‘negative virtues’ like docility and agreeableness were a way in which the male world ‘conspires to render the cultivation of understanding more difficult’ for women (*ibid.*: 129). Such claims are not confined to the early modern period. Even among contemporary epistemologists, such concerns recur. Miranda Fricker (2007: 49, 58) has claimed that subjection to epistemic injustice can ‘inhibit’ or ‘thwart’ the development of intellectual virtues, while José Medina (2013: 42) contends that those under conditions of oppression are exposed to ‘practices and processes that erode their epistemic character’.

---

<sup>12</sup> For a virtue- and vice-theoretic reading of Locke, see Yolton (1998).

For the most part, then, the ‘mistakes’ in belief-formation of interest to philosophers have resided in the character of epistemic agents. We might, therefore, expect contemporary virtue epistemology to have continued this trend. Alas, this has not been the case. Despite their preoccupation with good intellectual conduct, virtue responsibilists have spent little time considering the epistemic vices, the character traits, attitudes, and sensibilities that make us bad knowers.<sup>13</sup> Ironically, it was early virtue reliabilists who first discussed these qualities. Again, though, the vices in question were not unreliable cognitive impairments, such as deteriorating memory or poor eyesight. Sosa (1991: 229) cites haste and inattentiveness, Greco (1993: 415) identifies wishful thinking and superstition, and Goldman (1992: 162) picks out guesswork, wishful thinking, and ignoring contrary evidence. These are not just failings of intellectual character, but qualities for which people are often criticisable or blameworthy. Only within the past decade, however, has there been sustained interest in the intellectual vices, ushering in what Quassim Cassam (2016) named *vice epistemology* – the philosophical study of the nature and epistemological significance of intellectual vices.

There are two main normative camps in the growing vice-epistemological literature. Following in responsibilist footsteps, one camp takes intellectual vices to involve certain motivational states. So, just as responsibilists insist that intellectual virtues involve good epistemic motivations, or what Zagzebski (1996: 167) terms ‘cognitive contact with reality’, this camp understands vices to involve defective epistemic motivations. Unlike responsibilists, though, there is division among the ranks. Some theorists unpack “defective” in terms of ‘dis-valuable’ or ‘bad epistemic motives’ (Battaly, 2016: 106, 2017: 5), or ‘non-instrumental motives to oppose, antagonise, or avoid things that are epistemically good in themselves’ (Tanesini, 2018: 353, 2021). Others, meanwhile, understand “defective” to mean a ‘lack of motivation for knowledge’ (Zagzebski, 1996: 207), a ‘characteristic failure to attend to truth or truth-related considerations’ (Montmarquet, 2000: 138–139), or an ‘insufficient concern with knowledge’ (Baehr, 2010: 219, 2020). Following in reliabilist footsteps, the other camp takes intellectual vices to be qualities that reliably produce bad epistemic effects. This camp has a precedent in Julia Driver’s virtue-consequentialism (2000: 126, 2001, 2003), which states that ‘a character trait is an intellectual virtue iff it systematically (reliably) produces true belief. Accordingly, a reliabilist reading of Driver’s view would be something like: “a character trait is an intellectual vice iff it systematically (reliably) produces false beliefs. In several places, Cassam (2016, 2019a) has built on Driver’s work by introducing

---

<sup>13</sup> A number of earlier discussions are found in Fricker (2007), Roberts and Woods (2007), and Swank (2000).

what he calls ‘obstructivism’. Obstructivism is the view that a character trait, attitude, or way of thinking is an intellectual vice insofar as it systematically obstructs the gaining, keeping or sharing of knowledge (2019: 12). Unlike their virtue-reliabilist cousins, however, vice-reliabilists deal exclusively with features of intellectual character.<sup>14</sup> There has, moreover, been no explicit attempts to deploy the competence or achievement-theoretic framework of virtue reliabilism to the intellectual vices. As we will see in chapter 5, this is an unwarranted omission by vice epistemologists.<sup>15</sup>

It goes without saying that our intellectual vices do not manifest in the abstract. Just as Code emphasises the role of the community in cultivating intellectual virtues, so too can our communities encourage us to develop intellectual vices, thereby degrading our characters. Despite being of principal interest to the likes of Locke, Astell, and Wollstonecraft, such concerns have been few and far between within vice epistemology. Roberts and Woods (2007: 251) suggestively remark that an epistemic community can be ‘warped’ by vanity and arrogance, hyper-autonomy, and competitiveness. In this ‘fallen community’, they argue, some vices can ‘become more functional than their counterpart virtues’. The upshot is that there has been ‘some corruption in the epistemic environment’ (*ibid.*: 251). Furthermore, José Medina (2013: 72) claims that hermeneutically privileged interlocutors – those better equipped with more and better conceptual resources – can be ‘worse off because their epistemic characters tend to become more corrupted’.

Inspired by feminist epistemologists and ethicists, Ian James Kidd (2018a, 2019, 2020) introduced the concept of *epistemic corruption* into the vice-epistemological vernacular, as a means of capturing the deterioration of intellectual character.<sup>16</sup> Epistemic corruption thus appeals to the ordinary sense in which corrupting something involves making it worse off. Interestingly, the term ‘epistemic corruption’ seems to originate in Stephen Gardiner’s book *A Perfect Moral Storm: The Ethical Tragedy of Climate Change* (2011). However, the sense in which Gardiner uses this term differs from the sense in which Kidd uses it. For Gardiner, epistemic corruption has a narrow sense and a specific context: it names a tendency to ‘invoke...scepticism selectively against climate science’ despite not invoking the same degree of scepticism towards one’s other beliefs (*ibid.*: 462). It is this ‘epistemic double-standard’

---

<sup>14</sup> For the purposes of this thesis, I do not take a particular stance on these theoretical camps.

<sup>15</sup> In several places, Heather Battaly (2016, 2017, 2018b) has articulated a view she calls ‘personalism’, which is a hybrid between virtue responsibilism and reliabilism. However, her view does not seem to be widely embraced in the literature.

<sup>16</sup> Kidd’s contemporary inspiration came from the work of Claudia Card (1996), Lisa Tessman (2005), and Robin Dillon (2012), amongst others.

which Gardiner believes underpins epistemic corruption.<sup>17</sup>

Following Kidd, I will not be approaching epistemic corruption through this lens.<sup>18</sup> Instead, I understand the phenomenon in terms of those social environments which encourage us to develop intellectual vices or facilitate the loss of our intellectual virtues. In short, I am interested in how our social environments can degrade or damage our intellectual characters. To help better situate my project, we can, following Kidd, Battaly and Cassam (2020), distinguish three broad projects within vice epistemology:

1. *Foundational.*
2. *Applied.*
3. *Ameliorative.*

By developing the first fully-fledged account of epistemic corruption, my project here is a contribution to foundational vice epistemology. It seeks to study the nature and significance of intellectual vices not in the abstract, but by being sensitive to their social roots or *aetiology* (Kidd, 2016a). If the same intellectual vice can manifest across our social environments in markedly different ways, then my account of epistemic corruption is not only able to tell us something important about this vice, but it is able to reveal unique insights into the nature of intellectual vices more broadly. Needless to say, an account of epistemic corruption is going to have something to say about particular intellectual vices and failings. In this way, my project contributes to applied vice epistemology by offering analyses of specific intellectual vices, as well as identifying a range of other failings with which vice epistemologists ought to be concerned. It is all well and good identifying the ways in which our social environments can degrade our intellectual character, but it is equally, if not more, important to devise ways of preventing this from happening. To the extent that we are interested in what makes for bad intellectual conduct, we should be just as interested in avoiding it. Developing an account of epistemic corruption is the first step in achieving this. Ultimately, then, my project here is essentially ameliorative in its aims.

---

<sup>17</sup> See Gelfert (2013) for a rich discussion of Gardiner's notion of epistemic corruption.

<sup>18</sup> This is not to say that Gardiner's concept is entirely irrelevant. Elsewhere, I have deployed the concept of epistemic corruption to scepticism and knowledge. See Matthews (2022, 2023).

### 3. Thesis Overview

There are many moving parts to an account of epistemic corruption. In what follows, I will provide a brief overview of the chapters of this thesis.

I start in chapter 1 by providing an ontology of intellectual character. This is important because intellectual character is the site on which epistemic corruption occurs, eroding virtues or encouraging vices. Despite intellectual character forming the cornerstone of responsibilist virtue epistemology, contemporary theorists have said little about its nature. Instead, the concern has traditionally been with an agent's epistemic virtues (and vices). For the most part, these qualities have been modelled along Aristotelian lines, but as I argue in chapter 1, such a view faces difficulties on two fronts: on the one hand, evidence from social psychology suggests that our situational environments dictate our behaviour more than robust, stable character traits (Alfano and Fairweather, 2017; Doris 2001; Harman, 1999); on the other hand, the view that character traits are grounded in mental-state dispositions faces a dilemma I refer to as the *Ontological Problem* (Jeffrey and Beary, 2023). In this opening chapter, then, my task is to provide an account of character capable of overcoming both challenges. This is possible, I argue, by grounding character traits in *attitudes* (Tanesini, 2021).

With a clearer understanding of character to hand, I turn my attention to the focal point of this thesis: developing an account of epistemic corruption. This starts in chapter 2, where I distinguish three distinct, yet overlapping, forms of the phenomena. The first of these, which I term the *Aretaic* conception, is inspired largely by Kidd's work on the concept. After supplementing his conception, I identify two further forms, not present in his account, which I call the *Phronetic* and *Attitudinal* conceptions of epistemic corruption. I claim that the former turns on conditions that erode an agent's ability to exercise practical wisdom in deploying their intellectual virtues, as well as conditions that undermine the development of an agent's intellectual skills. Meanwhile, the latter focuses on the propensity for our social environments to encourage us to develop bad attitudes towards epistemic practices like inquiry and deliberation, which then feed the development of bad intellectual character traits. In light of these varieties, chapter 3 offers a critical analysis of what I call the "mechanics" of epistemic corruption, that is, the ways in which the phenomenon dynamically operates. In particular, I unpack a range of criteria we are likely to find under each of the varieties explored in the previous chapter. Once again, I take my lead from Kidd's work. While important in their own right, I argue that several of his proposals face problems in their existing guise. Though

largely diagnostic, then, this chapter provides a number of critical amendments to central features of any corruptionist critique, namely, the distinction between passive and active epistemic corruption, what Kidd calls the ‘axes’ of corruption, as well as the more general notion of epistemically corrupting conditions. In doing so, I not only provide a more robust framework for thinking about epistemic corruption, but one that is more informative from an ameliorative perspective.

Whereas chapters 2 and 3 examine the ways in which our social environments can degrade our intellectual characters, chapter 4 takes up the opposing question: how is it that our social environments can be hotbeds for epistemic corruption? Here, I ask if we can extend the concept of epistemic corruption to collective and institutional agents. An immediate question is whether such agents might even be said to possess epistemic virtues and vices. I demonstrate this by appealing to non-summativism about collective agency, the view that collectives can possess features that are irreducible to the features of their individual members. Drawing on Miranda Fricker’s (2013, 2020) account of ‘institutional ethos’, I distinguish two ways in which individuals might epistemically corrupt the collectives to which they belong: by passively corrupting its ethos and by actively corrupting it. While the notion of a *corrupted epistemic ethos* is a useful way of understanding the epistemic deterioration of collectives, I ultimately urge caution against drawing too close an analogy between collective and individual character.

Finally, in chapter 5, I turn my attention to a different analogy in virtue theory: the skill-analogy. This is the idea that virtue is (like) a practical skill. Not only has the skill-analogy influenced how virtue epistemologists conceive of epistemic virtues, I argue, but it has also shaped how vice epistemologists treat skill in their theorising of epistemic vice. Specifically, it has resulted in an assumption – that epistemic vices are characterised by an absence or deficit of skills – that I refer to as the *deficiency thesis*. I claim there are two ways of interpreting this assumption: (i) vices involve a deficit of the intellectual skills that characterise epistemic virtues; and (ii) vices involve a deficit of skill more generally. I argue that both interpretations are mistaken. Far from being characterised by a deficit of skill, I demonstrate how the very same theoretical resources which underpin the latter interpretation of the deficiency thesis actually allow us to model reliabilist epistemic vices on competences. I end by suggesting that these vices will often be auxiliary to responsibilist or character-based vices, which I illustrate through their role in epistemic corruption. I conclude this thesis by reflecting on the account of epistemic corruption arrived at, as well as three areas of future research that my account generates.



# 1.

# An Ontology of Character

## Introduction

This thesis aims to articulate an account of epistemic corruption. This is the broad claim that our social environments can facilitate the loss of our intellectual virtues or encourage us to develop intellectual vices. In the Introduction, I noted that the task of understanding how agents come to develop their intellectual vices has been a comparatively neglected topic in recent vice epistemology. Nevertheless, what discussion there is has tended to identify these virtues and vices as features of a person's *intellectual character*. In later chapters, we shall see how the remit of epistemic corruption can be broadened beyond character. Given how central this notion is to extant discussions, though, it will be worth getting clearer on what epistemologists mean when they talk of "intellectual character". This is the task of this opening chapter.

In some ways, this task is quite straightforward. Since virtue epistemology has been largely shaped by Aristotelian virtue ethics, the intellectual virtues have been widely treated as the epistemic analogue of the moral virtues.<sup>19</sup> This, in turn, has shaped how responsibilists have approached questions of intellectual character. When developing her early responsibilist approach to epistemology, for example, Lorraine Code (1984: 34) emphasises that the model of character central to her project was the "intellectual analogue of the stable virtues and

---

<sup>19</sup> While this is usually assumed, Alkis Kotsonis (2021) demonstrates that numerous Platonic dialogues provide rich material for thinking about intellectual virtues, including the *Republic*, *Symposium*, and *Phaedrus*.

dispositions'. James Montmarquet (1987: 488) notes that the epistemic virtues 'possess a kind of entrenchment Aristotle describes the moral virtues as having'. Linda Zagzebski (1996: 139) famously claimed in her *Virtues of the Mind* that the 'intellectual virtues ought to be treated as a subset of moral virtues in the Aristotelian sense'. If we are inclined to think that the former manifest "intellectual character", then it seems that this notion is simply reducible to "moral character". Jason Baehr (2011: 6, fn. 9) states from the outset of his *The Inquiring Mind* that the character model of intellectual virtue with which he is concerned 'resembles Aristotle's account of moral virtue'.<sup>20</sup>

Of course, if we think that there are indeed relevant similarities between the intellectual and moral virtues, then it would seem there is little need for us to embark on an independent study of intellectual character. By endorsing an Aristotelian picture of moral character, that is, we can arrive at an equally informative understanding of intellectual character. In short, we get two for the price of one. However, such a strategy would simply entertain what Christine McKinnon (1989: 326) calls a 'theory-by-default', which is to assume that an account of the latter neatly falls out of the former. This is not only something that many virtue epistemologists seem to have done, but also one that carries with it problematic implications for virtue theorists more generally. In a nutshell, the issue is that those who uncritically appeal to a broadly Aristotelian sketch of character face what I shall call the *Ontological Problem*.

My aim in this first chapter is to present an ontology of character that can avoid this problem and serve as an attractive alternative going forward in the rest of the thesis. I start in section 1, where I distinguish the notion of "character" from "personality". In section 2, I address recent empirical challenges to the notion of "character" advanced in situationist social psychology. This lays the foundations for section 3, in which I address the *Ontological Problem*. I then conclude in section 4.

---

<sup>20</sup> Tanesini (2021) is a recent outlier, offering an insightful discussion of intellectual character virtues and vices grounded in attitude psychology. For psychologically-informed discussions of character in virtue ethics, see, Fleeson and Jayawickreme (2015); Masala and Webber (2016); Miller (2013, 2014), and to some extent Wright et. al. (2020). In addition to Aristotle, the 'unity of virtue' thesis found in Plato's *Protagoras* (1996) also represents an attempt at conceptualising moral character.

# 1. Persons and Characters

What does it mean to talk of “character”? At first, the answer to this question might appear simple: it is what a person is generally and consistently *like*. However, it is common to hear people describe themselves and others as “agreeable”, “introverted” and “extroverted”, “conscientious”, “open”, or simply as “large personalities”. At the same time, those very same people might be described as charitable, timid, arrogant, humble, or curious. Is it not the case that all these qualities pick out what a person is like? If we are to adequately address my opening question, then, it seems that we also need to make sense of this second question.

In this first section, my aim is to get clearer on the notion of character by distinguishing it from its cognate in “personality”.<sup>21</sup> By doing so, we can capture the sort of qualities relevant to the later task of providing an ontology of character. Accordingly, in section 1.1, I unpack the notion of *personality*. I then distinguish personality from character by examining two issues: the role of responsibility for character (section 1.2) and the thickness of “characterological” concepts (1.3).

## 1.1 Personality and the Big Five

Let us start with the case of Robin. Robin is somebody who attends several sports clubs, regularly finds himself at parties and gatherings, and frequently engages in various other social activities. From this brief description, we would likely describe Robin as an extravert. Now, let’s suppose that when Robin attends these social activities, he regularly looks down on his teammates, thinks of himself as better than the rest, and defensively brushes aside criticism when challenged on his behaviour. It now appears that Robin is not just an extravert but that he is also arrogant. Yet, the qualities of “being an extravert” and “arrogant” seem to pick out very different behaviours. How can we vindicate this?

One way of attempting this is by saying that the quality of “being an extravert” is part of Robin’s *personality*. Being an extravert, or simply extraversion, is said to form one of the personality traits that make up the widely accepted ‘Big-Five’ model in personality

---

<sup>21</sup> In fact, some virtue theorists simply take character traits and personality traits to be interchangeable. See, for example, Battaly (2018a), Miller (2013, 2014), and Statman (1997). Given the differences I highlight below, I take this to be a misguided use of terminology.

psychology (Allport and Odbert, 1936; Goldberg, 1993; John and Srivastava, 1999).<sup>22</sup> In addition to extraversion, the so-called ‘five-factor’ model includes the traits of introversion, agreeableness, conscientiousness, openness, and neuroticism. As personality traits, these qualities are broadly held to involve causal dispositions to act, think and feel in response to socio-cognitive stimuli. Consider Robin again. Suppose he is invited to a party. In response, we would expect an extravert like Robin to be disposed to form certain beliefs regarding the party, such as who might be in attendance and who he could potentially meet. Furthermore, the beliefs and other mental states that Robin forms will tend to give rise to an emotional response: upon being invited, he is disposed to feel excitement at the prospect of meeting new people and socialising. Together, the mental and affective components of Robin’s trait of extraversion dispose him to act in a particular way: to accept the invitation and attend the party.

A similar story goes for the other traits that compose the Big Five. For instance, the highly introverted person is usually disposed to be shy or quiet around people; somebody low in agreeableness is often competitive and less cooperative with their peers; openness disposes one to be adventurous and creative; a highly conscientious person is disposed to be organised, mindful of details, and punctual (Picone et. al., 2021), whereas someone low in this trait is prone to procrastinate or be less organised; and those high in neuroticism are disposed to be moody or pessimistic. Of course, this is just a rough sketch, and the reality is that our personality traits come in varying degrees and overlap considerably. Nevertheless, even this brief picture reveals two important features of personality traits worth discussing.

First, the personality traits we possess are generally stable across time. Part of the reason why we might describe somebody like Robin as a “conversationalist”, for example, is because when exposed to extraversion-triggering stimuli – say, a party, wedding, or other social occasion – he is disposed to spend more time talking with others than to be alone (Mehl et.al., 2006, Carment et. al., 1965). Alternatively, we would typically expect somebody who registers high in neuroticism to be a “natural worrier” in the sense that exposure to stimuli like sitting an exam or attending an interview disposes them to feel anxious about these tasks. Now, this is not to say that personality traits are wholly fixed or that an extraverted person might not sometimes refrain from engaging in social activities; rather, the point is that our personality traits will tend to reveal a level of consistency to our actual behaviour when confronted with the relevant stimuli.

---

<sup>22</sup> While the ‘Big Five’ or five-factor model of personality traits is influential, it does have its detractors. For relevant discussion, see De Raad and Perugini (2002) and McCrae (2009). See, also, McAdams and Pals’ (2006) ‘New Big-Five’ model of personality.

This brings me to the second point. As the discussion above suggests, personality traits are often helpful indicators of behaviour. The extent of this influence can vary from the explicit to the implicit. At the more explicit level, there is growing empirical work in psychology that points to the different behaviour patterns of financial investors. According to one study, those who exhibited behaviour consistent with the trait of conscientiousness were far more likely to avoid incorporating recent market information into the stocks they traded than those who were less conscientious (Singh et. al., 2022). At the more implicit level, however, personality traits have been correlated with numerous cognitive biases. For example, studies have found that those high in extraversion and neuroticism are likely to exhibit overconfidence and herding biases (Durand et. al., 2013; Baker et. al., 2021). Thus, the effects of personality traits can vary quite markedly.<sup>23</sup>

What seems clear, regardless, is that personality traits involve dispositions to think, feel, and act in ways characteristic of the trait in question. But if this is the case, then we are no closer to figuring out how character traits differ from personality traits, especially since the former are widely held to exhibit the same features as the latter. Assuming that the two qualities do share these features, it is perhaps understandable why people conflate personality and character. Across the next two sections, I will offer two considerations as to how we might go about distinguishing the two.

## 1.2 Character and Responsibility

We can approach the first reason by reflecting further on Robin's extraversion. Much of the empirical research in personality psychology over the last thirty years has pointed to a strong connection between this trait and optimism (Marshall et. al, 1992; Sharpe et. al., 2011; Williams, 1992).<sup>24</sup> Considering this, let's now add that Robin is generally disposed to be optimistic in his outlook on life; he feels as though his future is bright and that even setbacks have their positives. Is Robin's optimism something we would be willing to hold him accountable for? Sure, the fact that he is optimistic could dispose him to engage in wishful

---

<sup>23</sup> For studies that questions the predictive potential of personality traits, see Kouchaki et. al. (2014), Mischel (1968) and Vernon (1963). Similarly, see Darley and Batson (1973), Doris (2002), and Harman (1999) for critical discussions on the role of situational factors on personality.

<sup>24</sup>This evidence also suggests that neuroticism, as well as openness, agreeableness, and conscientiousness, is an important factor in somebody being optimistic. See, for example, Costa jr. and McCrae, (1992); Lounsbury, Saudargas, and Gibson (2004); and Segerstrom, Castañeda, and Spencer (2003).

thinking about things he's interested in, or it could encourage him to be arrogant in his social life. But in itself, this provides no reason to think that Robin has any control over his optimism; at most, it shows that the psychological structure and mechanisms of his personality can, to paraphrase Nancy Snow (2013: 131), 'scaffold' the development of his traits. Another way to put the point is to say that Robin *just is* an optimistic person.

This conclusion, however, seems far less palatable when it comes to Robin's arrogance. Even if his arrogance extends from his extraversion, I think we are intuitively less inclined to accept that some people "just are" arrogant. When we admit that things "just are" a certain way, we usually imply that there was little or no control over the state of affairs turning out this way. A similar point extends to Robin. In saying that he "just is" an arrogant person, we imply that he had little control over becoming this way. Whether or not this actually is the case, what we are implicitly invoking here is a narrow sense of *responsibility* – or more aptly, a lack thereof.<sup>25</sup> We are suggesting that Robin is not responsible for his trait of arrogance. If he lacks responsibility for his arrogance, then his arrogance is on a par with his extraversion, for he presumably lacks responsibility for this personality trait. This seems to be the wrong verdict, which points to our first preliminary difference between character and personality: we are generally *responsible* for the development of our character traits in a way that is lacking when it comes to our personality traits.

The idea we are responsible for our character development has a venerable tradition in virtue theory. In the *Nicomachean Ethics* (bk. II, 1101a33), Aristotle famously identified traits of character to be the locus of moral virtue and vice. An important reason why he chose character traits, as opposed to natural faculties or temperaments, was because these qualities emerge through our own volition *via* a process of self-conscious habituation. As he put it, states of character arise out of similar activities, so that by, say, performing acts in the face of danger, and being habituated to feel fear or confidence, we *become* brave or cowardly (bk. II, 1103a33).<sup>26</sup> States of character, then, are qualities over which we have some degree of control to shape. David Hume (1960: 411, [1739]) similarly claimed that actions which 'proceed not from some cause in the character and disposition of the person who performed them' can 'neither redound to his honour, if good; nor infamy, if evil'. Hume's point is that our characters often causally determine our actions and so make us answerable – and hence responsible – for them. More recently, for Bernard Williams (1981: 130), 'if one acknowledges responsibility

---

<sup>25</sup> Of course, this is the case insofar as one is not sceptical about moral responsibility in the sense of accountability or control. For the purposes of this thesis, I set these worries aside. See, though, Pereboom (2014, 2021) for seminal discussions.

<sup>26</sup> All references to the *Nicomachean Ethics* (*NE*) are to the translation by Ackrill and Urmson (1998).

for anything, one must acknowledge responsibility for decisions and action which are expressions of character'. The general assumption is that we have a good deal of control over the development of our character traits, and *ex hypothesi*, our virtues and vices, which in turn makes us responsible for the actions that flow from them.

Now, it is worth emphasising that we do not automatically lack responsibility for our personality traits even if we lack full control over them. There is certainly a sense in which people can and do take steps to alter their personalities – say, an introverted person who tries to attend more social events in an effort to be more extraverted, or a pessimistic person who tries to begin seeing the positives in life. Cases like this are surely familiar; however, they are also very unlikely to be quick or overly explicit. Often, they will occur at a sub-personal level in ways that are not always immediately accessible (nor recognisable) to the agent. At best, any shifts in personality like this will lead us to evaluate agents in what Angela Smith (2008: 373) calls a normatively 'superficial way'; we will evaluate them in a way that warrants neither reproach nor fault. The same, however, cannot be said about Robin's arrogance, and the reason why this is takes us to the second difference between character and personality.

### 1.3 The “Thickness” of Character

An important reason why personality traits only merit the kind of 'superficial' evaluation just noted is because there is nothing pre-theoretically problematic about them. We might find Robin's extroversion impressive, but – again – this trait does not reveal much about him *qua* person; it is just the tip of the personal iceberg, as it were. Hence, personality traits like extraversion are normatively 'superficial' features of people at best. By contrast, there is something intuitively problematic about Robin's arrogance in a way that reflects something deeper about who he is as an agent. We can capture the intended contrast by saying that Robin's arrogance is conceptually “thicker” than his extravertedness (Foot, 1958; Hare, 1952; Williams, 1985). This speaks to the second difference between personality and character.

To unpack this contrast, it will help to first distinguish two sorts of concepts. Take something like the concept of “a field” or “playing cricket”. These are *descriptive* concepts, in that they serve to impart or describe some informational content about the concepts in question. Now take something like the concepts of “brilliant” or “terrible”. These represent *evaluative* concepts because they introduce value-judgements into our considerations. In saying that a piece of work is “brilliant”, that is, we are evaluating it in a positive light.

Bernard Williams (1985) introduces the notion of a 'thick' concept in order to

distinguish it from a ‘thin’ concept.<sup>27</sup> A thick concept is something like “cruel” or “selfish” whereas a thin concept is something like “bad” or “impermissible”. In this way, both the former and the latter issue evaluative judgements. Thick concepts, by contrast, are also descriptive in a way that thin concepts are not. Williams (*ibid.*: 140) makes sense of this difference by claiming that thin concepts are ‘action-guiding’ in the sense that an action being “bad” or “impermissible” only leads us to refrain from undertaking it. However, they do not tell us *why* an action is “bad” or “impermissible”. Here, a thick concept like “cruel” enters to provide a ‘world-guiding’ evaluation in the sense of providing others with reasons as to why something is bad or impermissible to start with (*ibid.*: 140-141).

Insofar as thick concepts are ‘world-guiding’, they generally offer a deeper sense of evaluation than thin concepts. As part of this, they evoke an emotional component that thin concepts alone seem to lack.<sup>28</sup> Williams offers a nice example of this affective force: “Of course, he went back on his agreement when he got to the meeting, the little *coward*” (Callcut, 2009: 212, my italics) [1973]. If we were to remove the thick concept “coward” from this sentence, Williams notes, it would take the form of: “As might have been predicted, he went back on his agreement at the meeting through fear; which he *ought not to have done* (or this was a bad thing)” (*ibid.*: 213). While both sentences convey moral judgements, it is clear that the second lacks the emotional force of the first. In Williams’ terms, it lacks the moral ‘overtones’ that reveal how the agent views the situation in light of their own outlook. Only a thick concept like “coward”, it seems, can do justice to the deeper sense of assessment that is lacking in the second sentence.

With this in mind, consider the following two sentences:

(1) “Of course Robin thinks he’s right; he’s so *optimistic!*”

(2) “Of course Robin thinks he’s right; he’s so *arrogant!*”

Clearly, both sentences convey some form of evaluation of Robin’s intellectual conduct. Yet, it is equally clear that (2) is far more descriptive than (1), and this is because “arrogant” is a thicker concept than optimistic. It not only provides a reason to think that Robin’s thinking is epistemically bad or defective, but it reveals a layer of value-commitments that are likely to

---

<sup>27</sup> Although Williams was the first to make this distinction in ethics, Väyrynen (2021) points out that he likely adopted it from Hare’s (1952) notions of ‘primarily’ and ‘secondarily’ evaluative words. It also has a precursor in the work of social anthropologist, Clifford Geertz (1973). Recently, the notion of ‘dual-character concepts’ has emerged, which are concepts whose descriptive and normative content are related but independent. For a philosophical overview, see Reuter (2019).

<sup>28</sup> Alistair MacIntyre (2013) also praises virtue/character concepts for their “thickness” by noting how they are richly rooted in our social practices and traditions.

tell us *why* his thinking is such.<sup>29</sup> This is because the value-commitments that ‘thicken’ his arrogance, as it were, inform the kind of reasons he takes for his actions.<sup>30</sup> These plausibly include reasons relating to his self-conception, to his intellectual peers, and ultimately to how he sees the world around him: Robin thinks he can win not because of the “positive feelings” that come with being optimistic but because he views himself as epistemically superior to those around him.<sup>31</sup> It is for this reason we might judge there to be something intellectually bad or defective about Robin’s epistemic conduct. In short, (2) imports stronger epistemic ‘overtones’ into our assessment of Robin *qua* cognitive agent than (1).

The fact that character traits are conceptually thicker than personality traits allows us to appreciate how their descriptive content – the values, desires, motivations, and sensibilities they capture – stands in close enough relation to us to reflect well or poorly on us as people (Sher, 2005: 57; Tanesini, 2021). By contrast, personality traits garner, at best, what I earlier referred to as normatively ‘superficial’ evaluation; they deserve neither reproach nor fault. To paraphrase Linda Zagzebski (1996: 85), character traits are more closely associated with one’s identity *as a person* than their personality traits.<sup>32</sup> This more general feature allows us to not only appreciate why we hold people responsible for the development of their character traits, but also why we are far more willing to invoke a range of reactive attitudes towards their manifestation than we do with personality traits.<sup>33</sup> Whereas we usually are not willing to blame or hold somebody in contempt for their pessimism or extraversion, we *are* often more than happy to attach the same reactions to their arrogance, cynicism, or close-mindedness.

Nevertheless, just because we are willing to attribute blame and praise to the character traits we develop, it does not follow that we should always invoke these attitudes. Given that many of us occupy various ‘epistemic predicaments’ – characterised by different obstacles,

---

<sup>29</sup> I should acknowledge that the contrast is not always so sharp, particularly in relation to a personality trait like conscientiousness. Unlike the other personality traits mentioned, conscientiousness seems to be a hard case that falls somewhere in the middle of the thin-thick spectrum. That said, I still think we can draw out the intended contrast by noting that conscientiousness itself picks out a range of more specific behaviours, including diligence, carefulness, and attentiveness, which do strike me as “thicker” concepts. In this way, these thicker traits once again fall out of a thinner trait. I thank Michael Hannon for raising this difficulty.

<sup>30</sup> This is not to discount the thickness of “optimistic” or other behaviours that stem from personality traits. Saying that somebody’s pattern of thought is “optimistic” is far more informative than describing it as “bad” or “right”. But this only speaks to the gradable nature of thick concepts themselves (Scheffler, 1987: 418).

<sup>31</sup> How we perceive the world is an important feature of our characteristic ‘sensibilities’ (McDowell; 1979; Fricker, 2007; Tanesini, 2020).

<sup>32</sup> Aquino and Reed (2002: 1424) offer a similar assessment with regards to what they call ‘moral identity’.

<sup>33</sup> For discussion of ‘reactive attitudes’, see Strawson’s (1963) *Freedom and Resentment*.

deficits, privileges, and luck – we will naturally be afforded varying resources on which we can draw to overtly control and determine the way our character turns out (Medina (2013: 28, 34; Dillon, 2012). For this reason, it will be worth attending to the very *aetiology* – the origins and causes – of our character traits (Kidd, 2020), and this means being sensitive to the kinds of reactive attitudes we wish to invoke.

## 2. Situationism and Empirically-Informed Character

So far, I hope to have persuaded you that we can distinguish the notion of *character* from that of *personality*. If we accept such a distinction, we find ourselves committed to what is often referred to as “realism” about character (Miller, 2018). This is the idea that character is a real entity that exists in the world, one which people not only possess but act from in the form of globally robust traits of behaviour. In Aristotelian terms, such traits are said to ‘proceed from a firm and unchangeable character’ (*NE*, bk. II). We can capture this mantra in the following two commitments, to which any theory of character traits is supposed to be faithful:

- 1) **Stability over Time:** To the degree that a person has a character trait *T*, she will exhibit characteristic behaviour of *T* in multiple cases of the same situation over time.
- 2) **Cross-situational Consistency:** If a person has a character trait to a high degree, she will manifest that trait in behaviour across different kinds of situations (Miller, 2013: 92).

In recent years, a number of empirical findings have brought into question the veracity of these two commitments. The bulk of this research has come from the Situationist Movement within social psychology, which has sought to cast doubt on the idea that our character traits are robustly stable, and therefore on the very utility of appealing to character traits in explaining our moral and intellectual conduct. Before I am in a position to address the *Ontological Problem* in section 3, then, I need to tackle this concern. In section 2.1, I unpack the ‘situationist critique’ as it has applied to virtue theory. In section 2.2, I consider the options with which the empirical evidence leaves us, before providing my own preferred empirically-informed response in section 2.3.

## 2.1 The Situationist Critique

The Situationist Movement took its starting point from a range of empirical results in social psychology conducted in the middle-latter half of the 20<sup>th</sup> century. Often, the studies cited are the infamous “Obedience to Authority” experiment (Milgram, 1974), the “Dime in the Phone Booth” case (Isen and Levin, 1972), and the “Lady in Distress” experiment (Latané and Rodin, 1969: 193–195; Latané and Darley, 1970: 60–63). In these experiments, researchers reported that they could induce certain behaviours in participants by altering specific situational features of the environment. The results indicated that situational factors led to increased cruelty from participants (Milgram, 1974), that participants in a good mood were more willing to help others (Isen and Levin, 1972), and that participants’ willingness to aid a screaming woman was dependent on how strangers behaved (Latané and Darley 1970: 60–63).

Obviously, these results were quite striking, both empirically and theoretically. On the empirical side, they not only suggested that a person’s behaviour was more often a function of their situational environments, but that situational influences played a *more salient* role in governing a person’s behaviour than something like robust “global” traits of character. After all, it seems reasonable to think that those with the trait of benevolence would reliably offer help in a trait-relevant situation like that of a woman screaming, contrary to what the Dime and Lady experiments found. That seemingly irrelevant situational factors had a significant effect on people’s moral conduct suggested that people often do not possess cross-situationally robust traits like benevolence. More worryingly, if situational factors do indeed play a greater explanatory role in our behaviour than global character traits, then this would seriously undermine the theoretical force of any philosophical view that places character at its centre, namely virtue theory.

Enter John Doris (1998, 2002, 2010, 2021) and Gilbert Harman (1999, 2000, 2009). Against the backdrop of these empirical results, both have long argued that we are justified in thinking that people do not possess what we think of as the traditional virtues or vices.<sup>34</sup> More schematically, we can reconstruct their argument in the following way (Miller, 2014: 192):

---

<sup>34</sup> Although Doris and Harman have been two of the most prominent situationists, it is worth noting that Flanagan (1991) raised earlier concerns about virtue theory that appealed to empirical psychology. For further discussion, see Merritt et. al., (2010).

**Premise 1:** If there is widespread possession of the traditional virtues and vices understood as global character traits, then systematic empirical observation using appropriate psychological experiments will reveal that most people behave in a certain kind of way.

**Premise 2:** Systematic empirical observation using appropriate psychological experiments fails to reveal that most people act in this kind of way.

**Conclusion:** Thus, there is not widespread possession of the traditional virtues and vices, understood as global character traits.

Now, some clarifications about this argument. First, it would be easy to think that the pair outright rejected the *existence* of character traits – indeed, they each make suggestive remarks to this effect: that, ‘people typically lack character’ (Doris, 2002: 2) and that ‘it may even be the case that there is no such thing as character’ (Harman, 2000: 165). If there is no such thing as character, then surely there are no such things as character traits. However, we should be careful here because theirs is a weaker claim against the existence of *global* character; that is, a property from which robust, cross-situationally stable character traits flow. In fact, both are open to the possibility of character traits being restricted to “local” situations (Doris, 2002: 23, 25, 64). The problem, then, boils down to what Ross and Nisbett (2011) term the ‘fundamental attribution error’. This refers to peoples’ ‘inflated belief in the importance of traits and dispositions, together with their failure to recognise the importance of situational factors affecting behaviour’ (*ibid*: 4).

Second, the argument is clearly aimed at virtue *ethics* as opposed to epistemology. Unfortunately, this does not put virtue epistemology in the clear since the general thrust of the situationist critique is quite easily extended to the intellectual virtues and vices, and hence to intellectual traits more generally. Indeed, this has been attempted by Mark Alfano (2012, 2013a; Alfano and Fairweather, 2017).<sup>35</sup> He cites studies such as the Duncker (1945) candle test, in which participants were tasked with fixing a candle to a vertical cork board in a way that prevented wax from dripping when the candle was lit. The experiment was a measure of one’s intellectual flexibility and creativity, but results showed that very few people were able to correctly construct the candle, matchbox and thumbtacks in a way that prevented wax dripping. Potentially more troubling were the results from Isen et. al.’s (1987) study, which found that trivial mood elevators such as viewing a comedy or giving participants sweets led

---

<sup>35</sup> See, also, Olin and Doris (2011) for discussion.

to significantly more success in the Duncker candle experiment.<sup>36</sup> Alfano marshals these results to support *epistemic situationism* about intellectual character traits, claiming that the studies highlight how people’s intellectual conduct is seemingly susceptible to trivial and epistemically irrelevant situational influences (2012: 232, 2013a). If we take epistemic situationism at face value, it suggests that situational factors play a more important role in our intellectual conduct than something like intellectual character. In turn, this casts doubts on the epistemological utility of such a notion.

## 2.2 Character Pessimism?

Given the evidence, it seems incumbent on any virtue epistemologist (and virtue theorist more broadly) to respond to epistemic situationism. The question now, though, is how should we respond? Perhaps the most straightforward response is to embrace the empirical results on offer. To emphasise, this need not commit us to thinking that character traits do not exist, since that would mean any articulation of intellectual character fails to get off the ground. Rather, it would mean abandoning the idea that our character traits are robustly global indicators of behaviour. It would still be perfectly consistent with the empirical evidence employed by situationists to view character traits as “local” properties of agents. As I noted above, this position is open to Harman and explicitly endorsed by Doris.

If we embrace the empirical data, the resulting ontology of character is one in which an agent’s traits are “fragmented” rather than parsimoniously integrated into a whole. For instance, it would be consistent to acknowledge that somebody possesses the trait of open-mindedness in their political dealings with others, but that their open-mindedness fails to extend to non-political matters. It does not follow from this, of course, that this person is thereby *closed-minded* in their non-political dealings – though this conclusion is interestingly compatible with viewing virtues and vices as local traits. More importantly, given the kind of predicaments I outlined at the end of section 1, it is often very difficult to cultivate and sustain character virtues (Cooper, 2018: 67). Indeed, Philippa Foot (2002: 9-10) [1978] follows Aquinas in conceiving of virtues as ‘correctives’ against the natural temptations towards badness in life. That virtues are “correctives” of this sort has been a theme of much work in recent feminist epistemology. Miranda Fricker (2007: 93) takes testimonial justice to be a

---

<sup>36</sup> It is also worth noting that when a group of people are in the same mood, one will tend to find that such elevators and inhibitors likely induce important differences in the behaviour of those *within* the group. This could be taken as further evidence against global virtue.

corrective, anti-prejudicial virtue. Nancy Daukas (2006: 112-116) identifies ‘epistemic trustworthiness’ to be a socially scaffolded virtue of agents, highly contingent on practices of credibility distribution. Heidi Grasswick (2017) claims that many epistemic virtues are correctives insofar as they are instantiated in response to situationally-dependent contexts of oppression. In light of this, an ontology of character premised on local traits would appear consistent with, and sensitive to, the non-ideal realities of life.

What’s more, notice that many of the scenarios which situationists are interested in are epistemically corrupting in the sense relevant to my thesis – they are situations in which a person fails to manifest a recognisably virtuous trait of character or where people come to manifest other, more recognisably vicious traits. In this way, a corruptionist critique is largely adjacent to, and compatible with, the situationist critique: what the latter finds troubling about the scope of global virtue and vice is precisely what the former recognises as the need for an ameliorative resource to rectify and improve our epistemic conduct.

Of course, though, if we accept this, then the next question is whether localised traits really are *traits*. If we are thinking of such traits as indexed to domain-specific areas of our lives, one might question the sense in which an agent is really *disposed* to act, think, and feel in the relevant sense? For example, compare two vases which are disposed to break when dropped on the relevant stimuli. Barring the potential for masking, suppose they exhibit this feature in most counter-factual situations. But now suppose one vase is disposed to break only when dropped on a particular surface, say, concrete – if it were dropped on any other surface, it would remain intact. Is the vase still disposed to break? In a narrow sense, perhaps. However, a further question now arises over how fine-grained we wish to “localise” the disposition. When it comes to our character traits, one might say, there is only so much fine-graining of the disposition we can make before it becomes *ad hoc* at best and non-existent at worst.

One way to push back against this would be to say that we can aggregate a person’s local traits, since this would still allow us to reliably stitch together a picture of how a person characteristically acts. As Doris (1998: 508) puts it, such a move could still allow us to ‘know quite a bit about how we could expect that individual to behave’. If so, it seems that a picture of character understood in terms of local traits is not immediately problematic. However, a potentially more troubling conclusion for this sort of proposal is the fact that much of the empirical data on which situationists have premised their critique of virtue theory has either failed to replicate or replicated poorly (Alfano, 2018; Klein et al. 2014; Klein et al. 2017; Ebersole et al. 2016). This problem appears to be symptomatic of a broader “replication crisis” currently engulfing much of psychology. Naturally, this has led several virtue theorists to

insist that a picture of character with something like robustly global traits is still very much on the table. Indeed, these theorists often acknowledge this replication crisis and then move quickly on to reject situationism (Tanesini, 2021; Wright et.al., 2020).

I would urge caution about such quick responses. For starters, only time will tell whether or not this replication crisis can be averted. If it cannot be (and I find this conclusion quite unlikely), virtue theorists would still do well to articulate empirically-informed accounts of character since this is likely to best reflect the realities of life as we find it. On the other hand, if this crisis can be averted then virtue theorists need to be prepared. No longer will they be in a viable position to deftly dismiss the evidence in an attempt to ‘move beyond’ this challenge (Wright et. al., 2020: 1, fn.1). After all, nothing would be stopping social psychologists and philosophers from conducting further experiments about the relationship between character and situational stimuli. Indeed, this just opens the door to new empirical data that could easily dampen the virtue theorist’s hopes.<sup>37</sup> Either way, then, it is beneficial to incorporate empirical evidence into one’s picture of character, including intellectual character.

## 2.3 Empirically-Informed Character

I want to end this section by examining an influential attempt at providing an empirically-informed picture of character proposed by Christian Miller (2013, 2014, 2017). In doing so, my aims are twofold. I first want to demonstrate that character realists can adequately respond to the worries articulated by situationists, regardless of whether the empirical results to which they appeal do in fact stack up. Second, I take initial inspiration from Miller’s account in developing my own ontology of character in section 3 below. As such, introducing his view allows me to lay the groundwork for my own account.

According to Miller (2013, 2014), our character traits consist of interrelated mental-state dispositions to form beliefs and desires relevant to the trait in question. He claims that whether these traits qualify as virtues or vices largely depends on whether they manifest the specific dispositions to form beliefs and desires relevant to the virtue or vice (2013: 38). In light of the empirical data from situationism, however, Miller takes seriously the possibility that people often lack the appropriate clusters of mental-state dispositions needed for their

---

<sup>37</sup> Some of this has plausibly already happened, such as the results from Baehr’s Intellectual Virtues Academy (2021).

traits to form traditional character virtues and vices.<sup>38</sup> Nevertheless, Miller attempts to reconcile the empirical findings with the two commitments I outlined at the beginning of this section with what he calls a ‘Mixed-Trait’ ontology of character. On his view, mixed-traits consist of mental-state dispositions that pertain to various moral domains.<sup>39</sup> For instance, what he calls the mixed-aggressiveness trait relates to how we interact with and treat other moral agents across ‘harming-relevant situations’, while the ‘mixed-cheating’ trait corresponds to situations involving honesty, lying, and integrity (2013, 2014: 47, 73). As with traits *simpliciter*, though, Miller contends that the possession of mixed-traits comes in degrees, such that two people might possess the same mixed-trait but exhibit different patterns of behaviour within the same trait-relevant situation. For instance, two people with the mixed-cheating trait might manifest the relevant mental-state dispositions. However, only one might exhibit behaviour consistent with cheating when the risk of being caught is high while the other refrains.

Now, if exhibiting behaviour relevant to a given mixed-trait can vary from one person to the other, depending on the situation one finds oneself in, one might argue that Miller’s account leads us to the very same picture of character presented above. That is, one which takes character to be highly fragmented and localised. Wouldn’t this show that a mixed-trait ontology of character is not faithful to the commitments above? This is too quick. While Miller acknowledges that mixed-traits do offer an initially fragmentary picture of character, he demonstrates that these traits nevertheless result in stable behaviour once we *aggregate* the average levels of trait-relevant behaviour across a given time period. Let us return to the mixed-cheating trait. Suppose we have two people who possess this trait. Even if the trait differs with respect to how it manifests within these two people, there will remain a sense in which their respective mixed-cheating traits will still exhibit stability within trait-relevant situations. If one person is disposed to cheat on exams even when the chances of being caught are high, while the other is disposed only when the stakes are low, then we can aggregate their behaviour across the relevant situations, and this will reveal that they exhibit a good deal of stability even if their respective mixed-traits differ in strength. Indeed, the fact that mixed-traits – and traits in general – come in degrees is partially why this stability arises. Thus, by aggregating a person’s mixed-trait behaviour, Miller’s ontology of character is faithful to **Stability over Time**.

---

<sup>38</sup> Miller is clear, however, that this conclusion still allows for some people to cultivate traditional virtues and vices, just that they will not be the norm (2013: 46).

<sup>39</sup> It is worth noting that the term ‘mixed’ is intended to capture the moral evaluation of these traits, in that the dispositions constituting them can take on morally positive and negative aspects.

What about **Cross-situational Consistency**? For starters, Miller is cognisant to the role of situational enhancers and inhibitors on the manifestation of such traits (2013: 250-253; 2014: 50-51). These can range from emotional states, such as guilt or remorse, social norms like standing in a crowd, to even affective states like being in a good or bad mood.<sup>40</sup> Despite their heterogeneity, enhancers and inhibitors are said to play an important role in influencing the *motivations* we draw upon when acting in trait-relevant ways. For example, studies suggest that feelings of empathy are a salient motivation in refraining from aggressive behaviour.<sup>41</sup> Importantly, some enhancers or inhibitors will be trait-specific, in which case they will be a function of the mental-state disposition that ground that trait – think of angry emotions enhancing aggression. But if mixed-traits are susceptible to these influences, in what sense can they be said to exhibit cross-situational consistency? Other things being equal, that is, shouldn't we expect someone with the trait of aggression to consistently act aggressively across situations, regardless of whether they happen to be in a good mood, angry, and so on?

Miller (2014: 54) admits that mixed-traits will not exhibit cross-situational consistency if we take the 'nominal' perspective, that is, the perspective of a third-party experimenter or onlooker. But if we consider what the psychologically salient features of a given situation might be relevant from the perspective of the subject, we will find cross-situational consistency. From this perspective, a particular perfume might remind the subject of a past partner and trigger a mental-state disposition to act in a patterned way that is not immediately recognisable to onlookers (Millar and Richardson, 2018). While another person might not respond to the perfume, its presence will be a psychologically salient feature to the specific subject. This will then cause the person to respond to and interpret like-situations in a consistent manner. More generally, Miller claims that this 'higher-order consistency' extends to mixed-traits more generally and is a function of their mental-states dispositions responding to stimuli and interpreting them accordingly. As a result, mixed-traits do lead to cross-situationally consistent behaviour.

Although Miller says that his account of mixed traits is focused only on moral character (2015: 167), I see no reason why we cannot try to extend his ontology to intellectual character. This would involve claiming that epistemic agents possess clusters of what we might call "mixed-intellectual traits". At this point, we can stipulate that these traits would be grounded in various mental-state dispositions to respond and interpret epistemically-

---

<sup>40</sup> For relevant discussion, see Geen (2001), Krahé (2001), Miller (2013, 2014), and Wilkowski and Robinson (2010) amongst many others.

<sup>41</sup> See, for example, Baumeister et al. (1996), Eisenberg (2000), Stuewig et al. (2010), Baron-Cohen (2011), and (Batson 2011).

charged situations. Under different conditions, these dispositions can strengthen, weaken, stabilise, and destabilise across the epistemic domains we encounter. For instance, “mixed-open-mindedness” might involve dispositions to form certain beliefs and desires (say, acquiring information, to change perspectives, and so on) but these dispositions will emerge across localised situations at first. If the conditions are congenial, the dispositions will strengthen and stabilise across a range of wider situations, eventually culminating in what we recognise as *virtuous* open-mindedness. This will be stable and cross-situationally consistent in the way just discussed. But if the conditions are unfriendly, then we can expect these dispositions to be inhibited and eventually wane. Worse still, they might actively enhance and stabilise “mixed-close-mindedness” in becoming an intellectual *vice*.

In chapter 2, I shall show more fully how this ontology of character complements the notion of epistemic corruption. For now, note that if the stability and strength of our mixed-traits is susceptible to situational enhancers or inhibitors, then this correlates with there being conditions that encourage the development of certain mixed-intellectual traits while facilitating the loss of others. In other words, it helps to vindicate the idea of there being epistemically corrupting conditions. Second, if we take for granted the idea that most people possess a range of mixed-intellectual traits as part of their intellectual characters, as opposed to traditional virtues and vices, then this makes the case for epistemic corruption all the more important. This is because corrupting conditions result in agents developing and manifesting recognisably vicious behaviour. But if vices are indeed uncommon, then this calls for a serious ameliorative resource in epistemic corruption to identify and correct for such conditions.

### 3. The Ontological Problem

Across the previous two sections, my aim has been to clarify the notion of character and defend appeals to character traits from recent challenges in social psychology. In this final section, I will put more meat on this skeleton of character by addressing what I have called the *Ontological Problem*. The task of addressing this problem is important for two reasons. First, metaphysical discussions of character have been ‘almost entirely neglected’ in virtue theory (Miller, 2014: 30). Indeed, the discussion in the previous section illuminated the potential problems that come from uncritically relying on a broadly Aristotelian framework. In this way, many recent attempts at articulating an ontology of character have been a direct response to the worries posed by situationism. Second, many of these empirically-informed

ontologies share a common commitment to the idea that mental-state dispositions *ground* character traits. As we shall see shortly, however, there is a potential dilemma awaiting these views, including Miller's. My aim is to provide an ontology of character that avoids this dilemma.

In section 3.1, I unpack a taxonomy of metaphysical views about character, which allows me to introduce the purported dilemma at the heart of the *Ontological Problem* in section 3.2 (Jeffrey and Beary, 2020). Not only does Miller's mixed-trait ontology of character fall squarely within this camp, but so does my own eventual view. Therefore, my own account is potentially vulnerable to the same dilemma. In section 3.3, I appeal to attitude psychology to side-step this dilemma by arguing that attitudes are the causal basis of the mental-state dispositions that form our mixed-traits. In doing so, I cast doubt on the force of this dilemma.

### 3.1 Taxonomising Character

A helpful way of introducing the dilemma at the heart of the *Ontological Problem* is by appealing to a recent attempt at taxonomising various metaphysical views about character. This taxonomy responds to two particular questions. The first is this:

**Grounding Question:** What is the order of grounding between dispositions to believe  $p_1 \dots p_n$  and desire  $d_1 \dots d_n$  and character trait  $T$ , where both are involved in the production of some behaviour typical of persons with  $T$  (Jeffrey and Beary, 2023: 188).

When it comes to answering this, perhaps the most intuitive response is 'reductivism' (2023: *ibid*). Character reductivism posits that a person's fine-grained mental-state dispositions *constitute* their character traits; that a person has these dispositions, in other words, explains or grounds their behaviour. This is precisely how Miller sees the relationship between mixed-traits and their dispositions. As he puts it, 'what *immediately* underlie a trait disposition are *further dispositions*'; that these dispositions to form certain trait-relevant desires and beliefs are the 'causal base of the disposition' and that Jones' compassion 'is itself *grounded* in underlying mental-state dispositions' (2013: 10-11, my italics, 2014: 26). Given that my piecemeal account of intellectual character largely employs Miller's mixed-trait ontology, it is worth stating up front that this commits me to the kind of character reductivism here.

A second response to the Grounding Question is 'non-reductivism'. On this family of views, the fact that a person possesses a trait  $T$  is what explains their fine-grained dispositions to form particular desires and beliefs. Thus, non-reductivism holds that Robin's trait of

arrogance is what explains or grounds his mental-state dispositions, not the other way around. This is not to say, of course, that these fine-grained dispositions play no causal role in his behaviour, since Jeffrey and Beary (2023: 188) are willing to accept that these ‘proximately cause behaviour’. Instead, the idea is that character traits *ultimately* cause behaviour. For this reason, non-reductivism posits that traits are not mental qualities that supervene on a cluster of more fine-grained dispositions; they are metaphysically basic.

I want to briefly flag an ambiguity with this view. According to non-reductivism, a character trait is what ultimately causes behaviour. As just noted, this does not mean that mental-state dispositions play no role, just that this role is limited to proximately causing behaviour. However, it is unclear how we should understand this relationship between proximate and ultimate causes of behaviour. One way could be as follows: as ultimate causes of behaviour, character traits are the *most salient* explanations of one’s conduct. As Jeffrey and Beary put it, ‘if Krista has the trait of aggression, she is disposed to form beliefs about needing to fight and to form desires to fight when someone teases her *in virtue of* having aggression’ (2023: 188). In other words, the main cause of Krista’s aggressiveness is not primarily because of her dispositions to form certain beliefs and desires, rather it is because of her *trait* of aggression.

But is this a coherent position to adopt? Notice how Jeffrey and Beary appeal to the term ‘trait of aggression’ in their example of Krista. In doing so, however, they seem to be begging the question regarding the order of metaphysical grounding. This is because they are presupposing that there is a “trait” of aggression that ultimately causes Krista’s behaviour to start with. Of course, if we presuppose that traits do exist, it might make sense to think that fine-grained dispositions flow from these coarse-grained qualities. Without presupposing this, though, the task of explaining this set of causal events looks problematic. To see why, consider this question: why create terms like “aggression” or indeed “trait” in the first place? If these qualities are indeed metaphysically basic, as Jeffrey and Beary claim, then, surely, we would not need to refer to fine-grained mental-state dispositions. A term like “aggression” would already be encoded with the relevant behaviours. This, however, seems mistaken. Suppose an alien was spying on a group of humans. Now, imagine that the alien notices that one of the humans always pushes their peers around, shouts at them, and is generally violent. If the alien has no concept of “aggression”, how would they describe the behaviour observed? Precisely in the ways just noted. Yet this would suggest that the mental-state dispositions are metaphysically basic, not the character traits. That is, one can seemingly grasp the former without reference to the latter, but one cannot understand the latter without appeal to the former. If so, then character traits are not metaphysically basic and the non-reductivist needs

a better way of establishing this.

The second question to which Jeffrey and Beary's taxonomy responds is this:

**Counting Question:** How many things are there in the final metaphysical account of character? (2023: 188-189).

Another way to frame this question is by asking what character is composed of. Here, Jeffrey and Beary offer three answers. The first is 'monism', according to which character traits just are *identical to* the dispositions that compose the trait. On this view, Robin's trait of arrogance just is identical to his mental-state dispositions. When considered alongside reductivism, monism offers an ontologically parsimonious account of character because it provides the simplest explanation of how such an entity is composed. In short, if all character traits are identical to their trait-dispositions, then monism holds that one's character just is identical to all the dispositions one possesses.

According to 'dualism', things are not quite so straightforward. Rather than take traits to be identical to fine-grained dispositions, it views character traits as being *constituted* by these dispositions. Depending on whether the dualist takes a reductivist or non-reductivist stance towards the Grounding Question, they will either consider character traits to exist *in virtue of* the dispositions but emerge as an independent quality over and above the sum of these dispositions (reductivist); or they will take the dispositions and traits to both exist but contend that neither is identical to, nor grounds, the other (non-reductivist) (Jeffrey and Beary, 2023: 189).<sup>42</sup>

Character monism is Miller's preferred view. This is evident from his claim that traits are 'simply identical to the mental-state dispositions which underlie them' (2014: 29). In defence of this claim, Miller emphasises that mental-state dispositions are the causal basis of character traits. But if each trait is metaphysically distinct from its underlying dispositions, as dualism seems to hold, then this implies that character traits are causally inert (2014: 30). Miller believes this conclusion to be 'extremely implausible', not least because it leads to causal over-determination about people's behaviour.

There are two points worth raising at this point. First, Miller's contention relies on

---

<sup>42</sup> Jeffrey and Beary (2023: 189) include a third response to the Counting Question, which they call 'hylomorphism'. This is the view that character traits exist in their own right, such that the mental-state dispositions associated with them 'exist as material parts' of the trait, and that 'the goals or orientations of the agent exist as formal parts' of the trait. For purposes of space, I omit this view from the discussion because it has no bearing on my eventual picture of intellectual character. That is not to suggest, however, that this view is unimportant.

mental-state dispositions actually being the causal basis of character traits. Below in section 3.3, I propose a view of character reductivism that denies this. Second, Miller leaves it an open question as to whether traits are identical to, or constituted by, dispositions (2014: 32). I argue below that dualism is not only far more plausible than Miller gives it credit, but that his monism is particularly vulnerable to the *Ontological Problem* in a way that my dualist ontology is not.

## 3.2 The Dilemma for Reductivists

I now turn to presenting Jeffrey and Beary's dilemma for character reductivism. The first horn of the dilemma turns on what they call the 'Differentiation Requirement', according to which an account of character must explain why some dispositions and not others constitute a particular trait (2023: 195). Take a trait like honesty. The reductivist will understand this in terms of its fine-grained mental-state dispositions, such as dispositions to form desires to avoid stealing even when one can get away with it or dispositions to form beliefs that one should not lie when it promotes good feelings (*ibid*). To satisfy the Differentiation Requirement, reductivists need to explain why *these* dispositions belong to, or supervene on, honesty and not some other trait. The problem with reductivism, they claim, is that it cannot satisfy this requirement without appealing to the trait in question, and hence the account is circular.

**First Horn:** If the reductivist appeals to the trait in their account of why certain mental-state dispositions supervene on the trait rather than some other trait, the account is circular (2023: 195).

But even if reductivism can avoid this circularity, Jeffrey and Beary contend, it will still face the second horn of the dilemma.

**Second Horn:** To explain or identify the relevant dispositions without appealing to the trait, reductivism will produce counter-intuitive verdicts in a range of cases where a person appears to exhibit a given trait but lacks one or more of the dispositions that compose it (2023: 196).

More specifically, they claim that reductivism leads to *false positives* and *negatives*: on the one hand, it results in cases in which a person possesses all the dispositions that comprise the

trait, yet who intuitively does not possess the trait. Similarly, they contend that reductivism leads to cases in which a person seemingly possesses the trait but does not have the specified dispositions associated with the trait. Jeffrey and Beary attempt to demonstrate this dilemma by considering Miller's (2018) account of generosity.

In keeping with his earlier commitment to reductivism, Miller breaks down the trait of generosity into a number of necessary component-dispositions. To possess this trait, he claims that one will be disposed to: (i) value the item or object that is bestowed by the action; (ii) be motivated by an ultimate desire that is altruistic or at least one that is primary in generating these actions in the absence of other motives; and (iii) perform these actions in a supererogatory way (2018: 231). Jeffrey and Beary acknowledge that Miller's account avoids the first horn of the dilemma because the dispositions make no reference to the trait of generosity. However, they argue that it falls prey to the second horn by generating false negatives and positives. An example of a false negative, they think, is a case like the following:

**Sergio:** A local restaurant owner, Sergio, is fiscally conservative with his business's earnings. He is committed to helping people in his community but wants to make sure he does not give so much of profits to local charity or individuals in need that it is costly to his business. Further, he does not himself value a life with luxury items, as he thinks a good life is one in which a person contributes significantly to the community. Sergio comes up with a plan to give small amounts on a regular basis, never giving an amount of great value to him or his business (2023: 199).

On Jeffrey and Beary's view, Sergio fails to satisfy conditions (i) and (ii) above. While he is committed to helping people in the community, he clearly values the restaurant's profits more and this directly informs how much money he gives and when. In turn, this makes it unclear whether his acts arise from an ultimate (or primarily) altruistic desire. Nevertheless, they think Sergio possesses the trait of generosity despite failing to satisfy the necessary dispositions. Thus, we have a false negative. Importantly, they note how any eventual reductivist account will have to stipulate sufficient conditions for possessing a trait in terms of the fine-grained mental-state dispositions. But once these conditions are set out, they contend, reductivist accounts will generate false positives like Dickens' Mrs. Jellyby from *Bleak House*.

**Mrs. Jellyby:** Mrs. Jellyby spends her time in various philanthropic pursuits, one of which is setting up a new colony in Africa. Meanwhile, her family is completely neglected. Her children are hungry, dirty, wild, and ignorant, and her husband is

driven almost to distraction by financial and domestic turmoil that stems from Mrs. Jellyby's "telescopic philanthropy" (2023: 199).

According to Jeffrey and Beary, Mrs Jellyby meets all the necessary conditions of Miller's account of generosity, but we would likely refrain from describing her as "generous". So, his account seems to generate a false positive. In sum, they believe that any reductivist account of character will inevitably face either of the two horns of the dilemma, and therefore will fail to provide an adequate, non-circular view of behaviour. Considering that I am relying on Miller's reductivist mixed-trait model of traits, the question facing us now is whether this dilemma would afflict my account of intellectual character. I argue that it does not and propose how character reductivism in general can avoid this dilemma.

### 3.3 An Attitudinal Framework of Character

The apparent dilemma facing character reductivism results from its commitment to viewing character traits as being *reducible to* or *grounded in* the fine-grained mental-state dispositions that compose the trait. The best way to avoid this dilemma, therefore, is to offer a revised story about the causal basis of character traits. In this section, I will begin this by arguing that attitudes are the causal basis of character traits. This involves first showing how attitudes share a number of qualities with character traits, which requires us to unpack the notion of an attitude in some detail.

#### 3.3.1. Introducing Attitudes

Within social psychology, attitudes are typically treated as summary evaluations of objects that manifest as either likes, dislikes or ambivalence. The objects in question are those things that can be experienced, thought about, including even some propositions. In this way, psychologists tend to understand attitudes as psychological states of people that consist of one or more valence-associations, which is to say that they consist of positive and negative affects towards the target object when construed as a mental representation (Eagly and Chaiken, 1993; Fazio and Olson, 2007; Maio and Haddock, 2018). In the epistemic realm, the objects in question will often be practices of questioning, explaining, evaluating, interpreting, or states like truth, knowledge, and understanding. For instance, when I mentally represent

the idea of doing research for this thesis, this triggers in me a positive association about a particular object (research) to the effect that I like this activity. In short, I have a positive attitude towards research.

What accounts for this attitude? One factor is its *content*. Attitude content can be understood broadly as the informational basis from which the attitude is derived, which itself consists of the positive and negative considerations I have weighed up in relation to undertaking research. My attitude is just a summary of these considerations. In this way, it acts as a cognitive shortcut by removing the need for us to constantly re-consider a given object (Banaji & Heiphetz, 2010; Fazio & Olson, 2007; Tanesini, 2021). According to Maio and Haddock (2018: 25-26), the content of an attitude involves three components: *cognitive*, *affective*, and *behavioural* (CAB). The cognitive component refers to the beliefs, thoughts, and attributes we associate with a given object; the affective refers to the feelings and emotions evoked by the object; and the behavioural element refers to the past experiences and behaviours related to the object. So, the information basis of my attitude towards research is informed by my evaluative belief that research is a rewarding task, feelings of success and accomplishment when finding things out, and recalling times when research has afforded me these experiences.<sup>43</sup>

I noted above how attitudes can serve as cognitive shortcuts by removing our need to constantly re-consider particular objects. Given the time and effort this saves, this is just one of the ways in which attitudes function to serve an array of human needs. We can now distinguish three further functions that attitudes are said to play (Tanesini, 2021: 54).<sup>44</sup> The first is ‘knowledge’, which turns on how we form and revise attitudes to help us make sense of the world. If we understand attitude function in motivational terms, then the content of our attitudes is a result of our motives to have an accurate view of target objects (*ibid*: 54). The second function is ‘instrumental’, which attitudes serve when their formation, revision, and sustenance is a response to an object’s utility. Since an attitude is a summary of its informational basis (evaluative beliefs, affective states, behaviours), the motives relevant here will likely influence the evidence on which we base our summary evaluations of an object’s utility (Kunda, 1990; *ibid*, 55). The third is ‘value-expressive’, which reflects the role or purpose an attitude plays in giving expression to a set (or sets) of values I hold dear. The

---

<sup>43</sup> As I note below, the content of my attitude is not static; like attitudes more generally, it can strengthen or weaken in response to new experiences, information, or evidence.

<sup>44</sup> Here, I borrow Tanesini’s (2021: 54-58) way of carving up attitude functions. Note, however, that I omit several other attitude functions to which she appeals, including object-appraisal, ego-defensive, and social adjustive, since these are less essential to my aims here. Though, it will become clear that my eventual ontology owes much to Tanesini’s.

object need not be an explicit value such as my strong attitude towards women having bodily autonomy; it can be as simple as my positive attitude towards my bicycle, part of which is generated by the values it represents such as freedom, exercise, and adventure.<sup>45</sup>

The last feature of attitudes I shall mention is their *strength*. In chapter 2, I will elaborate more fully on this dimension as it relates to epistemically corrupting conditions, but for now we can identify two factors. First, attitude strength can be a measure of their *accessibility* to people. At least with respect to explicit attitudes, this accessibility is a function of the association between its valence(s) and our mental representation of the object. If a person responds quickly to coffee with repulsion, for example, this suggests that one's negative attitude is strong because it evokes a rapid, accessible negative association. The second measure of attitude strength is *centrality*, which refers to how close an attitude is to one's professed values, commitments, and identity (Clarkson et al., 2008; Zunick et al., 2017). This clearly bears significant parallels to my discussion of character traits being conceptually "thick" in the sense of them being 'world-guiding' in expressing our value-commitments and reasons for action.

### 3.3.2. Attitudes as the Causal Basis of Traits

With a clearer understanding of attitudes to hand, I will now demonstrate how they can be seen as the causal basis of character traits. Above, I noted that attitudes can function to express our values in a way that parallels how character traits exhibit the kind of conceptual "thickness" discussed in section 1.2. Very often, these attitudes will be highly accessible to the person possessing them, triggering valence-associations when confronted with an object that aligns or conflicts with their deeply-held values and commitments. Indeed, on the view articulated here, the values and commitments expressive of one's character traits start life as attitudes towards the relevant object, whose content congeals over time and is scaffolded by functional roles.

On this view, the trait of open-mindedness would emerge as an initial like or preference to listening to others' views and finding out information, encountering similar occasions that increasingly evoke stronger affective, cognitive, and behavioural responses of this sort, and this positive attitude serving various functions in one's life (see, Baehr, 2011; Hare, 1987; and Riggs, 2010). These might plausibly involve generating motives to help one

---

<sup>45</sup> It goes without saying that attitudes can serve more than one function.

make sense of others' arguments and gaining an accurate understanding of the world (knowledge), affecting the information and evidence one selects or avoids when considering the utility of a particular epistemic object like an argument, proposition, or piece of testimony (instrumental), and the positive attitude one holds towards their local debating club being an expression of the value they attach to open-minded discussion (value-expressive). It is not unreasonable to think that these attitudes would go on to form this person's self-conception or identity, with the trait of open-mindedness being a central part thereof.

Now, even if one were to accept this story, it would only show that the *content* of (intellectual) character traits could be based on attitudes, not that the traits themselves are *causally connected* to the attitudes in any way. Well, we can establish this by noting that attitudes can exhibit cross-situationally consistent and stable features. Since attitudes can function to express our values, we might expect them to be good predictors of behaviour, and this is indeed what the evidence suggests (Maio and Haddock, 2018: 73; Tanesini, 2021: 63; Watt et al., 2008). Of particular relevance here is the evidence suggesting that attitudes are good predictors of behaviour for those with a high need to engage in cognitive activity. This seems especially apt if we consider that those who engage in such activity will manifest character traits with a distinctive intellectual or epistemic guise.

That said, attitudes appear to be less effective at predicting the behaviour of those with a lesser or lack of a need to engage in cognitive activity (Cacioppo and Petty, 1982; Cacioppo et al., 1996). This might suggest a *prima facie* reason for thinking that attitudes cannot predict behaviour we associate with traits like intellectual laziness or close-mindedness. However, we can explain this result by appealing to a mixed-trait ontology of character. Suppose somebody dislikes engaging in cognitive activities. As a result, we should expect their attitudes towards practices like inquiry, investigations, or asking questions to be correspondingly low. Moreover, since attitudes can bias the reasons we offer in support of our behaviour (leading to motivated cognition), they can have a negative impact on the cognitive, affective, and behavioural associations we have towards such activities. If this is correct, though, then we can predict from this that such a person will not engage as effectively in epistemic practices compared to those with a high need to engage in cognitive activities.

Using the mixed-trait model, we can now say that these negative attitudes towards such practices are only likely to strengthen under conditions that do not encourage their engagement, forming a range of mental-state dispositions to think, feel, and act in particular ways towards these practices. The stronger these attitudes become, the more prevalent these thoughts, feelings, and actions will manifest towards the epistemic practices in question. Simultaneously, the more prevalent these dispositions become, the less reason one will have

to engage in these practices. In turn, this will strengthen the agent's negative attitudes towards the need to engage in cognitive activities, which again leads the relevant mental-state dispositions to congeal, and so on and so forth. If we accept this cyclical process, then what we have is behaviour that is likely to be both stable and cross-situationally consistent in a way that reflects traits like intellectual laziness or carelessness. As such, attitudes can effectively predict both good and bad epistemic behaviour.

Attitudes also allow us to make sense of the motives and emotions characteristic of many traits. Recall that attitudes are summary evaluations of objects. Part of these evaluations involve weighing up considerations relevant to how we appraise the object in light of the goals or desires we are motivated by. To borrow Tanesini's example (2021: 66), suppose a person's goal is to feel good about himself. An effective way of achieving this goal is to defend his ego against threats to the positive picture he has constructed of himself. In light of this, the person will often evaluate objects according to whether they pose a risk to his self-conception or enhance it, even if they do not fully acknowledge this fact about themselves. Here, the person bases his attitudes solely on information that is relevant to these concerns, and this generates motivations that are consistent with the mental-state dispositions associated with traits like arrogance, narcissism, or vanity. Furthermore, attitudes allow us to make sense of the affective dimension of character traits. Tanesini (*ibid*, chapter 5) argues that the trait of intellectual arrogance is based on feelings of anger, which triggers as a response to ego-defensive attitudes. On the general view offered here, then, the emotions that accompany character traits will be derived from the affective content of the attitudes on which they are partially based.

Pulling these pieces together, we can establish the following chain of causation: attitudes are summary evaluations that strengthen to form mental-state dispositions to act, think, and feel in ways relevant to the object(s) in question. As these dispositions strengthen, they cause agents to engage in or refrain from these objects, which in turn strengthens those initial attitudes. Together, this process underpins the existence of our character traits. For ease of explication, here is the causal chain of events:

Attitudes  $\Rightarrow$  Mental-State Dispositions  $\rightarrow$  Character Traits.

In this way, it is not primarily the case that mental-state dispositions cause the behaviour manifested by our character traits; instead, the mental-state dispositions *themselves* are grounded in prior attitudes as occurrent states, and it is these attitudes which are ultimately responsible for the behaviour instantiated by our traits. Thus, attitudes are not just the causal

basis of our character traits; they are the causal basis of our character development more generally.<sup>46</sup>

### 3.3.3 Solving the Dilemma

Not only is the view offered above still reductionist in nature, but it is importantly *dualist*. Since attitudes are composed of occurrent cognitive, affective, and behavioural states, they are not identical to the resulting mental-state dispositions formed. Even if we are inclined to think that traits are identical to their mental-state dispositions, they would not be identical to attitudes. The upshot is that character traits, while existing *in virtue of* these attitudes, emerge as metaphysically distinct qualities over and above the attitudes on which they are based.<sup>47</sup> This dualism, along with the view that attitudes are the causal basis of traits, allows me to sidestep the dilemma facing character reductionism. Let us reconsider the first horn of the dilemma.

Recall that this horn holds that character reductionism fails to offer an adequate response to what Jeffery and Beary term the ‘Differentiation Requirement’, which requires reductionists to explain why certain dispositions are associated with traits in a non-circular fashion. By grounding character traits in attitudes, however, my view can do just this. To borrow their example, we can explain Mother Teresa’s trait of honesty by first pointing to the fact that she has a negative attitude towards objects like thieves and liars, whose content consists of various occurrent cognitive, affective, and behavioural states: she might believe that lying is bad, feel sadness or disappointment about those who steal, and past experiences with thieves and liars may evoke a negative association towards these objects. Second, these negative attitudes will likely play a number of functions in her life, principally expressing her value-commitments to telling the truth and duly paying for goods. Over time, these attitudes will cluster into mental-state dispositions to form beliefs that telling the truth is better than

---

<sup>46</sup> The account offered here bears some resemblance to Machery’s (2016), who argues that attitudes are traits such that the causal base of the latter is found in the content of the former. However, it differs in a crucial respect. His view ultimately takes traits to be identical to attitudes because traits derive their causal basis from attitudes. I claim that traits are constituted by, but non-identical to, attitudes because the content of the attitude is fundamentally different from the mental-state dispositions which comprise the trait. This is important because it avoids circularity in defining the trait, something to which Machery’s account is vulnerable. Second, if attitudes are identical to traits, his view will have a hard time explaining which attitudes pick out the relevant traits without referencing the traits themselves. Hence, it leads back to the first horn of the dilemma above. For these reasons, my account offered here seems preferable.

<sup>47</sup> We might think of this relationship in terms of supervenience, that is, any change in the relevant attitudes necessitates a change in mental-state dispositions, and ultimately the character trait.

lying and desires to avoid engaging in these activities. These mental-state dispositions then correspond to the trait we call honesty.

Notice that this explanation of Mother Teresa's behaviour does not appeal to the trait of honesty. Instead, it understands her trait as multi-layered, beginning with the summary attitudes that strengthen to form the resulting mental-state dispositions. This is important because Jeffrey and Beary (2023: 195) claim that an account of character should be able to 'provide a principled way to partition the fine-grained dispositions in the reduction base of a trait, as well as an explanation for the partition'. By appealing to attitudes, we can see that the view offered here satisfies both of these requirements. More specifically, it partitions a trait's fine-grained mental-state dispositions on account of the summary attitudes from which they are formed, and this partition is explained in terms of the cognitive, affective, and behavioural content of those attitudes respectively. The result is that my view offers a non-circular response to the Differentiation Requirement and therefore avoids the first horn of the dilemma.

What about the second horn? Again, this says that any reductionist account which identifies dispositions without reference to the trait is bound to generate false negative and positive counter-examples (2023: 196). Interestingly, this horn does not immediately arise for my account because it does not ground traits in dispositions. What matters on my view is whether an agent has formed the occurrent attitudes towards a particular object, since these are what ultimately lead to the formation of the relevant mental-state dispositions. As such, my account is not beholden to whether an agent possesses one or more dispositions necessary to have a trait, nor whether an agent intuitively possesses the trait but lacks the relevant dispositions. In postulating that character traits are constituted by, and not identical to, attitudes, my reductionist account of character side-steps intermediary questions about dispositions. This reveals two important things. First, Jeffrey and Beary's dilemma only really affects forms of character reductionism that are committed to monism, such as Miller's. Any account that places attitudes at the core – and hence is dualist – will avoid both horns. Second, it raises the question of whether there really is a dilemma confronting character reductionism. If the arguments in this section are correct, there is not.

## 4. Conclusion

My aim in this first chapter has been to familiarise the reader with what is arguably the focal point of responsibilist virtue epistemology: intellectual character. I started in section 1 by contrasting the notion of “character” with that of its cognate in “personality”. The discussion there revealed the former to be more closely related to responsibility, as well as being a normatively ‘thicker’ concept than the latter. I then ended by briefly sketching Aristotle’s view of moral virtue as the basis for thinking about character.

In section 2, I turned my attention to recent empirical evidence from social psychology that has been somewhat of a thorn in the side of virtue theory. This “situationist” evidence attempts to cast doubt on the idea that we possess anything like robustly global traits of character, suggesting instead that our moral and epistemic conduct is more heavily influenced by situational factors. After considering the balance of evidence, I suggested virtue theorists should attempt model character in an empirically-informed way and I offered Miller’s Mixed-Trait ontology as an example.

With these preliminaries in order, I then tackled what I called the *Ontological Question* in section 3. At its core, this question dealt with a metaphysical dilemma facing any account of character that is reductivist in nature, that is, one which claims that traits are reducible to mental-state dispositions. To avoid this dilemma, I argued that attitudes are the causal basis of character traits rather than mental-state dispositions. The upshot was that my account was not only able to side-step this dilemma, but that an attitude framework of character brought the very existence of this dilemma into question.

In a sense, this opening chapter is the most positive in its outlook: it provides the basis for thinking about how we might go about theorising character in relatively ideal circumstances. In the remaining chapters, this positive, idealised way of theorising gives way to a non-ideal approach to how different social conditions can degrade, malform, and ultimately corrupt our intellectual characters. This starts now with a critical discussion of the phenomenon with which this thesis is concerned: *epistemic corruption*.



## 2.

# Varieties of Epistemic Corruption

## Introduction

The previous chapter was largely theoretical in presenting an empirically-robust ontology of character. To borrow Robin Dillon's phrase, the ontology I offered provided us with insights about 'what is inside the individual' (2012: 90). There, I argued that our (intellectual) character traits are grounded in attitudes. One reason why it is worth asking abstract questions about the ontology of character is because it can help virtue epistemologists understand what goes into the development of intellectual character, and especially the development of *virtuous* intellectual character. Another reason, however, is to understand the ways in which this development can go wrong. As Dillon (2012: 104) is keen to stress, our characters are not simply a matter of what is *inside* of us, but also how cultural, social, and political structures, practices, ideologies, relationships, and activities shape and constitute our characters. In short, character is as much external to individuals as it is internal, and a good character theory will transcend this distinction.

Having pieced together a picture of intellectual character in the previous chapter, I now turn my attention to the task which will occupy me for the remainder of this thesis: articulating an account of *epistemic corruption*. As the title of this chapter suggests, there is more than one way of understanding this phenomenon. Indeed, contemporary virtue epistemology contains the resources to distinguish at least three varieties of epistemic corruption. My aim in this chapter is to distinguish and identify these varieties from each other. By doing so, not only will we be better placed to appreciate the dynamic ways in which

epistemic corruption can operate; we will come to see just how susceptible our intellectual characters are to this phenomenon. This will be particularly helpful in setting up the finer mechanics of this notion in the next chapter.

In section 1, I start by focusing on what is arguably the most dominant understanding of epistemic corruption in the literature to-date, namely the idea that epistemic corruption involves on the erosion of an agent's intellectual virtues or the development of their intellectual vices. I refer to this as the *Aretaic conception*. In section 2, I draw attention to a variant of epistemic corruption that concerns an agent's (in)ability to exercise what Casey Johnson (2020) has called 'epistemic phronesis'. For this reason, I refer to this variant as the *Phronetic conception*. Finally, in section 3, I identify a hitherto unexplored way of understanding epistemic corruption that falls out of my ontology of character from the previous chapter. Since it trades on an agent's intellectual traits and attitudes, I call it the *Attitudinal conception*. In section 4, I conclude.

## 1. The Aretaic Conception

For the most part, discussions of epistemic corruption have served to pick out or identify the range of social conditions that either (i) encourage the development and exercise of an agent's epistemic virtues or (ii) facilitate the loss or erosion of their epistemic virtues (Kidd, 2017, 2018, 2019, 2020). Indeed, one would be easily forgiven for thinking that the notion exists purely as a univocal phenomenon. As we shall see shortly, however, initial appearances can be misleading; epistemic corruption can take on a host of different, if not overlapping, forms. In this section, I shall isolate and supplement what is arguably the most dominant way of understanding the phenomenon.

In the introduction to this thesis, recall, I noted how the concept of epistemic corruption was first introduced into the vice-epistemological vernacular by Ian James Kidd. Since then, Kidd has come to refer to it as 'character corruption' or a 'character-based' account of corruption (2021). This employs a sense of corruption that refers to the 'gradual deterioration or destruction of the moral and intellectual character of human beings' (*ibid*: 331). The kind of deterioration or destruction in question is a product of social conditions that lead to the erosion of a person's virtues and the 'facilitation of their vices' (*ibid*: 331). As this suggests, the value-neutral nature of 'character corruption' allows Kidd to adopt a pluralistic stance towards the kind of virtues and vices deteriorated and facilitated. That is,

‘character corruption’ is attentive to the loss and gain of both moral and intellectual virtues and vices.<sup>48</sup> Since my focus here is on the *epistemic* version of this phenomenon, I shall largely bracket concerns about moral character. Though, this need not commit us to abandoning moral values entirely.

While ‘character corruption’ might be an intuitive term, I think there are at least two reasons why a more fitting name for Kidd’s variant should be the Aretaic conception. The first is that he cashes out epistemic corruption primarily in terms of intellectual virtues and vices. When describing how the manufacture of doubt can be epistemically corrupting, for example, he observes how this process can create conditions ‘conducive to the development of *epistemically vicious dispositions*’ (Biddle, Kidd, and Leuschner, 2017: 173). In his ‘Deep Epistemic Vices’ paper, epistemic corruption is defined in terms of ‘experiences or activities that promote the development and exercise of *epistemic vices* and/or fail to encourage the cultivation and exercise of the *epistemic virtues*’ (2018: 48-49). Educational systems, for instance, are described as epistemically corrupting insofar as they ‘create conditions that are conducive to the development and exercise of *epistemic vices(s)* by agents whose formation and agency is shaped by those conditions’ (2019: 224). Socially oppressive conditions can be epistemically corrupting, Kidd emphasises, when an agent’s character ‘comes to be damaged due to one’s interactions with persons, processes, doctrines, or structures that facilitate the development and exercise of *epistemic vices*’ (2020: 71).

The second reason trades on the *kind* of epistemic virtues and vices salient to the Aretaic conception. Recall from the Introduction that there are two broad approaches to conceptualising epistemic virtues and vices. Virtue responsibilists conceive of them as praiseworthy/blameworthy features of intellectual character whilst virtue reliabilists understand them as reliably truth-conducive/inhibiting qualities of agents. Of these qualities, though, reliabilist virtues and vices tend to be cognitive faculties like good eyesight or bad hearing. Whilst Kidd maintains that epistemic corruption is ‘neutral’ with regards to these conceptualisations (2019: 225), his decision to employ terms like ‘character corruption’ suggests that it is actually partial to the former.<sup>49</sup> What Kidd seems to have in mind when he talks of ‘reliabilist’ virtues and vices are those character traits or attitudes that either reliably

---

<sup>48</sup> It is plausible to think that ‘character corruption’ might also extend to aesthetic, spiritual, political, and prudential virtues and vices, too. It can also extend to hybrid virtues and vices such as intellectual-ethical (Fricker, 2007; Williams, 2002) and ethico-epistemic (Daukas, 2018).

<sup>49</sup> An additional factor is that Kidd’s concern for epistemic corruption emerged from a prior concern with ‘epistemic edification’ – the way in which our social environments can encourage the development and exercise of epistemic *character* virtues. See Kidd (2016b) and Cooper (2010) for discussions of edification.

produce good epistemic effects (true beliefs, knowledge etc.) or reliably produce bad epistemic effects (false beliefs, obstruct knowledge, etc.). In other words, the kind of reliabilist virtues and vices of interest to him are those found in consequentialist analyses of virtue and vice such as the work of Julia Driver (2000, 2001) and Quassim Cassam (2016, 2019a).

For epistemic corruption to be neutral in this regard, it must be able to accommodate the fact that reliabilist virtues and vices are primarily cognitive faculties and defects. An initial worry on this front, though, is that we seem to have little control over the exercise of our cognitive faculties in the way we do with our character traits. This makes it harder to capture the sense in which our social environments might “encourage” the development of bad eyesight or “facilitate” the loss of good hearing, for example.

However, I think this concern is misplaced. On the one hand, it is often a matter of ‘constitutive luck’ which character traits we end up forming (Card, 1996; Nagel, 1979; Williams, 1981). As Heather Battaly (2016: 100) notes, young men who grew up under the influence of the Taliban regime were at particular risk of becoming dogmatic and indoctrinated. Yet the fact these men became dogmatic was hardly an active *choice* they made; it was something largely beyond their control.<sup>50</sup> In this way, the suggestion above presupposes too much control over our character traits. On the other hand, there do seem to be clear cases in which our social environments erode the functioning of our cognitive faculties. Consider a faculty like good memory, which is a paradigmatic reliabilist virtue. There is substantial evidence to suggest that exposure to highly stressful conditions can significantly impair one’s memory (Joëls et. al., 2006; Schwabe and Wolf, 2010; Vogel and Schwabe, 2016).<sup>51</sup> If stressful conditions impair our memory, there is not only a sense in which they facilitate the loss of a reliabilist virtue in good memory; there is a further sense in which they encourage poor memory. Given that poor memory typically leads one to reliably form false beliefs about events, those conditions also foster the development of a reliabilist epistemic vice.

While this strictly takes epistemic corruption beyond intellectual character, it is worth emphasising the interlocking relationship between responsibilist and reliabilist virtues and vices. By impairing our memory, stressful conditions are also likely to have a negative effect on our capacities to exercise traits like attentiveness, epistemic conscientiousness, and intellectual carefulness to name a few. To the extent that our memory fails us, we are less

---

<sup>50</sup> This does not mean, of course, that we cannot criticise these men for their dogmatism. See Cassam (2019a) and Battaly (2019) for insightful discussions about the prospects of holding agents responsible for their epistemic vices.

<sup>51</sup> For an overview of these empirical findings, see Piefke and Glienke (2017).

likely to pay attention to certain people or propositions, or to recall important details and information. Indeed, we might even find ourselves overlooking things in a way that borders on epistemic recklessness. By accommodating responsibilist and reliabilist virtues and vices, then, the Aretaic conception not only proves itself to be a truly neutral variant of epistemic corruption; it begins to capture the important interplay between these two sets of qualities.<sup>52</sup>

## 2. The Phronetic Conception

The Aretaic conception conceives of epistemic corruption in terms of social conditions that facilitate the loss or encourage the development of an agent's epistemic virtues and vices. However, there is far more to a person than their virtues and vices, which suggests that there are more ways in which epistemic corruption could theoretically occur. Across the next two sections, I will consider two such variants.

The first of these has recently been identified by Casey Johnson (2020: 235), who claims that we can 'understand contemporary schooling in the United States as epistemically corrupting'. Whereas Kidd's focus is on conditions that are conducive to the development of epistemic vices and the loss of epistemic virtues, Johnson is more interested in how education practices can 'discourage epistemic phronesis' (2020: 227, 235). For this reason, I refer to it as the Phronetic conception. To fully appreciate this variant of epistemic corruption, it will help to briefly unpack the notion of *phronesis*, or practical wisdom.<sup>53</sup>

Aristotle traces practical wisdom back to his separation of the soul into the rational and appetitive parts. Within the rational part, he makes a further distinction between the *scientific* (theoretical) and *calculative* (practical) parts (*NE*, VI, 1138b-1139a22). Whereas the theoretical dimension relates to universal truths such as those in the sciences or mathematics, the calculative aspect concerns particulars of a situation and involves a 'reasoned state of capacity to act with regard to the things that are good or bad for man' (*NE*, VI, 1140a20-b6). Part of this calculative capacity is the power to *deliberate well*, which involves making the correct choices not only about one's ends, but also concerns the means by which one achieves these ends (*NE*, VI, 1142b11-b30). The ability to do this is what Aristotle calls *phronesis* or the state of being practically wise. Given that practical wisdom involves a deliberative capacity to identify the 'correct' means for action, it is reserved solely for virtue; it 'aims'

---

<sup>52</sup> In chapter 5, I shall return more explicitly to the relationship between reliabilist and responsibilist vices.

<sup>53</sup> Henceforth, I use the two terms interchangeably.

towards virtue by helping a person discern which aspects of a situation will help them achieve their end and those they are better off not pursuing. In this way, practical wisdom guides or acts as the ‘rational principle’ for our actions across situations, helping us to consistently hit the mean of virtue and achieve well-being (Roberts and Woods, 2007: 306).<sup>54</sup>

An important ingredient of practical wisdom is a perceptual acuity to identify the best course of action to take. This is not simply the ability to visually see something or some particular, but rather is a trained and calibrated sensitivity to salient features of a morally- or epistemically-charged situation. These salient features cause us to act in a particular way. For instance, those who are benevolent are likely to spot opportunities to help others, while those lacking this trait may forego similar opportunities if they arise. Often, this sensitivity is directly informed by past experiences and a range of background interpretive values that contribute to our moral and ‘epistemic socialisation’ (Fricker, 2007: 72). Moreover, this perception is likely to be non-inferential such that we automatically form evaluative beliefs about what to do without considering it. In Nancy Sherman’s words (1991: 40), we engage in the process of ‘seeing as’ – interpreting the relevant circumstances that lead to spontaneous action and judgement. As with much of our moral and intellectual conduct, though, this perceptual capacity is not purely a cognitive ordeal; very often, our emotional investments or vulnerabilities will determine what we see – or fail to see – in a given situation (Bommarito, 2017; Brady, 2014).<sup>55</sup>

Building on this, Johnson takes educational conditions to be corrupting insofar as they discourage epistemic phronesis, which she says involves discouraging ‘the kind of accurate perception and experiential reflection required for someone to develop *phronesis*’ (2020: 235). Educational conditions can bring this about, she contends, by failing to provide students with the opportunities to recognise that education is a context in which intellectual virtues can be brought to bear. In turn, this inhibits those students from manifesting various intellectual virtues within the classroom, even if they might otherwise instantiate them outside of the classroom. As Johnson puts the point:

---

<sup>54</sup> As Aristotle conceives of it, practical wisdom is subsumed within the broader faculty of *cleverness*, which generically enables us to ‘tend towards the mark we have set before ourselves, and to hit it’ (*NE*, VI, 1144a13-1144a18). I return to consider the role of cleverness in chapter 5.

<sup>55</sup> My construal of this perceptual component is largely inspired by Nancy Sherman’s (1991) account of practical wisdom.

Many students don't realise that educational contexts are appropriate for curiosity, humility, and other epistemic virtues because in their test-based experience, that development is neither recognised nor rewarded (2020: 234).<sup>56</sup>

In this way, educational settings – and specifically test-based settings – are epistemically corrupting, but not because they encourage students to develop intellectual vices, nor because they facilitate the loss of any intellectual virtues. Instead, they are corrupting insofar as they prevent students from perceiving that education is a context in which it is *suitable* to exercise and manifest their intellectual virtues.<sup>57</sup> So, even though students might have at their disposal the intellectual skills and traits to manifest a given virtue, their educational settings stop them from putting two and two together and making virtue, as it were. In Johnson's terms, these students lack epistemic *phronesis*.

At this point, we need a number of clarifications about the Phronetic conception. First, given the close connection between phronesis and virtue, one might question whether what I am calling the Phronetic conception simply boils down to the Aretaic conception. More specifically, if epistemic corruption occurs on this variant when conditions prevent agents from identifying when to exercise their epistemic virtues, is this not simply a case where conditions facilitate the loss of virtues? This is a complex question, in part because the relationship between practical wisdom and virtue itself is contested. Some theorists take the former to be both necessary and sufficient for the latter (McDowell, 1998); others, meanwhile, consider practical wisdom to be necessary but insufficient for virtue (Bloomfield, 2013). Whether phronesis is both necessary and sufficient for moral virtue, it seems insufficient for *epistemic* virtue. While a person needs to have good judgement as to when to manifest a given epistemic virtue, they also need to be motivated by epistemic goods. If they are not, they are, at best, only manifesting the relevant intellectual traits. In this way, the Phronetic conception only focuses on an aspect of virtuous activity, not virtue itself.

Second, Johnson's variant resembles a proposal that Kidd himself anticipates. One of Kidd's general observations about the concept of corruption (and edification) is that it is 'not confined to epistemic virtues and vices' (2019: 223). Instead, it might be that an educational system can 'produce students with a proper sense of the nature and significance of epistemic virtues but fail to instil a further commitment to be epistemically virtuous' (*ibid*: 223). What Kidd is suggesting here is that educational systems can be corrupting in the sense that they

---

<sup>56</sup> Not just that, they are entirely absent from schools in that they have no interest in students developing virtues.

<sup>57</sup> One example is Jason Baehr's Intellectual Virtues Academy, designed to cultivate virtues and edify students' characters. See his *Deep in Thought* (2021).

encourage *akrasia* or weakness of will: students *know* what is needed to manifest a given virtue but are ‘not yet disposed to pursue or manifest it’ (*ibid.*: 223). Although similar, Johnson’s claim is arguably stronger than this. Her worry is not so much that educational settings will induce a kind of weakness of will in students, since this implies that the students are aware or at least recognise that education or schooling is a place in which they ought to exercise their intellectual virtues but fail to do so. Instead, Johnson’s principal concern is that particular educational settings can fail to make students even *aware* that they ought to be able to exercise intellectual virtues in these environments. This much takes the Phronetic conception beyond Kidd’s remarks.

Third, we might well wonder what work the “epistemic” qualifier is doing in Johnson’s notion. I raise this query because *phronesis* is said to derive from the rational part of the soul, which would squarely make it an intellectual virtue on Aristotle’s view. However, Aristotle is also clear that character virtue is a ‘state involving rational choice, consisting in a mean relative to us and determined by reason – the reason, that is, by reference to which the *practically wise* person would determine it’ (*NE*, bk. II, 1106-1107a). If character virtue is determined by practical wisdom, which itself is an intellectual virtue, then it looks as though character virtue requires an intellectual virtue. Indeed, Aristotle highlights this connection when he writes that ‘we cannot be really good without practical wisdom, or practically wise without virtue of character’ (*NE*, bk. VI, 13, 1144b). This would suggest that “epistemic” *phronesis* is a partially redundant term because *phronesis* is precisely the intellectual virtue which governs the responsibilist character virtues of interest to Johnson. Thus, “epistemic” *phronesis* is simply *phronesis*.

Now, I suspect Johnson includes this qualifier because she attempts not only to capture responsibilist character virtues, but reliabilist virtues as well. For example, she claims that a person will be an epistemic *phronimos* on a virtue reliabilist understanding just in case she ‘knows which belief-forming process to use in which context’ (*ibid.*: 229). She cites as reliabilist virtues belief-forming processes like ‘attentiveness, good memory, and acute reasoning’ (*ibid.*: 229). Mirroring her treatment of the responsibilist virtues, her contention would be that educational conditions are epistemically corrupting insofar as they prevent agents from recognising when to exercise their attentiveness, good memory, or acute reasoning when they are well-placed to do so. In this way, Johnson’s variant would accommodate reliabilist virtues much like my expansion of Kidd’s Aretaic conception.

While this is a highly innovative move, I think we should be careful about how we set it up. Although I expressed doubts about the extent to which we can control the character traits we develop, I do nonetheless accept that the exercise of these traits requires a level of

volition and deliberation that is lacking when it comes to qualities like good memory or acute reasoning. This partially underscores why we are often hesitant to blame people for their poor memory or reasoning abilities. Equally, it seems to make little sense to think that we can know when to use these belief-forming processes. After all, I cannot choose which cognitive faculty will provide me with true beliefs if I am listening to music – the very act of listening to music dictates that my reliable hearing does this. This suggests that the exercise of reliabilist virtues like good hearing is passive and automatic unlike the responsibilist virtues.<sup>58</sup> For this reason, I am unsure whether the former qualities can be subject to practical wisdom in a way the latter qualities are, and therefore whether the Phronetic conception can accommodate them.

Although this might rule out cognitive faculties from the Phronetic account of epistemic corruption, it need *not* rule out reliabilist virtues in principle. This is because reliabilist virtues, recall, are simply any truth-conducive dispositions of agents, and our cognitive faculties are but one class of truth-conducive dispositions we possess. Interestingly, Casey seems to recognise this. When it comes to the belief-forming processes relevant to the ‘virtue reliabilist’s understanding of *phronesis*’, the processes in question quickly turn out to be ‘questioning and challenging’, both of which are intellectual *abilities* (*ibid.*: 231). This is further supported by her contention that epistemic phronesis puts one in a position to ‘better know which intellectual skills and traits to employ’ and that ‘a person might have the skills involved in an epistemic virtue ... yet lack *phronesis* and so fail to deploy these virtues in formal educational contexts’ (*ibid.*: 229-230). Intellectual skills like being good at asking questions or challenging one’s interlocutors are not only qualities we can develop over time, but whose exercise we have control over.

If we understand reliabilist virtues as intellectual skills, then, the Phronetic conception of epistemic corruption can certainly accommodate these qualities. An upshot of this view would be a kind of “skill-based” variant of epistemic corruption, according to which conditions are corrupting insofar as they discourage agents from developing and exercising the intellectual skills related to epistemic virtues.<sup>59</sup> Worse still, these conditions might never present agents with opportunities to develop these skills. Indeed, they could lead these agents

---

<sup>58</sup> See Baehr (2011, ch. 3) for an in-depth discussion of the structural relationship between responsibilist and reliabilist virtues.

<sup>59</sup> One might worry that such a conception takes the focus away from an agent’s intellectual character, since skills are not, strictly speaking, a feature of a person’s character as much as they are developed capacities. However, this worry is overly restrictive. As many have observed, the development of various intellectual skills often occurs in tandem with our epistemic traits, and the cultivation of these traits is only complemented by these skills (Baehr, 2011, Watson, 2018; Zagzebski, 1996). As such, a focus on an agent’s skills is also an indirect focus on their character (mal)development.

to exercise the intellectual *failing* of ‘bad questioning’, an inability to either elicit worthwhile information or direct their questions at the right people (Watson, 2020). Here, the environment erodes a reliabilist epistemic virtue and facilitates the development of a truth-inhibiting *reliabilist vice* in bad questioning. In chapter 5, I shall return to develop this theme in greater detail by drawing attention to the kind of skills that distinctly attach themselves to various epistemic vices, or what I call *vice-indexed skills*. For the meantime, let us consider a third conception of epistemic corruption that emerges directly from my discussion in the previous chapter.

### 3. The Attitudinal Conception

So far, I have examined two variants of epistemic corruption that have received varying attention in the virtue-epistemological literature. In this section, I shall consider one that has not yet been explicitly addressed. Interestingly, if the ontology of character I developed in the previous chapter is correct, then it will turn out that this variant underpins the two conceptions already considered. I refer to this variant as the Attitudinal conception.

Whereas the Aretaic and Phronetic conceptions focused on an agent’s propensity to develop and exercise epistemic virtues and vices, this conception shifts the focus to one’s intellectual *traits*. In many cases, after all, the loss of an epistemic virtue or the development of an epistemic vice will correspond to the weakening or strengthening of a particular intellectual *trait*. For this reason, we can think of the Attitudinal conception as taking a zoomed-out approach to epistemic corruption. One benefit of taking such an approach is that it allows us to come to terms with the idea that developing and manifesting globally robust epistemic virtues and vices is often a tough ask. As I argued in the previous chapter, this possibility need not commit us to thinking that people cannot possess epistemic virtues and vices – far from it, in fact. But rather than jump to the conclusion that a person’s intellectual conduct is automatically a product of their epistemic virtues or vices, it does mean that we first step back and consider their broader character development and social context.

It is at this point that the Attitudinal conception comes into its own. Instead of taking epistemic corruption to involve conditions that encourage the development of intellectual vices or discourage the exercise of intellectual virtues, it understands the phenomenon in terms of those conditions that either encourage the development or loss of intellectual traits. In light of the ontology of character I offered in the previous chapter, we can now appreciate why I refer to this variant as the Attitudinal conception. Insofar as attitudes are the causal

basis of character traits, any social conditions that encourage or discourage the development of intellectual traits will first encourage or discourage the development of prior *attitudes* towards particular objects – be that values like truth or practices such as inquiry or education. As one continually engages with these goods or practices (or fails to engage with them), one’s summary evaluations of these objects will strengthen (or weaken), which then leads the agent to form the corresponding mental-state dispositions to think, feel, and act. In turn, these dispositions provide agents with recourse to engage with (or avoid) these values and practices. The result is a set of behaviours that we correlate with recognisable (intellectual) character traits.<sup>60</sup>

If we think that epistemic virtues and vices often derive from intellectual character traits, then notice that the Attitudinal conception explains what is at stake in both the Aretaic and Phronetic conceptions of epistemic corruption. For the former, the problem is not so much about social conditions facilitating the loss of an agent’s virtues or encouraging the development of their vices; it is a matter of such conditions facilitating the loss and gain of certain intellectual character traits, and in turn attitudes. In the latter, it is not primarily an issue of conditions preventing students from recognising that education is somewhere they ought to exercise epistemic virtues; it is that these conditions fail to encourage in students the *appropriate attitudes* towards education. Not only does the Attitudinal conception better warrant the label ‘character corruption’, then, but it importantly serves as a broader ameliorative resource for those of us who are somewhere on the character spectrum: yet to develop fully-fledged epistemic virtues but not quite manifesting epistemic vices.

## 4. Conclusion

My aim in this chapter was to introduce the concept of epistemic corruption by mapping the contours of this phenomenon. This involved distinguishing three varieties of epistemic corruption from each other. The first of these, which I call the Aretaic conception, emerges when our social environments facilitate the loss of an agent’s epistemic virtues or encourage the development and exercise of epistemic vices. Going beyond Kidd’s proposals, I argued that this variant is able to accommodate both responsibilist and reliabilist virtues and vices,

---

<sup>60</sup> Such a view finds some uptake in the virtue-ethical and -epistemological literature. Thomas Hurka (2001) presents a ‘recursive’ account of moral virtue and vice, according to which these qualities are attitudes towards goods or evils that are, themselves, either intrinsically good or bad. Alessandra Tanesini (2021) argues that the causal basis of what she calls the intellectual ‘virtues and vices of self-evaluation’ is rooted in attitudes.

ensuring that the Aretaic conception can capture the interplay between our intellectual character virtues and our cognitive faculties.

While the Aretaic conception dominates the small literature on epistemic corruption, I argued that there is a second variant that should be of equal concern to virtue and vice epistemologists – the Phronetic conception. I suggested that this conception can, itself, emerge in two ways. Following Casey, I noted that it can occur when our social environments discourage the development and exercise of epistemic phronesis, preventing agents from recognising that a particular occasion is suitable for manifesting their intellectual character virtues. Second, it can emerge when those environments erode or discourage the exercise and development of an agent’s intellectual skills, a theme to which we shall return in the final chapter of this thesis.

Finally, I identified a hitherto unexplored view of epistemic corruption that turns neither on an agent’s epistemic virtues and vices, nor on their ability to exercise epistemic phronesis. It focuses on the ways in which our social environments affect the kind of attitudes we form towards epistemic objects. Specifically, the Attitudinal conception emerges when our social environments either encourage the development of bad attitudes towards epistemic goods or practices or erode good attitudes towards them. Ultimately, I claimed that this conception allows us to see what is at stake in both the Aretaic and Phronetic conceptions. Once we attend to these variants, not only are we able to see connections between all three, but we are also able to appreciate how they make our epistemic lives complex and susceptible to corruption in the first place. With this in mind, let us now turn to the substantive question of *how* this corruption might occur in practice.



# 3.

## The Mechanics of Corruption

### Introduction

Far from being a univocal concept, the previous chapter revealed that epistemic corruption can take on a myriad of forms. Specifically, it can occur when our social environments facilitate the loss of epistemic virtues or encourage the development of epistemic vices (Aretaic conception); it can arise when those environments discourage the exercise and development of practical wisdom in deploying one's epistemic virtues and intellectual skills (Phronetic conception); and it can also take place when our social environments encourage the development of bad attitudes towards epistemic practices and objects, or erode good attitudes thereof, leading agents to form bad intellectual traits or find their good intellectual traits dwindle (Attitudinal conception).

While the previous chapter was largely diagnostic, this chapter is more normative. It asks how we *ought* to think about epistemic corruption if we are to make it an effective ameliorative resource in vice epistemology. This requires us to now focus on what I earlier referred to as the *mechanics* of epistemic corruption. More precisely, the task before us is to examine how exactly our social environments might give rise to the phenomena captured by Aretaic, Phronetic, and Attitudinal conceptions. Building on the previous chapter, this task will provide us with a comprehensive understanding of epistemic corruption as we proceed in this thesis.

Accordingly, I start in section 1 by unpacking what I have loosely been referring to

as epistemically corrupting conditions, what these consist in, as well as the relevant features present under such conditions. In section 2, I critically discuss what Kidd has referred to as the ‘axes’ of epistemic corruption. Sections 3 and 4 are then devoted to his distinction between ‘passive’ and ‘active’ epistemic corruption. Finally, I conclude in section 5. Throughout this chapter, I raise a number of challenges for existing ways of theorising epistemic corruption and attempt to provide solutions to these difficulties. However, since I am recommending an expansion and revision of epistemic corruption here, it is largely beyond the scope of this chapter to consider whether my proposals are consistent with what Kidd (2020, 2021) has elsewhere called ‘critical character epistemology’. That said, if the revisions I propose do make epistemic corruption more effective as an ameliorative resource, this would go a long way in making my account continuous with his framework.

## 1. Corrupting Conditions

Scattered throughout this thesis so far, the reader will have noticed the phrases “epistemically corrupting conditions”, or “corrupting conditions” crop up frequently. Despite these conditions arguably forming the bedrock of any corruptionist critique, I have largely proceeded without going into detail as to what they might amount to. Given this, “corrupting conditions” will be the first mechanic of epistemic corruption I address.

At a general level, epistemically corrupting conditions can take on a host of forms. They can be *practical* (such as a lack of resources or time), *social* (including chilly environments or hostile and toxic cultures), *psychological* (stress, low morale, etc.), or a combination of all of these.<sup>61</sup> On a more specific level, though, we can extract a number of criteria that all of the above conditions will contain. The first of these is what Kidd refers to as the *corruptee(s)*. This is the agent(s) who is being subjected to the epistemically corrupting conditions: they are ‘the person or thing being corrupted’ (Kidd, 2020: 71). In many cases, the relevant corruptee will be an individual who operates under such conditions. These might include students in a particular school, civil servants in a government department, or employees within a corporation, to list just a few.

By identifying the relevant corruptee(s), however, we will often be better placed to pinpoint the corresponding *corruptor(s)* – the person or thing doing the corruption. Here, we might want to distinguish a corruptor, on the one hand, from the source of corruption, on the

---

<sup>61</sup> Of course, this list is far from exhaustive.

other. In the case above, the general source of corruption might be the school, government department, or corporation. At a more specific level, though, the relevant corruptor could be a policy or strategy that individuals like teachers, civil servants, or employees are tasked with carrying out. When the UK Home Office implemented its infamous ‘Hostile Environment’ policy towards British-Caribbean people, for example, it was the source of epistemic corruption, but the policy was the primary corruptor insofar as it encouraged civil servants to be cynical and obfuscate evidence (Gentleman, 2018; Williams, 2020).

From this brief gloss, it might be tempting to think that the corruptor(s) within epistemically corrupting conditions will always be some kind of institution, collective, or structural feature. While this might often be the case, it is worth noting that under such conditions the lines between the corruptee and corruptor can often become quite blurred. As Kidd (2020: 71) acknowledges, somebody who finds themselves under corrupting conditions can easily become complicit in their own ‘epistemic self-corruption’. Those complicit in this process will not likely set out to damage their own intellectual character; rather, it will typically occur as something the vices enable them to do – a lazy person does not intentionally seek to be lazy, they want what laziness can afford them, and this helps to sustain or entrench their vice. This helps us see why Kidd (2019: 224) emphasises that epistemic corruption has no ‘intentionality condition’, by which he means that conditions can be corrupting without that being an intention or aim of their design. For this reason, he distinguishes between conditions that are in fact *intentionally* corrupting – those purposefully designed to inculcate epistemic vices in agents or erode their virtues – and those that are *inadvertently* corrupting – those which accidentally achieve this.<sup>62</sup>

Now, it is one thing to say that conditions are epistemically corrupting, and quite another to identify the relevant qualities encouraged or discouraged under those conditions. Kidd’s prime focus, we have seen, is on an agent’s epistemic virtues and vices. If an account of epistemic corruption is to remain ‘neutral’ with respect to virtue-responsibilism and -reliabilism, though, this should be reflected in the corruptionist critiques we are offering. This requires us to carefully supplement the original criteria found within epistemically corrupting conditions. So, while responsibilist and reliabilist virtues and vices might be relevant to the Aretaic conception, they are not all accommodated by the Phronetic conception. Instead, the qualities that matters when identifying epistemically corrupting conditions will be *epistemic*

---

<sup>62</sup> Not all corrupting conditions will equally affect those subjected to them. It is an empirical question as to whether inadvertently corrupting conditions outpopulate intentionally corrupting conditions.

*phronesis* and intellectual skills, as well as the extent to which social conditions discourage agents from developing these. However, if the Attitudinal conception is anything to go by, then the relevant criteria will be an agent's intellectual traits and attitudes. Unsurprisingly, then, the relevant quality to identify under epistemically corrupting conditions will shift depending on the variant we are interested in.

Although we are clearer on this front, it still remains to be seen *how* exactly epistemically corrupting conditions might encourage or discourage the kinds of qualities above. To address this, we can turn to what Kidd calls the *modes of epistemic corruption* (2020: 72). Here are two such modes:

- (I) *Acquisition*: a corruptor can enable the acquisition of novel epistemic vices, bad attitudes towards epistemic goods and practices, or “vice-indexed” skills, ones not previously a feature of the corruptee's character. Think of a student with no prior tendencies to intellectual arrogance who, under the influence of arrogant teachers within a school that rewards arrogance, comes to develop this vice.
  
- (II) *Activation*: a corruptor can activate dormant epistemic vices, latent poor epistemic attitudes, or under-exercised skills. Imagine a different student with some underlying dogmatic tendencies that have, as luck has it, been kept below the surface. Upon transferring school, though, the student is exposed to an extreme interpretation of “independence of mind”, which activates their latent dogmatic attitudes, facilitating the development of full-blooded dogmatism.

As their respective names suggest, these modes operate when a corruptee has no pre-existing epistemic vices, bad attitudes, or vice-indexed skills. In other cases, though, a corruptee might already possess one or more of these qualities. Recognising this, Kidd articulates the following four modes of corruption.

- (III) *Propagation*: a corruptor can increase the extent to which an epistemic vice, bad epistemic attitude or trait, or vice-indexed skill affects a corruptee's conduct. Suppose Ben is a good bullshitter, but that he only bullshits to his parents. After being praised for this skill, he gradually begins bullshitting to more and more people around him and he is consumed by the web of bullshit he has woven.
  
- (IV) *Stabilisation*: a corruptor can increase the *stability* of a vice, poor attitude, or vice-indexed skill by reducing its chances of being susceptible to disruption or the development of countervailing virtues and good attitudes. When alone, Mary's

boss often reprimands her for being intellectually lazy. Due to social pressures, though, this laziness is absent in the presence of her colleagues. Unfortunately for Mary, her job is moved entirely online and she comes to work exclusively at home on her own. No longer facing any pressure, she embraces her intellectual laziness and it stabilises in the absence of her colleagues.

- (V) *Intensification*: a corruptor can increase the *strength* of a vice, bad attitude, or increase the occasions on which one exercises vice-indexed skills. Adam has a poor attitude towards all things mathematical. When confronted with equations and formulas, he shrugs his shoulders in despondency. Over time, things get worse: instead of simply shrugging, the sight of algebra makes Adam feel stressed, he comes to believe that algebra is useless, and he begins skipping his mathematics classes. While initially weak, his poor attitude turns into a stronger despondent disposition.

It goes without saying that this list is not exhaustive. In addition to those identified above, there are likely to be a myriad of other ‘modes’ by which epistemic corruption can take hold. Indeed, as we shall see now, another mode of corruption falls out of the discussion below.

## 2. The Axes of Corruption

The previous section unearthed one important layer of epistemic corruption in the kind of conditions that the phenomenon can give rise to. However, anyone wishing to offer a sufficiently robust corruptionist criticism will need to be more specific than this in their claims. To be as specific as possible, one will inevitably have recourse to the second mechanic of epistemic corruption worth discussing, namely, the *axes of corruption*.

Above, we distinguished corruptees from corruptors. Just as equally, we can distinguish between what we might call the ‘axis of corruptees’ and the ‘axis of corruptors’. Here, we can take our initial lead from Kidd, who draws attention to three particular dimensions that would characterise the former axis: the *vices* encouraged, the *agents* involved, and the *domains* affected (2019: 225). These, respectively, allow us to ascertain the strength of the corrupting conditions. Thus, Kidd claims that ‘weakly corrupting systems’ are those which ‘tend to entrench *some* vices in *some* of its members’, even if those vices are confined to

certain aspects of their character. To borrow his example, certain conditions might make me more dogmatic about political matters but leave my musical tastes unscathed (2019: 225). By contrast, ‘strongly corrupting systems’ are those that tend to ‘entrench *many* vices in *many* of their members’.

It is unlikely that epistemically corrupting conditions will fall cleanly into one or the other category. While an empirical matter, I suspect that conditions will be largely multifaceted, by which I mean they will manifest as stronger or weaker conditions. More often than not, the strength or weakness of a given set of epistemically corrupting conditions will only be as strong or weak as those subjected to them. Consider an analogy with alcohol. The strength of an alcoholic beverage is determined, in part, by the consumer. Somebody who rarely drinks will be more susceptible to a pint of beer than somebody who drinks a pint of beer every day. Just as manufacturers include general claims about a beverage’s alcohol content, so too can we make general claims about epistemically corrupting conditions. In many cases, though, we will need to investigate the conditions through a particularist lens depending on the corruptees in question.<sup>63</sup> We shall return to this point shortly.

Given the emphasis on a corruptee’s vices, Kidd’s focus is clearly on the Aretaic conception. With the different variants from the previous chapter to hand, we can now supplement these dimensions and thereby update our understanding of this axis. In addition to their epistemic virtues and vices, then, we can add a corruptee’s ability to exercise epistemic phronesis, their intellectual skills, and their traits and attitudes to the list of qualities that can be encouraged or eroded under epistemically corrupting conditions. This, in turn, alters how we might measure the strength of epistemic corruption. Accordingly, weakly corrupting conditions would now include environments that tend to discourage epistemic phronesis in some of its members, the exercise of some intellectual skills in some of its members, and some good attitudes in some of its members. Conversely, strongly corrupting conditions would be those which tend to discourage epistemic phronesis in many people, the exercise of many intellectual skills in many people, or discourage/encourage good/bad epistemic traits in many people.

Once again, we can only really get a grip on the axis of corruptees by considering the axis of corruptors. In this respect, Kidd identifies at least four such axes, some of which correspond to the modes of corruption noted above. First, there is the *scope* of influence that a given corruptor yields within a given socio-epistemic environment – that is, the range of corruptees

---

<sup>63</sup> Cassam (2022) adopts a similar strategy for explaining terrorism, distinguishing generalist from particularist explanations.

affected by the corrupting tendencies at play. Second, corruptors can be judged by their *stability*. The more stable the corrupting conditions, the likelier they are of encouraging or facilitating the range of epistemic qualities discussed above. Third, there is a question of a corruptor's *specificity*. Initially, Kidd distinguishes this axis by reference to the kind of epistemic vices encouraged or virtues eroded. With our expanded repertoire of qualities, however, we can supplement this distinction. So, *specifically* corrupting conditions not only consist of, say, a set of alethic vices or traits like dogmatism or close-mindedness, but include specific intellectual skills that are at risk of being eroded. Meanwhile, *generically* corrupting conditions are those which tend to facilitate a broader range of vices, traits, or failings, including those that straddle the 'intellectual-ethical' (Fricker, 2007: 120) or 'ethico-epistemic' divide (Daukas, 2018).<sup>64</sup>

Finally, we can identify a corruptor by reference to its *strength*. It is this last axis that I want to examine in closer detail. At an intuitive level, we might expect a stable corruptor with a wide scope to correspond to what we earlier referred to as 'strongly corrupting conditions'. Indeed, we have seen how Kidd suggests this by characterising such conditions as those which 'entrench many vices in many subjects' (2019). If corrupting conditions are unable to do this, or only in some subjects, the result is 'weakly corrupting conditions'. However, Kidd (2020: 74) later goes on to describe the strength of corrupting conditions by a different criterion: whether they are 'more capable of reliably corrupting subjects'. For reasons that will shortly emerge, I think neither of these definitions is fully adequate, which not only requires us to amend how we measure the strength of epistemic corruption, but also the more general axis of corruptors.

As we have seen, Kidd talks a lot about the kind of conditions we might find at work under epistemic corruption. Unfortunately, he does not say a great deal about those subjected to this phenomenon, and this leaves us with only half an understanding of how exactly agents might be corruptible. This seeps into his initial distinction between weakly and strongly corrupting conditions, which is my first concern. The concern is not so much that we cannot neatly carve up epistemically corrupting conditions in this way; rather, it is that Kidd's distinction leaves us no room to appreciate his observation that epistemic corruption is a 'dynamic set – or set of processes' (2019: 224). If this is true, we should expect the strength or weakness of corrupting conditions – and those within them – to be just as dynamic. This calls for the recognition of what I will call *moderately* corrupting conditions, which recognise

---

<sup>64</sup> Fricker's concept of *testimonial injustice* (2007) and Daukas' notion of *epistemic trustworthiness* (2006, 2018) would be two primary examples of such virtues and vices, insofar as they trade on both truth and justice. Roberts and West (2020) also detail the ethico-epistemic dimensions of honesty.

how some conditions might dynamically encourage/discourage many of the qualities above in only some agents or encourage/discourage only some of the qualities in many agents.

By recognising the existence of moderately corrupting conditions, we can better understand how agents might be susceptible to epistemic corruption. Take, for example, a social environment which encourages an array of epistemic vices like dogmatism, intellectual arrogance, and intellectual laziness; discourages epistemic phronesis or the exercise of many intellectual skills; or facilitates the loss/development of numerous good/bad epistemic traits. As it happens, though, only a handful of the people exposed to these conditions *do* go on to develop many of the vices or traits, to lack epistemic phronesis, or find the quality of their intellectual skills dwindles. In Kidd's words, it turns out that many under these conditions 'have sufficient awareness and strength of mind to resist the corrupting tendencies' (2019: 225).

Now imagine a case in which many people come to develop only a select epistemic quality. Consider Norman Malcolm's memoirs of Ludwig Wittgenstein (2001). In the winter of 1946-1947, Malcolm recounts how Wittgenstein began considering the prospect of resigning his chair at Cambridge, in part, because he saw himself as a 'dreadful teacher' (*ibid.*: 53). What brought Wittgenstein to this conclusion was a realisation that his teaching 'produced a tendency in his students to assume that precision and thoroughness were not required in their own thinking'. In turn, his students were prone to slavishly repeat his remarks, imitate his style of speaking, and produce 'slovenly philosophical work' (*ibid.*: 53). Although Wittgenstein prized independence of mind, the irony is that his teaching unintentionally instilled in his students the kind of thinking we might associate with vices like intellectual servility. At least on Malcolm's description, this was a case in which many subjects (Wittgenstein's students) came to develop the features reminiscent of a select vice (intellectual servility). In other words, Wittgenstein's influence was moderately corrupting.

If we draw a sharp distinction between strongly and weakly corrupting conditions, not only are we unable to appreciate the strength of character certain people might exhibit under epistemic corruption, but we cannot accommodate important, dynamic cases like the one above. Recognising the existence of moderately corrupting conditions allows us to make good on the dynamic nature of epistemic corruption. This is particularly important if the ontology of character I offered in chapter 1 is correct. For if epistemic virtues and vices manifest as intellectual *traits* of character, then these traits will emerge from prior attitudes. Those who manifest epistemic virtues and vices, then, will come to develop them in a gradual, patchwork fashion. Furthermore, the stronger the attitudes become, the likelier an agent is to develop the skills associated with the resulting trait. In sum, there would be nothing

incoherent about *some* people developing *many* of the qualities just noted. Just as equally, we should expect cases in which *many* people only ever develop *some* of these qualities. After all, nobody's character is the same.

My second concern picks up from the first. If there can be *moderately* corrupting conditions, this brings into question Kidd's claim that the strength of such conditions is a function of how 'reliably capable of corrupting subjects' they are. This is because there can now be corrupting conditions that are perhaps better described as 'moderate' in strength, but which are in fact reliably capable of corrupting those subjected to them. Indeed, the case of Wittgenstein above seems to suggest this possibility. One of the reasons why Wittgenstein believed his influence as a teacher was 'largely harmful' was because his domineering style led many of his students to become fawning or servile (Malcolm, 2001: 53).<sup>65</sup> Over time, this effect became so discomforting for Wittgenstein that it played a major role in his eventual resignation from Cambridge.

Whatever else characterised Wittgenstein's teaching, it seems fair to say that he was reliably capable of inducing a kind of intellectual servility in his students – his decision to stop teaching is a testament to his reliability at inadvertently creating these conditions. Yet it is equally true that he only encouraged this servility in a very narrow domain of their epistemic lives *viz.*, their academic conduct, as opposed to their more general intellectual conduct. But if we measure the strength of a corruptor by reference to whether they are 'reliably capable of corrupting', we would be forced to classify Wittgenstein as a strong corruptor. This, I think, is the wrong verdict for the reason that only a narrow domain of his students' intellectual characters was affected by his teaching despite the reliability with which this occurred.

What's more, if we premise the strength of a corruptor on their capability to 'reliably corrupt', it will force us to conclude that conditions or persons who are *unreliable* at doing so ought to be considered weakly corrupting. However, this conclusion would now cause us to overlook cases in which a corruptor incurs a significant characterological shift in another despite being largely unreliable. Imagine a revised story about Wittgenstein's teaching. Suppose his influence on students was sporadic, such that the majority of them did not fall prey to the kind of intellectual servility discussed above. As it happened, though, Wittgenstein found himself teaching a class in which there was one highly impressionable student. Suppose further that this impressionability was the result of a desire to fit into the intellectual life at

---

<sup>65</sup> Iris Murdoch was a notable outlier here. Wittgenstein's approach to philosophy was quite austere; he encouraged those who were not good at philosophy to give up philosophy and abandon their careers as such.

Cambridge. Whatever views were “trendy” at the time, the student sought to endorse; whatever the student perceived to be “intelligent”, they actively tried to emulate it. In the student’s eyes, Wittgenstein was the paragon of intelligence, and thus they mimicked what they saw in front of them. Unfortunately, this not only made the student increasingly intellectually servile, but this servility manifested in them defaulting to others’ views more generally, regardless of their epistemic merit.

Again, Wittgenstein’s corrupting tendencies seemed to be inadvertent and entirely unintentional. But unlike the previous case, the intellectual servility his teaching encouraged is not simply confined to this student’s academic life; rather, it comes to engulf and guide their epistemic affairs at large. Moreover, given that the majority of Wittgenstein’s students are not susceptible to his influence, we would be hard pressed to describe his corrupting tendencies as anything but reliable. Indeed, Wittgenstein’s failure is an indication of his relative *unreliability* as a corruptor. Yet this sort of case now strikes me as an instance of strongly corrupting conditions, precisely because of how deep-seated the characterological effects on this student are.

Not only do cases like this reinforce the point I made earlier about the strength of corrupting conditions only being as strong as those subjected to them; they further reveal that the strength of epistemically corrupting conditions can come apart from their reliability. Just because a corruptor is *reliable*, we should not automatically infer from this that it is *strong*. The two are independent variables. In light of this, I suggest that we distinguish a fifth axis of corruptors: *reliability*. In light of this axis, we can add a sixth mode of corruption to this list above, one which captures the reliable capability of different corrupting conditions.

- (IV) *Authentication*: a corruptor can authenticate an epistemic vice(s), poor attitudes, or vice-indexed skills by increasing the *reliability* with which they manifest or are exercised across differing socio-epistemic contexts.

Corrupting conditions are reliable insofar as they consistently damage an agent’s intellectual character regardless of how deep-seated the damage is. Of course, it goes without saying that the erosive effects of a reliable corruptor are likely to be cumulative, such that, over a prolonged time, the relevant corruptee(s) becomes more susceptible to developing epistemic vices or losing any epistemic virtues, phronesis, or intellectual skills they possess. Still, the best way of proceeding in these matters is by considering the nature of the corruptees in question and the wider socio-epistemic contexts in which they operate.

The idea that we could measure a corruptor by reference to its reliability brings my discussion very close to suggesting that epistemically corrupting conditions involve a ‘success

condition', which is something that Kidd explicitly excludes from his account (2019: 225). I take it that a success condition would amount to saying that corrupting conditions *must*, in some sense, damage the intellectual characters of those subjected to them. Without this, the conditions are not corrupting. As Kidd reminds us, though, a system or set of conditions can be corrupting 'even if, as a matter of fact, it fails to corrupt some of its members' (*ibid.*: 225).

Is a reliability axis just a success condition in disguise? I do not think so. As I am thinking of it, reliability serves to indicate the counterfactual robustness of a given set of epistemically corrupting conditions, which allows us to then accommodate the idea of there being such conditions even if, on occasion, they fail to affect anyone. To make this point clearer, take the example of damp gunpowder. Although damp, this feature does not detract from the overall reliability of gunpowder as an explosive – were the conditions drier, it would reliably detonate if ignited. Here, gunpowder has a disposition to explode, which it reliably does.<sup>66</sup> What the dampness decreases, however, is the *success* with which the gunpowder explodes.

A similar point extends to epistemic corruption. Much like the damp gunpowder, it could be that the epistemically corrupting effects of Wittgenstein's teaching fail to "ignite". To use the relevant terminology, we might say that he finds himself in conditions that 'mask' his reliability – in this case, teaching students whose characters are less susceptible to his influence or who have powerful resistant tendencies. But it does not follow from this that he thereby loses his reliability *qua* corruptor, just that he is in unsuitable conditions to bring this about. We can repackage this point in modal terms by saying that there remains a *close possible world* in which Wittgenstein's influence does adversely affect the epistemic conduct of his students. Imagine a world in which Wittgenstein remained at Cambridge teaching because he encouraged his students to be intellectually autonomous. While a nice prospect, such a world bears little resemblance to the actual world in which Wittgenstein found himself. Since this modal world is far removed from his actual world, he would still remain a reliable corruptor, even if there were occasions on which his influence failed to affect any students.

Notice that this way of characterising the reliability axis allows us to nicely accommodate Kidd's contention that a system would still count as epistemically corrupting even if it failed to affect anyone. What matters when determining the reliability of a corruptor is the similarity of the possible worlds in which they fail to corrupt, relative to the actual world in which they exist. The further the modal distance, the less reliable the corruptor is.

---

<sup>66</sup> For discussion of 'dispositional properties', see Gilbert Ryle's (1949) *The Concept of Mind*. This example is inspired by David Armstrong's (1980) example of glass' disposition to shatter when dropped.

The closer the possible worlds, the more reliable they are. Not only is a reliability axis importantly different from a success condition, then, but its inclusion makes any account of epistemic corruption better equipped to explain why conditions that, on occasion, fail to corrupt anyone might still merit the label corrupting.

### 3. Passive Epistemic Corruption

We can glean several things from the discussion so far. First, epistemic corruption is a dynamic set of processes that takes many forms and can damage a person's character from multiple angles. Second, the modes under which this can occur – and which agents are affected – are equally varied, operating along a range of axes that we can attempt to track. Together, these features have provided us with a fairly robust understanding of how epistemic corruption can operate. Nevertheless, the picture is still lacking an important feature that has lurked in the background of the discussion. This is the subtle distinction Kidd draws between the damage inflicted on a person's character by eroding epistemic virtues and the damage caused by encouraging them to develop vices. These, respectively, refer to what he calls *passive* and *active* epistemic corruption. Let us start with passive corruption.

Many instances of epistemic corruption I have considered in this chapter so far have been passively corrupting: they have been conditions which 'fail to adequately facilitate or properly encourage the exercise of virtues' (Kidd, 2019: 229), that lead to the 'erosion of conditions that encourage the cultivation and exercise of epistemic virtues' (*ibid.*: 231), or which involve the 'failure of the subject to develop an epistemic character characterised by virtues and integrity' (Kidd, 2020: 71). We can be more concrete about these kinds of conditions by examining David E. Cooper's (2008) critique of certain educational practices.

Cooper takes issue with how certain kinds of education can, as he puts it, 'marginalise, demote, or impugn' teachers' ability to exercise professional virtues like truthfulness. By 'truthfulness', Cooper has in mind the dispositions that comprise what Bernard Williams (2002: 44) refers to as the 'virtues of truth'.<sup>67</sup> These include the virtue of 'accuracy' – a disposition to ensure that one's beliefs are true – and the virtue of 'sincerity' – a disposition to express 'what one actually believes' in spite of background incentives to do otherwise (*ibid.*: 87, 96). For Cooper (2008: 80), schooling can erode these virtues in several ways. First,

---

<sup>67</sup> In addition to accuracy and sincerity, Cooper (2008) adds fidelity and transparency.

truthfulness cannot be measured in ways that grades, pass rates, and employability can. Unlike quantifiable criteria, truthfulness is not something that is measured via assessment. Second, schooling can erode truthfulness if it fails to help service an array of economic and social interests (*ibid*: 80). If making students more employable is the aim of education, teachers will inevitably become encouraged to be 'economical' with the truth and honesty. As Cooper puts it:

Efficiency might require, for example, that students, rather like soldiers going into battle, are instilled less than honestly with over-confidence in their prowess, their chances of success, and the satisfaction that success might bring (2008: 80).

An efficient educational system renders teachers vulnerable to insincerity about their students' abilities, which in turn means that students acquire a less than accurate sense of themselves. Finally, truthfulness can wane if teachers are required to pretend that a certain minority view is equal in credibility with an established view, even when, in their judgment, it is not. To borrow Cooper's example, an educational system might avoid discussions of evolution in biology classes so as to not upset Fundamentalist Christians, say (*ibid*: 80). By remaining neutral on this topic, such an educational system not only does its students a disservice by preventing them from learning about this aspect of science, but it creates conditions that frustrate accuracy and sincerity for teachers if asked about the origins of life.

The kind of conditions that Cooper worries about are certainly captured by Kidd's notion of passive epistemic corruption. However, I want to address two general issues. The first of these is relatively minor and it is that Kidd's focus is exclusively on what I have been calling the Aretaic conception of epistemic corruption. This explains why the emphasis on passively corrupting conditions are those which erode epistemic *virtues*. As I noted above, though, there is far more to a person's intellectual conduct than their epistemic virtues, and so a more inclusive view of what counts as passively corrupting should reflect this reality. Accordingly, we can expand passive corruption to include the deterioration of an agent's cognitive faculties, the erosion of their intellectual skills, their capacity for epistemic phronesis, and their good attitudes towards epistemic goods and practices. This way, passive corruption has a far greater reach in its ameliorative potential.

My second reservation about passive corruption is more substantial. While I agree that there exists what Kidd calls 'passive corruption', I am less confident that we can straightforwardly identify such conditions if we endorse his characterisations. In essence, passive corruption amounts to either a failure to facilitate the development of various epistemic qualities or the erosion of conditions that encourage these qualities. While there

might not appear to be much hanging on these two interpretations, the former leads to a rather extreme conclusion, if we take it at face value. This is because it would render almost all of our social environments “passively corrupting”, given that they frequently fail to facilitate the development of epistemic virtues, intellectual skills, traits and so on. Imagine, for example, that I am a member of a local cricket club. Amongst other things, the club encourages teamwork and friendship, but it does not cater to intellectual curiosity, flexibility and so on. Does this mean that we should classify it as passively corrupting? If we strictly interpret passive corruption as Kidd intends, not only would we be forced to make such a concession, but we would have to classify many of our social environments as passively corrupting on account of their failure to facilitate the development of epistemic virtues, etc.

This strikes me as incorrect, and I suspect Kidd would agree by emphasising that a cricket club is not the sort of social environment that we *should* expect to *properly* (or at all) facilitate epistemic virtues and so on (2019: 231).<sup>68</sup> As such, it is not the kind of environment that we should label passively corrupting. Of course, this raises the following question: *how* should we know which environments to expect to properly facilitate the development of epistemic virtues? Perhaps this is an empirical question, but a sensible strategy might be considering how central epistemic values and practices like truth, understanding, inquiry and education are to the environment in question. Such a strategy would suggest that we index passive corruption to environments in which the development and exercise of epistemic virtues, intellectual skills, phronesis, and good epistemic attitudes are explicitly called for. Where they are of central importance, we should not hesitate to characterise conditions that erode them as passively corrupting.<sup>69</sup>

Nevertheless, this strategy also seems to run into a problem. Consider the qualities that might reasonably be encouraged on a journalism course, which seems to be a putative example of a social environment that trades on epistemic values. Among the virtues encouraged, we might expect things like accuracy, inquisitiveness, intellectual courage, etc. However, it might be the case that these virtues come into conflict with other, recognisable virtues like intellectual creativity or humility. Unless one subscribes to a particularly strong version of the “Unity Thesis” – the idea that possessing one virtue entails possession of all the virtues – the upshot is that the journalism course heavily encourages the former virtues

---

<sup>68</sup> The caveat here would be a board meeting *within* the cricket club, since this is likely to involve deliberations between members. However, this simply emphasises the importance of specifying corrupting conditions.

<sup>69</sup> In fact, this proposal seems to have been implicit in much of the existing applications of epistemic corruption, which have mostly focused on educational settings (Battaly, 2013; Forstenzer, 2018; Kidd, 2019; Kidd et. al., 2021; and Monypenny, 2021).

over the latter.<sup>70</sup> Insofar as the course fails to place emphasis on the latter virtues, though, it risks eroding conditions that would encourage their development. Understood so, the course would count as passively corrupting.

Again, though, I would refrain from labelling it ‘passively corrupting’. While the course erodes the potential for students to cultivate intellectual humility or creativity, the reality is that certain occasions of our epistemic lives call for different qualities. This might mean bringing a particular set of intellectual character traits, virtues, or skills to the table on one occasion, whilst refraining from exercising a different set on another. Even if this does erode our potential to develop and manifest a certain cluster of traits, virtues, or skills in one instance, it does not follow from this that another cluster is failing to flourish under those same conditions; they are just more suited to the environment in question. If we understand passive corruption in terms of an erosion of conditions that facilitate the development of epistemic virtues, we not only risk treating “epistemic virtue” as a single category that can be developed and exercised without tension; we risk overlooking this important feature of our epistemic lives. Therefore, we are no closer to having a general procedure for identifying passively corrupting conditions.

This leaves us understanding passive corruption in terms of a ‘failure of the subject to develop an epistemic character characterised by virtues and integrity’ (Kidd 2020: 71). Unfortunately, if the ontology of character I proposed in chapter 1 is correct, then this casts doubt over the prospects of us developing such epistemic characters. On the view I sketched, recall, our characters emerge and develop in a piece-meal fashion from the summary attitudes we form. As these congeal, they form the mental-state dispositions of behaviour we associate with character traits. As I noted above, though, it will often be a matter of constitutive luck which resources we have at our disposal, which in turn can dictate the kind of attitudes we form. This can lead us to form different, sometimes conflicting mental-state dispositions. If I am raised in a family that places an emphasis on puzzles and games, I might come to develop positive attitudes towards solving them. But if I dislike getting my hands dirty, my attitude towards fixing my car will be negative. On my view, it would not be inconsistent if such a person found the characteristic perseverance they manifest when solving puzzles lacking when it comes to all things mechanical. Indeed, it would be perfectly consistent to find this person exhibiting the dispositions we associate with intellectual laziness. While people are certainly capable of developing a character ‘characterised by integrity and virtue’ in theory,

---

<sup>70</sup> Some contemporary theorists attempt to defend a version of this thesis, including Toner (2014) and Vaccarezza (2017). See Badhwar (1996) for an overview.

in practice this will be much harder.

How should we proceed? One approach suggested by the discussion above would be to concede that virtually all of our social environments are passively corrupting. Even if this is an extreme conclusion to draw about our susceptibility to epistemic corruption, it would certainly be no less valid simply on account of this. After all, if there are at least three variants of epistemic corruption, why should we expect ourselves to be overwhelmingly immune from their effects? If passively corrupting conditions are as widespread as this suggests, perhaps this just reflects the messy reality of our socio-epistemic lives. However, there is a way of coming to terms with this prospect. To see this, consider the following analogy. Think about the number of cracks there are on the roads we drive down. Passively corrupting conditions are like these cracks; they might populate our roads, but we can afford to largely ignore them. What we really worry about are the *potholes* that damage our cars when driven over. The environments which severely undermine the development of epistemic virtues, intellectual skills, epistemic phronesis, and good epistemic attitudes are the potholes that damage our intellectual characters. These are the kind of passively corrupting conditions we should be concerned about.

Across his work on epistemic corruption, Kidd helpfully identifies a number of features we are likely to find under these sorts of passively corrupting conditions. However, he does not group them in a way that makes this explicit. Therefore, we can make a taxonomy of passively corrupting conditions. These will variably include:

(I) *An absence of virtuous exemplars*—under passive corruption, those with a track-record of exhibiting behaviour we would describe as epistemically virtuous will either be removed or hidden from sight, so that there are fewer opportunities for others to learn from them.

(II) *Denigrating virtuous behaviour*—passively corrupting conditions might consist of people chastising fair-mindedness or intellectual charity. For example, fair-mindedness might be spun as weak-mindedness and intellectual charity as an over-sensitivity to others' arguments, claims and so on. Those who manifest epistemic virtues under these conditions are asking for trouble.<sup>71</sup>

---

<sup>71</sup> Although epistemic virtues might be sufficiently robust to withstand criticism, this might not be the case with a person's good attitudes towards the truth or inquiry, which seem more vulnerable to exploitation. If so, it would suggest that denigrating features could operate with varying success depending on the quality involved.

(III) *Establishing conditions that increase the exercise costs of epistemic virtues* – although there might be evidence of institutional malpractice, that institution might have policies which make investigating this matter come at a high cost to one’s livelihood. If the cost of exercising intellectual courage becomes too high, for example, it is only natural for one to refrain from doing so. In turn, the conditions suppress the development and exercise of this epistemic virtue.

A limitation of Kidd’s account, we have seen, is its vice-centric approach to epistemic corruption, and this is clearly present in these three features. If I am right that there are at least two other variants of epistemic corruption, then we need to add the following options:

(IV) *Denigrating highly skilled epistemic conduct* – just as passively corrupting conditions might involve chastising people for exhibiting epistemic virtues, so too might they involve undermining people who are skilled at asking good questions, making incisive challenges to authority, and so on.

(V) *Frustrating the development and exercise of intellectual skills* – a company might implement policies or practices that make it harder for people to develop, refine, or exercise an array of intellectual skills. Alternatively, it might remove programmes that educate for these skills.

(VI) *Discouraging the development and exercise of epistemic phronesis* – an institution might remove ways for its employees to develop the good judgement as to when and where they should exercise any epistemic virtues they possess. Upon discovering that an employee is expressing concerns about the impact of a product, for example, the manager might question why they are concerned or insist that work is the wrong place to express this sentiment.

(VII) *Tolerating the loss or violation of epistemic and/or ‘zetetic’ norms* – suppose there is a gradual acceptance of intentionally false assertions with a company, a failure to base beliefs and statements on the truth, or the introduction of irresponsible processes of gathering and responding to evidence (Friedman, 2020; Flores and Woodard, 2022).

Many of these features will operate in tandem with others. For example, under conditions in which there is a clear absence of virtuous exemplars, it would not be surprising to find that the exercise cost of virtue is much higher, particularly since there would be no inspiration for people to overcome the difficulties that make exercising epistemic virtues costly. Moreover, if there is a close connection between the exercise of epistemic virtues and intellectual skills,

then conditions which denigrate the former are also likely to denigrate the latter. Relatedly, conditions which denigrate a person's epistemic virtues and intellectual skills are more likely to succeed at discouraging epistemic phronesis. When combined, then, the features present under passively corrupting conditions are likely to be more potent than the individual sum of their parts. Regardless of which features we do find under these conditions, though, it will be harder for agents to develop epistemic virtues, phronesis, good intellectual traits, and intellectual skills. In sum, the taxonomy above provides us with a richer picture of the mechanics of passive epistemic corruption.

## 4. Active Epistemic Corruption

Things are slightly more intuitive when it comes to the idea of active corruption. As Kidd describes it, active corruption takes place against the backdrop of 'experiences or activities that promote the development and exercise of epistemic vices' (2018: 48), environments that 'promote, fuel, or reward the exercise of vices', those in which there is a 'deterioration of any pre-existing virtues and integrity already present in the subject's character' (2020: 71); and it occurs when 'one or more epistemic vices are acquired, intensified or entrenched as a result of the agent operating within or being subjected to environments that facilitate their development and exercise (2021: 155). Once again, an example will help illustrate this form of epistemic corruption.

Heather Battaly (2013) has drawn attention to the problematic nature of the Seven Stars Solutions policy that operated in Texan higher education for many years. The problem she identifies is the policy's potential to encourage a vice she calls 'epistemic insensibility' in academics – a failure to 'desire, consume, engage in, or enjoy appropriate epistemic objects; on appropriate occasions; at appropriately frequent intervals' (2013: 268). Battaly claims the policy encourages this vice because of how it narrowly defines epistemic goods in terms of their instrumental value. A degree or research area is seen as epistemically valuable *iff* it contributes to social or economic concerns. Those which do not are taken to be less valuable and potentially devoid of any epistemic value. If those operating under this policy came to develop this kind of epistemically insensible view towards their research, then it would be an example of active corruption at work.

The encouragement of epistemic vices is surely a feature of actively corrupting conditions. However, such conditions will not just trade on epistemic vices. In addition, they

will likely involve (i) the explicit development of poor attitudes towards epistemic goods and practices, which lead to the development of bad intellectual traits; (ii) conditions that entrench intellectual failings like bad questioning or poor reasoning; (iii) and encourage the development of what I will refer to in chapter 5 as *vice-indexed skills* such as being a skilled liar, skilfully obfuscating evidence, and being a skilled bullshitter. With these qualities in our arsenal, we can better identify a range of corruptors and corruptees we are likely to find under active epistemic corruption.

Much like its passive counterpart, we must broaden the remit of active corruption if it is to operate as part of an effective ameliorative strategy. That said, we should be careful how wide we cast the net of active corruption. This is particularly the case with regards to Kidd's contention that active corruption occurs when there is a 'deterioration of any pre-existing virtues and integrity already present in a subject's character' (2020: 71). Here, we can extend two of my earlier concerns about passive corruption to this characterisation. Recall my example of the journalism course from above. There, the course placed great emphasis on cultivating a particular set of epistemic virtues and intellectual skills. Although it does not follow from this that any pre-existing virtues and skills the agents possess will deteriorate, it does plausibly weaken other virtues and skills that are perhaps not as salient or required. Over time, we might reasonably expect the strength of these qualities to decrease, if not deteriorate.

Despite actively encouraging a cluster of epistemic virtues and intellectual skills, the journalism course once again creates conditions that risk deteriorating some of an agent's other pre-existing virtues. In light of this, the course would be guilty of active corruption on Kidd's characterisation. Again, though, this diagnosis seems to run counter to the plausible idea that our virtues can come into conflict. Jason Baehr (2011: 63) recognises this when contrasting the reliability of faculty virtues with responsibility character virtues. Unlike the former, which I noted are often indexed to particular environments, he notes that 'the applicability of a character virtue to a particular proposition or field of propositions usually depends in *a very deep way on contingent features of the person or situation* in question (*ibid*: 63, my italics). The upshot is that certain clusters of virtues tend be reliable in one situation, yet unreliable in another. As such, we are often forced to make trade-offs regarding the qualities best suited for an occasion. As he observes:

The virtues of intellectual caution and carefulness, for instance, might be required in one situation to reach the truth about a proposition which in another situation could be known only via an exercise of intellectual courage and perseverance (*ibid*: 63).

If somebody wants to reliably uncover information or break stories as a journalist, the former virtues are unlikely to be as effective in this endeavour as the latter. Instead, being intellectually courageous and persevering seem to be the kind of qualities that would reliably lead one to the truth in these matters. It would be no surprise, therefore, if the course encouraged intellectual courage and perseverance at the expense of intellectual caution and carefulness. If these qualities were present in a particular student's intellectual character, the lack of emphasis or encouragement for these virtues is likely to result in their strength or effectiveness waning. This is not to suggest, of course, that these virtues will disappear entirely; rather, it is that their quality will degrade over time if they go unexercised. However, this just speaks to a more general point about the realities of life. I think few, if any of us, could ensure that the character traits, skills, and capacities we have developed in our lives remain at the same quality throughout. Like most things, their quality largely depends on how frequent and infrequent we exercise them. If so, this characterisation of active corruption not only fails to honour the realities of life, but also generates the rather ironic conclusion that environments which nevertheless encourage certain virtues could warrant the label of actively corrupting.

Now, even if it is granted that this conclusion is ironic, I suspect Kidd would simply consider it to be a more accurate depiction of the fragility of our intellectual characters. Just because a given environment encourages a cluster of epistemic virtues, Kidd might say, this should not rule it out from still deteriorating other pre-existing virtues. While logically possible, this kind of response poses its own problem for Kidd. For if certain environments can encourage epistemic virtues whilst simultaneously deteriorating others, then in what sense can a person possibly cultivate a character with any 'pre-existing integrity'? Perhaps this is possible with regards to particular domains of our intellectual conduct. Still, it would remain unclear whether this would even amount to anything like "integrity" of character in any meaningful sense, particularly if the virtues required in one domain risk deteriorating those in another. So, it looks as though Kidd cannot embrace the kind of actively corrupting conditions just noted, whilst at the same time characterising this form of epistemic corruption in terms of conditions which deteriorate any pre-existing virtues and integrity present in an agent.<sup>72</sup>

For these reasons, I recommend that we retain Kidd's initial characterisation of active

---

<sup>72</sup> There are two additional worries worth noting on this point. First, by characterising the kind of environments discussed as actively corrupting, Kidd risks placing all cases of active corruption on a par, which could undermine the severity of those environments that are intuitively more deserving of the label "actively corrupting" – especially those in which epistemic vices are involved. Second, it would seem to make epistemic corruption an inescapable phenomenon we could never fully ameliorate.

corruption in terms of conditions which facilitate the development and exercise of bad epistemic qualities. To put some more meat on the bone of this suggestion, let me offer a similar taxonomy of the features we might expect under actively corrupting conditions. Again, Kidd (2020: 76-77) identifies a number of initial features but they are not explicitly associated with active corruption. Thus, they variably include:

(I) *Valorising vicious conduct and exemplars* – a corporation might place a high value on its CEO’s ardent arrogance, holding them up as a paragon of corporate success to be emulated by other employees.

(II) *Rebranding vices as virtues* – in valorising the CEO’s arrogance, the corporation could attempt to be disguising it as confidence as a means to boost productivity or efficiency, rather than have its employees spend time deliberating about important decisions.

(III) *Rewarding the exercise of vices* – a tobacco company might reward its scientists who dishonestly publish in journals that discredit links between smoking and cancer (Proctor, 2011). An oil company could introduce financial incentives for its employees to undermine evidence concerning anthropogenic climate change (Oreskes and Conway, 2012).

(IV) *Establishing policies which require developing and manifesting vices* – to develop their careers under the Seven Stars Solutions policy, recall, academics had to insensibly abandon projects that lacked “value” in favour of those with economic and social outputs (Battaly, 2013).

Once again, we can supplement these features with at least a further four:

(V) *Valorising poor attitudes towards epistemic goods and practices.*

(VI) *Disguising poor attitudes as good ones* – rebranding close-minded attitudes as intellectually courageous ones.

(VII) *Establishing policies that encourage vice-indexed skills* – by incentivising intellectual dishonesty, an oil company might provide fertile conditions for agents to refine and practice their skills of lying or obfuscating evidence.

(VIII) *Creating institutional/organisational cultures that encourage a flagrant disregard for epistemic/zetetic norms* – an institution might refuse to open inquiries into malpractice

when they ought to, work cultures might chastise the double-checking of work, or introduce time-constraints that encourage poor handling of evidence, etc.<sup>73</sup>

Again, it goes without saying that these features will largely overlap. The oil giant which rewards dishonesty is going to encourage its employees to develop their lying or obfuscatory skills. Where there is a flagrant disregard for epistemic norms, conditions could encourage people to develop poor attitudes towards the truth or practices like inquiry. Another way of valorising epistemic vices is by rebranding them as epistemic virtues, and so on.

Though I have presented two distinct taxonomies of passive and active corruption, this is largely for stylistic purposes and not because the two are mutually exclusive. From the discussion, I hope it is clear just how intertwined the two can be. It is not difficult to imagine that the kind of conditions which gave rise to the actively corrupting tendencies of the Seven Stars Solutions policy were brought about by some of the passively corrupting features we have discussed: educational authorities denigrating research agendas that embed intellectually creative approaches to abstract questions; failing to reward or incentivise intellectual perseverance; and perhaps most obvious of all, eroding the virtue of *epistemic temperance* – a desire to ‘consume, engage in, and enjoy appropriate epistemic objects, on appropriate occasions, in appropriate amounts, and frequencies’ (Battaly, 2013: 267).

However, it is just as possible for active corruption to give rise to passively corrupting conditions. For instance, the more entrenched the Seven Stars Solutions policy became, the harder it plausibly became for academics to cultivate certain positive intellectual traits like intellectual creativity, perseverance, and epistemic temperance. The policy likely frustrated peoples’ ability to refine and develop intellectual skills relevant to those areas of academia that were seen to not contribute to social and economic concerns – think of a logician’s inability to practice first-order logic – and the policy was likely to warp epistemic norms around the kind of truths people ought to have valued (and devalued) in academia. In this way, the establishment of actively corrupting conditions loops back to sustain and foster passively corrupting conditions, just as much as passively corrupting conditions give rise to actively corrupting ones. Unsurprisingly, then, the two kinds of epistemic corruption will often be inseparable.

---

<sup>73</sup> There is also the possibility of there being meta-failings, where people are not honest with themselves about their limitations and failings. Somebody who considers themselves intellectually courageous, for instance, might be in denial about the extent of their closed-mindedness, and so opt to characterise themselves as courageous over closed-minded. I suspect this dimension contributes to making certain epistemic vices ‘stealthy’ (Cassam, 2019a).

## 5. Conclusion

In this chapter, I set out to provide the reader with a substantive understanding of what I have called the *mechanics* of epistemic corruption. The first of these was the idea of epistemically corrupting conditions, which I unpacked in greater detail. I then examined the relationship between those subjected to epistemically corrupting conditions – the corruptees – and those who do the epistemic corruption – the corruptors. Finally, I critically discussed Kidd’s distinction between passive and active epistemic corruption.

Though I took my lead from Kidd’s existing work on the phenomenon, this chapter cast a critical eye over many of his claims and contentions regarding how we should theorise epistemic corruption. For starters, I argued that we should not characterise the strength of corrupting conditions by reference to their reliability. Indeed, I noted how these two parameters can, and often do, come apart. By treating the strength of a corruptor as synonymous with its reliability, an account of epistemic corruption cannot be dynamic in the way Kidd envisages. Accordingly, I proposed and defended the inclusion of a reliability axis to ensure that my account overcomes the difficulties faced by Kidd’s. Moreover, I offered an improved way of thinking about passive and active corruption, a theme that arises in the next chapter. What this chapter ultimately sought to do, then, was produce a more dynamic and rounded view of a concept with which we can examine our socio-epistemic environments.

However, there is a further question that remains unanswered: if there can be epistemically corrupting conditions, how do these conditions emerge in the first place? It is to this question that we now turn.



## 4.

# Epistemic Corruption and Collective Agents

## Introduction

In the previous two chapters, I introduced the concept at the heart of this thesis: *epistemic corruption*. My aim there was to scrutinise the existing literature on the phenomenon, focusing on Kidd's original account. Although his insights provide a useful *vice-theoretic* approach to understanding how our social environments can degrade the quality of our intellectual characters, I argued that they have their limitations. This led me to amend the scope of epistemic corruption, in an effort to fully realise the potential of the corruption concept as an ameliorative resource in vice epistemology. I hope, therefore, that the account I have begun developing can live up to this task.

Towards the end of the last chapter, I left an important question unanswered. The keen-eyed reader may have noticed that many epistemically corrupting conditions tend to manifest within an institutional or collective context. Amongst those that Kidd (2019: 225, 2022) highlights are 'punishing schedules', 'chilly climates', 'hostile cultures' and 'cultures of speed'. Indeed, as I observed, much of the theoretical and empirical work on epistemic corruption has focused on educational institutional environments. The more general point is that the phenomenon of epistemic corruption seems to partially depend on there being institutions, groups, and structures that instantiate the sort of corrupting conditions to which individuals can be subjected. If this is the case, though, how is it that these institutions or

collectives come to be hotbeds for epistemic corruption in the first place?<sup>74</sup>

In this chapter, I address this question by attempting to provide an account of how collective agents might, themselves, be susceptible to epistemic corruption. So, just as we previously saw how our socio-epistemic environments can degrade the intellectual character of individual agents, I now want to consider the extent to which *individuals* can degrade the quality of the institutions and collectives to which they belong.<sup>75</sup> In doing so, the hope is to provide a symbiotic picture of the relationship between the epistemic corruption of individual and collective agents.

I start in section 1 by motivating the claim that collectives and institutions can be adequate *corruptees* – that is, objects relevant to a corruptionist critique. To flesh this out, I draw on Miranda Fricker’s influential account of ‘institutional ethos’ (2020). By appealing to Fricker’s concept, section 2 develops two ways in a collective or institution can be subjected to epistemic corruption and acquire what I call a *corrupted epistemic ethos*. While Fricker’s account is a useful starting point, I argue in section 3 that her analogy between individual and collective character only goes so far. Specifically, I draw attention to two limitations that cast doubt on the analogy. Insofar as the analogy breaks down, I claim, we should not expect a neat account of collective character to fall out of individual character. Nor, by extension, should we expect the epistemic deterioration of collective agents to be fully captured by an account of epistemic corruption. This leads me to conclude, in section 4, with a methodological point: we should be cautious about relying too closely on analogies when doing vice epistemology.

## 1. Institutional Ethos

Epistemic corruption trades on the notion of intellectual *character*. In chapter 2, I distinguished three ways of specifying this general claim, distinguished by the features of the individual agent that is being corrupted. The Aretaic Conception emphasises an agent’s epistemic virtues and vices. The Phronetic Conception attends to their capacity to exercise epistemic *phronesis* and intellectual skills. The Attitudinal Conception on an agent’s epistemic

---

<sup>74</sup> In this chapter, I try to avoid using the terms ‘collective’ and ‘institution’ synonymously, since most institutions are collectives but not all collectives are institutions. For instance, a school might count as an educational institution, but its English Department is not an institution in itself, but rather a collective within the institution, namely a department or group.

<sup>75</sup> As this suggests, the sort of collective agents with which I am concerned here are ‘jointly intentional agents’ (List and Pettit, 2011) as opposed to ‘mere populations’ (Lahroodi, 2007).

traits. However, all of these accounts confine the scope of epistemic corruption to individual agents. This is a problem, insofar as (a) there are collective epistemic agents like universities, research bodies, etc., and (b) we often speak, in everyday language, of collective epistemic agents undergoing negative changes consistent with the concept of epistemic corruption: that universities are becoming less open-minded and are more hostile to certain ideas (Haidt and Haslam, 2016), or how London’s Metropolitan Police Service displays ‘too much hubris and too little humility’ in its day-to-day functioning (Casey, 2023: 12).

In this section, I start by motivating the idea that collective and institutional agents can possess the sort of qualities relevant to a corruptionist critique, that is, epistemic vices, traits etc. (section 1.1). If collectives can possess these qualities, can they also possess something akin to character? In section 1.2, I turn my attention to this question by examining Miranda Fricker’s (2020) recent attempt at modelling the institutional analogue of individual character, what she terms *institutional ethos*.

## 1.1 Collective Vice Epistemology

We are often quick to point out the epistemic deficiencies of the institutions and collectives that are part and parcel of our social environments. For those familiar with academia, it is all too common to hear people bemoan the arrogant treatment of staff by universities, the narrow-mindedness that has engulfed particular funding bodies, or lament the departments that have lost their characteristic rigour and creativity.<sup>76</sup> How should we respond to this kind of language? One option is to say that the statements are purely rhetorical or reflect what Kidd (2021) terms ‘merely rhetorical vice charges – they are a means to let off steam or communicate one’s discontent. Another, though, is to see this kind of discourse as reducing to *individuals*. What people mean is that the members of the university, funding body, or department exhibit (or fail to exhibit) the relevant qualities. On this way of seeing things, the features attributed to the collective are simply reducible to the aggregate features of the individuals who comprise the collective. In effect, these responses presuppose a *summativist* approach to collective agency (Quinton, 1976; Lackey, 2016, Cordell, 2017).<sup>77</sup> According to

---

<sup>76</sup> For various discussions that invoke this kind of language in regard to educational contexts, see Battaly (2013), Berg and Seeber (2016), Kidd (2023a), and Forstenzer (2018). A classic account is Allen Bloom’s, *The Closing of the American Mind* (1988).

<sup>77</sup> The term ‘summative’ is from Quinton (1976). Fricker (2010, 2020), Gilbert (1989, 2000), and Lackey (2014, 2016) adopt the term *summativism*, Lahroodi (2007) uses *individualism*, whereas List and Pettit (2011) prefer *eliminativism*. While subtly distinct, each trades on the more general idea above.

summativism, a collective  $C$  exhibits a given quality  $Q$  only if a sufficient number of members exhibit  $Q$ , where this can be understood in terms of most or all of the members exhibiting  $Q$ . Thus, if the funding body consists of fifty members, and twenty-six of them exhibit narrow-mindedness, then this suffices for the funding body to warrant the label ‘narrow-minded’.

Summativism about collective agency has its advantages. First, it coheres well with *methodological individualism*, the claim introduced by Max Weber (1922: 13) and subsequently taken up in the philosophy of social science, that social phenomena must be explained by showing how they arise from individual actions (Hollis, 2002).<sup>78</sup> Second, it seems to capture a good range of our evaluative reactions to collective bodies. For instance, suppose I hold a group of my students *blameworthy* for failing to submit their assignments by the deadline. In making this attribution, my intention is to level blame at the *individual* students who comprise the group. This is made clearer by reflecting on the fact that, had the group submitted their assignments on time, I would not have labelled it blameworthy in the first place. Third, and most importantly, summativism seems to offer us an intuitive way of thinking about the epistemic corruption of collective agents. It seems reasonable to think that if a group or collective has come to exhibit an epistemic vice, this must be because the individual members of that collective also possess the vice.

For example, suppose there is a group of scientists who are developing a new pharmaceutical drug. At  $T^1$  the group is comprised of intellectually diligent scientists who repeat their experiments and check their results. Summativism tells us that this group is intellectually diligent *qua* collective. For various reasons, however, several of these scientists leave at  $T^2$  and are replaced by intellectually sloppy scientists. At  $T^3$  the group is now mostly comprised of intellectually sloppy scientists, who fail to check their results and repeat experiments. The presence of these sloppy scientists not only helps to deteriorate the group’s prior diligence, but it infects its epistemic conduct by making it intellectually sloppy. This is captured by summativism. Since the group is comprised mostly of intellectually sloppy scientists, summativism correctly tells us that the group is intellectually sloppy *qua* collective. A once virtuous collective agent has been epistemically corrupted by its members and become a vicious one. So far, so good for summativism (Lahroodi, 2008, 2018).

For all its merits, though, there are some now-standard compelling objections to these kinds of summativism. Indeed, I do not think summativism can fully do justice to my aims in this chapter. First, there seem to be cases in which a collective is comprised of individually

---

<sup>78</sup> It needn’t follow, of course, that because social phenomena arise from individual actions, they must therefore be reducible to those individuals.

vicious agents, yet the collective does not possess this vice. These sorts of cases are likely to occur under what Jeroen de Ridder (2022: 372-373) calls the ‘interaction model’ of collective agency, a form of which involves individuals interacting with a collective’s formal procedures and structures.<sup>79</sup> Miranda Fricker (2010: 236-239) offers an example like this where a debating society is comprised of individually prejudiced members, but due to their interactions with various internal mechanisms, each members’ prejudice is cancelled out in proceedings. As a result, the *group* ends up displaying impartiality in its debates.

This kind of case is worrying because it casts doubt on the initial hypothesis above that a collective’s epistemic vices must be the product of its individual members’ vices. By contrast, it suggests that a collective can operate without any epistemic vices despite being comprised of individually vicious agents. If this is the case, though, it seems difficult to see how we could motivate the idea that a collective could be epistemically corrupted by its constituent members if the individual’s vices cannot translate to the collective. Of course, one might argue that this example does not rule out summativism as a candidate for thinking about the epistemic corruption of collectives, since it still offers us the correct verdict in the earlier example of the group of scientists. This case seemed to be a clear instance of individuals epistemically corrupting a collective from the inside. As such, summativism should still very much be on the table.

Even if we accept this, though, it does not insulate summativism from potentially more damaging kinds of cases: ones where a collective is comprised of mostly virtuous individuals, yet the collective comes to exhibit an epistemic vice. For example, Reza Lahroodi (2007: 287) presents the case of a church committee comprised of six individuals. Five are open-minded about gay rights while the other is homophobically narrow-minded. However, none of the five knows that the sixth is narrow-minded. As it happens, the narrow-minded member comes to have a great deal of influence over the proceedings. For fear of intimidation from this member, the other five toe the church’s line on these issues and begrudgingly vote against gay rights issues. Given the committee’s record of voting against gay rights issues, the five come to justifiably (though falsely) believe each other to be narrow-minded. This results in the committee becoming narrow-minded about gay rights *as a group*.

By intimidating the five, one member is able to steer the committee down a narrow-minded path and the other five resign themselves to this state of affairs. This kind of case is

---

<sup>79</sup> de Ridder’s (2022) focus is on collective epistemic virtues. In addition to the interaction model, he identifies a further two ways in which collectives can instantiate epistemic virtues: the ‘addition model’, where virtuous individuals are added to a collective body; and the ‘emergence’ model, where collective epistemic virtue arises from non-virtuous or even vicious individuals.

important. In many collectives and institutions, it is not the case that *every* individual exhibit epistemic vices. Often, in fact, they display the contrary. The problem, again, is that summativism tells us that the church committee is open-minded because the majority of its members are open-minded. However, this seems to be the wrong conclusion to draw. For starters, we would not expect an open-minded committee to vote down gay rights issues, except in highly unusual circumstances, but it does just this. Second, since the majority of the members believe each other to be narrow-minded, summativism tells us that the committee believes it is narrow-minded *qua* group.<sup>80</sup> However, it seems difficult to reconcile the committee being open-minded with it also possessing a collective belief that it is narrow-minded. In fact, if summativism is correct, then the committee cannot even possess the collective belief that it is open-minded since most of its individual members believe precisely the opposite. When the actions of a collective clash with its professed values and beliefs, something has gone wrong somewhere. The failure of summativism to capture this aspect should lead us to consider an alternative view.<sup>81</sup>

If summativism is inadequate, what are our options? One thing the discussion above reveals is that there can be a clear divergence in the traits exhibited by a collective and the individuals who comprise it. In this way, the cases above serve as *divergence arguments*, a term introduced by Margaret Gilbert (1989), which appeal to divergences between individual- and collective-level behaviours and characteristics. These arguments are used to cast initial doubt on the prospects of summativism, and to help motivate the case for *non-summativism* about collectives.<sup>82</sup> Miranda Fricker (2010: 137-138) enumerates at least two types of non-summativism, one of which we have already encountered. Recall the example above of a debating society comprised of individually prejudiced members, but whose prejudices cancel each other out to generate neutrality. Here, individual characteristics cancel each other out at the collective level. Alternatively, individuals might adopt multiple *practical identities*, which entail commitments to accept or resist certain attitudes, behaviours, and goals that one would not otherwise do as a *private individual*.<sup>83</sup> For example, a Member of Parliament might privately accept the reality of global warming but find themselves bound by loyalty and political considerations to publicly toe their party's sceptical stance.

---

<sup>80</sup> For discussions of collective belief, see Gilbert (1994, 2004).

<sup>81</sup> I leave it an open question whether summativism can be an adequate approach to understanding collective agency *more generally*.

<sup>82</sup> For seminal work on non-summativism about collective agency, see Bratman (2014), Byerly and Byerly (2016), Fricker (2010; 2020), Gilbert (1989; 2000; 2013), Lahroodi (2007), List and Pettit (2011), and Tuomela (2013).

<sup>83</sup> Fricker borrows the notion of a 'practical identity' from Christine Korsgaard's *The Sources of Normativity* (1996).

What all non-summativist accounts share in common is the view that a collective  $C$  can possess a quality  $Q$  without all (or any) of its individual members possessing  $Q$ . I shall have recourse to a particular non-summativist account shortly but notice that it would correctly attribute narrow-mindedness to the church committee despite the majority of its members being open-minded. This is important because without this attribution, we cannot make sense of the plausible idea that the committee's conduct has undergone some form of epistemic deterioration. In light of this, non-summativism looks to be a better way of motivating the sort of statements I canvassed at the start of this section – and with it, a basis for thinking about the epistemic corruption of collective agents.

## 1.2 The Core of Ethos

Having seen that collective agents can possess the sort of qualities relevant to a corruptionist critique, we can now begin asking whether the *source* of those qualities is the same for both individuals and collectives. In other words, if the source of an individual's traits, virtues, and vices is their character, to what extent can collective agents be said to possess something like a character? In several places, Miranda Fricker (2013, 2020) has taken up this question, and in doing so offered an account of what she calls *institutional ethos* (or simply *ethos*).<sup>84</sup> According to Fricker, ethos is meant to serve as the 'institutional analogue of an individual's character' (2020: 90). For this reason, she understands an institution's ethos to be the 'collective motivational dispositions and evaluative attitudes within the institutional body' (*ibid.*: 91). How does an institutional body instantiate these features? On Fricker's view, there are two 'core' features that allow for this, each of which I shall address in turn.

The first core feature of institutional ethos is its reliance on Margaret Gilbert's *joint commitment* model of collective agency (1989, 2000, 2013). A joint commitment emerges when two or more agents, under conditions of common knowledge, commit to some end together (2000: 21). This crucially involves a willingness or expression to commit that is 'out in the open' as far as the committing parties are concerned (Gilbert, 2023: 219). Unlike a personal commitment, a joint commitment is not something that members can abandon individually but rather is akin to a joint enterprise, where the committing agents enter into a position that

---

<sup>84</sup> The reason why I focus on Fricker's account is because it is the most detailed in the literature on collective virtue epistemology. That said, de Rooij and de Bruin (2022) have recently offered a 'functionalist' account of collective character. As such, my arguments here extend only to Fricker's account.

can only be rescinded ‘together’ (2000: 22). Joint commitments help form the basis of what Gilbert terms a ‘plural subject’ (1989). Accordingly, institutional ethos emerges when those party to an institutional body jointly commit, under conditions of common knowledge, to a set of values and attitudes.

Fricker argues that Gilbert’s model of joint commitments best captures the spirit of ethos for three reasons.<sup>85</sup> First, jointly committing to an institutional body often involves a *binding* commitment to its values – be that via a legal, moral, or contractual commitment. Part of this binding aspect is cognitive, in that jointly committing to an institution entails that each member ‘take on’ a responsibility to carry out that institution’s ends, which itself involves being aware that one is jointly committing (2010: 245). This can happen, Fricker says, when participants act in ways that are separate from their ‘private identities’ as individuals. Instead, they come to act in accordance with their *practical identities*. Consider somebody who works as an executive for a corporate oil company. As a private individual, they may sincerely believe in global warming, but as an executive feel compelled to vote against environmentally friendly measures on account of their practical identity. By adopting this practical identity, the executive expresses a willingness under conditions of common knowledge to jointly commit to the oil company’s climate-sceptical values. It is this feature of joint commitments, Fricker claims, that allows institutional ethos to be a coherent entity (2020: 91).

Second, there is the *rebuke* feature of joint commitments. This is a largely practical undertaking, such that if one attempts to unilaterally renege on a joint commitment, one becomes an adequate object of criticism or blame. Fricker adopts this condition to ensure that people cannot stray from the values expressed by an ethos without some form of punishment or cost to themselves. For those in precarious positions, the normative force of an ethos can be powerful and difficult to resist, even if a consequence is one’s complicity in certain institutional behaviour. This is especially relevant for those Fricker (2010: 248) calls ‘passengers’ – those who simply ‘go along’ with the ethos’ values *qua* members of the institution, due to their need for financial stability or support. Fricker claims that this feature of joint commitments helps to ensure that institutional ethos is ‘temporally and counter-factually stable’ in a way analogous to individual character (2020: 95).

Third, joint commitments can be entered into by *implicit agreement*. If one fails to speak up against a certain practice or policy, or if one goes along with existing practices in the sense

---

<sup>85</sup> Fricker acknowledges that summative views will sometimes be necessary to create an ethos (2020: 93).

of being a ‘passenger’, one thereby implicitly agrees to the content of the joint commitment.<sup>86</sup> As Fricker (2020: 97) puts it, joint commitments can be entered into ‘frighteningly easy’. The upshot of this implicit agreement condition is that ethos does not require explicit affirmation from those jointly committed to it. This has the benefit of allowing ethos to be something that is quite organic, in that it can undergo incremental changes without demanding the wholesale affirmation of the values from members (Gilbert, 2000: 5). By simply going along with the status-quo, one expresses an implicit willingness to jointly commit to the new values, goals, and attitudes expressed in the ethos.<sup>87</sup> In sum, Fricker claims that the joint commitment model allows us to appreciate how institutions can possess stable values and motives that we might similarly chalk up to an individual’s character.

The second ‘core’ feature of institutional ethos is its bifurcated structure. In furnishing us with this concept, Fricker is keen to not just construct the institutional analogue of individual character, but the analogue of *intellectual* character. To make good on this, she makes a distinction between what she calls the ‘inner ethos’ and its ‘outer performative’ element (2020: 98-99). This way of carving up ethos stems from her decision to adopt Linda Zagzebski’s (1996) view of intellectual virtue, according to which virtuous conduct requires good epistemic motivations and reliable success in implementing the ends of those motivations.<sup>88</sup> Accordingly, the ‘inner ethos’ contains the epistemic ends to which members of an institutional body jointly commit. Some institutions, of course, are distinctively epistemic, such as those in science, education, or the media, but most will have some epistemic function and ends. If those party to an institution jointly commit to good epistemic ends – like a concern for the truth or a desire for understanding – then its inner ethos will be epistemically good. Conversely, if there is an ‘inadequate commitment to good epistemic ends’, then the inner ethos will be epistemically bad (Fricker, 2020: 98).

Sometimes, however, an institution’s conduct is not best measured by the quality of its inner ethos. Indeed, Fricker (2020: 99) claims that an organisation can be epistemically dysfunctional despite its inner ethos involving a joint commitment to the truth. Over a period of time, imagine that a company’s attempts at realising its good epistemic motives fall flat: the department tasked with rolling out a certain policy has failed to communicate with

---

<sup>86</sup> Not all cases of silence indicate a willingness to jointly commit. Alessandra Tanesini (2018b) has drawn attention to what she calls ‘eloquent silences’, where silence acts as a form of dissent or protest to some action.

<sup>87</sup> The way Fricker construes this willingness is similar to David Lewis’ (1979) notion of ‘accommodation’. See, also, Langton (2018).

<sup>88</sup> It is worth noting that this way of conceiving of epistemic virtue has its detractors, including Montmarquet (1987), Riggs (2001), and Baehr (2011).

another department, which itself has failed to receive the results from another department. And the institutional snowball grows. Here, the problem is primarily external, namely, that the company fails to implement the policies or practices that would enable it to make good on its epistemic motives. This reflects a failure in the ‘outer performative’ element of its ethos (Fricker, 2020: 99). Together, this suggests that a *virtuous epistemic ethos* is one in which individuals party to a collective jointly commit to good epistemic motives in the inner ethos and reliably implement those motives in practice. Though Fricker does not use this term, this seems to ultimately be what is at stake.

## 2. Corrupted Epistemic Ethos

Fricker is keenly aware that institutional ethos is a fragile entity. One consequence of framing the concept in terms of joint commitments, she admits, is that it reveals how institutional bodies can ‘act in ways that depart from, and help deteriorate, a pre-existing epistemic ethos’ (2020: 97). If this occurs, Fricker says, there is room for what she refers to, in passing, as the *corruption of institutional character* (2020: 98, my italics). In this section, I will flesh out these remarks by articulating two ways in which this can occur. In section 2.1, I draw attention to what I refer to as a *passively corrupted ethos*, whilst in section 2.2 I supplement this picture with what I call an *actively corrupted ethos*.

### 2.1 Passively Corrupted Ethos

If the corruption of institutional character turns on institutional bodies acting in ways that ‘depart from, and help deteriorate, a pre-existing epistemic ethos’, how might this state of affairs arise? We can address this initial question by supposing that the ‘pre-existing epistemic ethos’ in question is virtuous. As I noted above, an institution’s epistemic ethos will be virtuous to the extent that those party to the institution jointly commit to good epistemic ends in the ‘inner ethos’, and the institution is able to reliably implement those ends in its ‘outer performative’ element. In this way, a virtuous epistemic ethos is one that not only enjoys systematic harmony between its internal and external components, but one from which distinctly institutional epistemic virtues flow. Hence, if those party to an institution jointly commit to being sensitive to evidence or attentive to truth-related considerations, and the institution succeeds in making good on this, then the institution is virtuously epistemically

conscientious *qua* institution.

Insofar as a virtuous epistemic ethos – and institutional epistemic virtue – is a matter of systematic harmony between the inner ethos and the outer performative element, it suggests that systematic *disharmony* is what renders an institution or collective vulnerable to the corruption of its character. Indeed, Fricker’s own examples seem to confirm as much. In one example (*ibid.*: 96-97), she asks us to imagine a branch of government whose individual officers all endorse, or at least go along with, an ethos of truthfulness and honesty. Despite this, there is suspicion surrounding one particular officer that he is engaging in precisely the opposite sort of behaviour; they are lying and bullshitting for their own political gain. As parties to the joint commitment underpinning this ethos, the individual officers have the standing to rebuke this member. Instead of doing this, though, they are silent; nobody is willing to address the suspicions openly and there comes to exist a pretence that everything is as it should be. By failing to call out this behaviour, these officers not only express an implicit willingness to preserve the status quo, but in doing so they generate a new joint commitment to opacity and dishonesty, both of which are in tension with the values in the inner ethos. Their willingness to tolerate this state of affairs means that the officers end up ‘betraying whatever may be left of the decent epistemic ethos of truthfulness’ (2020: 97). The upshot, Fricker says, is that the good epistemic values in the inner ethos are likely to be ‘eroded in the medium-to-long-term’ (*ibid.*: 98).

If this is how we are to understand the corruption of institutional character, then one thing should become immediately clear: it bears striking resemblance to what we know as passive epistemic corruption. Just as the latter can occur when social conditions erode an individual’s epistemic virtues, so the former takes place when those party to an ethos facilitate the erosion or loss of good epistemic values found in the inner ethos. Although Fricker does not make any explicit connection between the two, this seems to be what is going on throughout her discussion. Interestingly, Kidd (2021) has picked up on this connection, sketching an account of what he calls *institutional epistemic corruption*. He provides the following general definition of this phenomenon:

[E]pistemic corruption occurs when the character or *ethos* of an agent or *institution* is damaged by conditions, events, or processes that tend to facilitate the development and exercise of epistemic vices (*ibid.*: 353, my italics).

Kidd characterises this in terms of the ‘corruption of collective motivations that guide the institution and/or the corruption of its performance, of its processes and their outcomes’ (*ibid.*: 352). As this makes clear, the notion of ethos is central to Kidd’s account of institutional

epistemic corruption. We can assume, therefore, that the relevant *corruptee* will be an institution's ethos, and more precisely the two components above: its inner ethos and/or its external processes and outcomes. Similarly, we can expect the *corruptors* to be individuals or groups within an institution who facilitate the damage or decline of its ethos. Ironically, this is not what we find. Instead, Kidd observes that the corruptees will be 'politicians, civil servants, non-governmental agencies, and so on', while the corruptors turn out to be 'conditions, norms, or arrangements' (*ibid*: 353). Amongst other things, these include 'the absence of exemplars of virtue', the 'suppression of critics', and the 'valorisation of vicious conduct' to name just three (*ibid*: 354).

While the conditions and arrangements just noted might be epistemically corrupting in their own right, Kidd's account does little to tell us *how* they might have come about. Though he says that these conditions are largely a matter for historical and sociological investigation, these conditions, norms, arrangements are presumably the product of individuals within the institution. But this would then make the likes of individual politicians, civil servants and so on the *corruptors*, contrary to what Kidd says. Moreover, if individuals constitute these conditions within an institution, and these conditions are epistemically corrupting, then this explains why Kidd opts to class 'politicians, civil servants, non-governmental agencies, and so on' as *corruptees*, for they will be the one's subjected to the corrupting effects of these conditions, not the collective motivations and processes central to the ethos. This suggests that Kidd's focus is still on the epistemic virtues and vices of *individual* agents and not those manifested by an institution. If this is how we are to understand institutional epistemic corruption, though, then not only is it unclear how this differs from the Aretaic conception of epistemic corruption we identified in chapter 2; it is unclear what makes this phenomenon *institutional*.

On the view I wish to vindicate, by contrast, the focus is on how individuals party to a collective or institutional body can epistemically corrupt its ethos through acting (or failing to act) in certain ways. What initially emerges is a way of appreciating how the culpable failure of individuals can erode the good epistemic values found in an institution's inner ethos, thereby damaging its status as a virtuous institutional epistemic agent. This occurs because of what we might call a *passively corrupted ethos* – a drift from virtue:

*Passively Corrupted Ethos*: when members of a collective or institution jointly commit to good epistemic values in the 'inner ethos' but persistently fail to implement those values in the 'outer performative' element.

There are three points I would like to draw attention to regarding this notion. First, the processes that lead to this state of affairs need not materialise through any explicit action contrary to the epistemic values contained within the inner ethos. As the example shows, passively corrupting an ethos can simply be a matter of failing to speak up against behaviour that runs counter to the initial joint commitment to good epistemic values. Second, a passively corrupted ethos is likely to operate cyclically, such that tolerating lapses in the outer performative element erode the good epistemic values found in the inner ethos. By failing to implement their commitment to honesty and truthfulness in the example above, members in the branch of government made it easier for good epistemic values to wane in the inner ethos. Third, the notion of a passively corrupted ethos complements passive epistemic corruption. By tolerating lapses in the outer performative element, that is, the harder it will be for individuals in an institution to develop and manifest individual epistemic virtues. If a government department has diverged from its commitment to honesty, and the employees tolerate this state of affairs, they not only signal their willingness to go along with a new joint commitment contrary to honesty. In doing so, they help create conditions *within* the department that make it harder for them to cultivate the virtue of honesty.<sup>89</sup>

Given the parallels between passive epistemic corruption and a passively corrupted ethos, it will be no surprise to find that the latter emerges in similar ways to the former. Though Kidd does not quite capture my idea of a passively corrupted ethos, he does helpfully articulate a number of ways in which I think such an ethos might emerge (2021: 354-355). These include:

(I) *An absence of virtuous exemplars within an institutional body* – suppose members of an institution jointly commit to good epistemic values, but none of those members actually display these values in their own intellectual conduct. In the absence of these ‘exemplars’, not only will there be few, if any, opportunities for other members to shore up the good epistemic values; it will be easier for those members to tolerate failures to implement them.

(II) *Increasing the exercise cost of good epistemic values* – if it becomes increasingly difficult to act on the good epistemic values within an institution’s inner ethos, one alternative will be to resist doing so. This notoriously occurred in the aftermath of the so-called

---

<sup>89</sup> This reveals the dynamic and diachronic nature of a passively corrupted ethos, which is arguably absent from Kidd’s account of institutional epistemic corruption above.

*Climategate* scandal, when climate sceptics intimidated scientists for their perceived attempts at inflating the significance of climate change (Biddle, Kidd, and Leuschner, 2017). Despite their commitment to climate research – and the good epistemic values that went along with it – it became increasingly harder for the scientists to carry out this research for fear of intimidation (Mann, 2013).

To these three, we can add a further two:

III) *The absence of skilled epistemic conduct* – in order to implement the epistemic values found in an institution's inner ethos, members may require certain intellectual skills. For example, if a team of journalists jointly commits to an ethos of intellectual curiosity, it will be incumbent on them to ask good questions. But if none possess this intellectual skill, then it will be difficult for them to perform this activity and therefore live up to the intellectually curious values in the ethos.

(IV) *Frustrating the development and exercise of intellectual skills* – a journalism school might jointly commit to values of intellectual curiosity, rigour, and carefulness, yet make it difficult for its students to actually develop the intellectual skills associated with these values. It might, for example, cancel a course that teaches students how to ask good questions or discontinue a course that explains the importance of cross-checking stories. Without these courses, it will be harder for students to implement the good epistemic values found in the ethos.

As before, these features are likely to operate in tandem with each other. Where there is an absence of virtuous exemplars within an institutional body, for example, it might be more costly to exercise the good epistemic values contained in the inner ethos, and *vice versa*. After all, if it is difficult to implement these good epistemic values, and there is no inspiration for members to do so from exemplary figures, then there will be a drain on these values in the long-run. Likewise, an institution that fails to support the development of its members' intellectual skills will frustrate the exercise of any good epistemic values associated with those skills in the inner ethos. Thus, a passively corrupted ethos can emerge in many overlapping ways.

## 2.2 Actively Corrupted Ethos

We have just seen how a passively corrupted ethos emerges when those party to an institution or collective jointly commit to good epistemic values, but persistently fail to implement those values. I now want to suggest that there is another way of understanding the corruption of institutional character, one that Fricker does not explore, but which nonetheless captures a range of cases in which a collective or institutional body degrades, epistemically-speaking. This occurs when members of an institution or collective body actively encourage the uptake and performance of epistemically bad values. Let us refer to this state of affairs as an *actively corrupted ethos* – a drive to vice.

Before I cash out this variant, it is worth noting that there are at least two ways of thinking about it. On the one hand, it could involve individuals jointly committing to the presence, preservation, or sustenance of epistemic ‘bads’ such as ignorance, falsehoods, and so on. On the other hand, it might involve individuals culpably failing to commit to some epistemically good values, such as a commitment to the truth or a commitment to be sensitive to evidence. I raise this issue because Fricker contends that the former kind of joint commitments will be ‘unusual at best’ (2020: 99). This seems *prima facie* plausible. After all, it seems dubious to say that those party to an institution could jointly commit to bad epistemic ends like being informationally-disjointed or highly bureaucratic (Holroyd, 2020). Institutions *just are* organised in these ways regardless of any commitment. For this reason, Fricker claims that any motivational disorder constituting an institutional epistemic vice will take on the negative form of an ‘inadequate commitment to good epistemic ends’ (2020: 99). Accordingly, if an actively corrupted ethos does involve a drive to institutional vice, then we might wish to conceive of the ‘bad epistemic values’ in these terms.

However, this avenue is not exactly straightforward either. As Jules Holroyd (2020) has pointed out, a failure to commit to a virtue need not always (if ever) signal the presence of a vice. To borrow her example, an organisation might not jointly commit to intellectually courageous ends, but this failure alone should not lead us to think of it as intellectually cowardly (*ibid.*: 137). Similarly, a research institution might fail to jointly commit to open-mindedness, but it would be too quick to say that it possessed the vice of close-mindedness. Very often, she notes, failures to jointly commit to some good epistemic ends just ‘signify that a group has other priorities’ (*ibid.*: 137): the research institution might simply prize intellectual rigour over open-mindedness. In reality, institutions will often have a myriad of competing priorities, a great deal of which will be primarily or secondarily epistemic. On Fricker’s view, though, culpable lapses in the inner ethos – i.e., inadequate commitments to good epistemic

ends –which arise from any of these priorities would be sufficient for institutional epistemic vice (2020: 101). Not only does this make her account highly demanding, but it would put most institutions at risk of possessing an actively corrupted ethos.

This conclusion is reminiscent of one we encountered in the previous chapter regarding the remit of passive epistemic corruption. If it turned out that many, if not all, of our social environments were passively corrupting, recall, we would seemingly diminish the value of a corruptionist critique. To avoid this, I suggested that we focus on environments in which there is a concerted effort to erode people’s epistemic virtues, traits, skills and so on. I think a similarly irenic conclusion is warranted with regards to an actively corrupted ethos here. This means understanding an actively corrupted ethos in terms of joint commitments to falsehoods, deception, ignorance and so on. Thus, we would arrive at the following definition:

*Actively Corrupted Ethos:* when those party to an institution or collective jointly commit to epistemically bad values in the inner ethos and/or reliably implement the ends of these values in its outer performative element.

While a joint commitment to these values might be rather rare, the reality is that they often reflect the most serious cases of institutional epistemic deterioration. In his book *Golden Holocaust*, for example, Robert Proctor (2011) observes how many Big Tobacco companies in the twentieth century actively undermined scientific evidence concerning the health and environmental effects of smoking, in order to keep people ignorant and committed to buying tobacco products. Similarly, Naomi Oreskes and Eric Conway (2012) persuasively show how oil companies were effective ‘merchants of doubt’, undermining data about the extent to which oil contributes to anthropogenic climate change. In both cases, the officials party to these institutions were clearly committed, explicitly or implicitly, to what sociologists of knowledge refer to as *agnotology*, or the deliberate manufacture and deployment of targeted doubt and ignorance.<sup>90</sup> This, itself, is an expression of what Jason Baehr (2010: 191) calls ‘impersonal epistemic malevolence’ – an active drive to ‘stop, diminish, undermine, destroy, speak out, or turn others’ against the epistemic good.

These cases reveal an important difference between a passively corrupted ethos and an actively corrupted one. The former, we noted, can occur through the inaction of members

---

<sup>90</sup> A more recent, but no less sinister example, is that of Purdue Pharmaceutical company, which obfuscated evidence and suppressed knowledge about the harmful effects of its notorious ‘painkiller’, OxyContin. See Meir (2018). Cassam (2019a) notes that doubt-mongers are also guilty of epistemic malevolence.

party to an ethos going along with performative failings. The latter, meanwhile, depends on members taking an active role in encouraging epistemically bad values. Of course, though, just because there are members who encourage these values in the inner ethos, it does not follow that they will be taken up in the outer performative element. It is reasonable to think that anyone with minimally good epistemic values might attempt to frustrate the spread of bad epistemic values within an institution. However, this just speaks to the second difference between the two forms of a corrupted ethos.

A passively corrupted ethos does not require members to hold bad epistemic values for it to occur – in fact, it need not even require that they hold good epistemic values, just that they are willing to go along with good epistemic values, regardless of whether they recognise them as such. By contrast, an actively corrupted ethos will most likely depend on there being individual members who possess a range of epistemic vices for its continued existence. In this way, an actively corrupted ethos can emerge by replacing existing personnel with those who are willing to act on its values. This occurred, for instance, during the roll-out of the UK Home Office’s infamous ‘hostile environment policy’. Though targeted at illegal immigrants, in general, it is remembered for its specific treatment of Britain’s Caribbean population in what came to be called the Windrush Scandal. Prior to the policy, an ex-Home Office employee reported how her colleagues were typically sensitive to the evidence of each immigration claim. To ensure that the policy was successful, though, the Home Office quickly and deliberately began recruiting people who ‘were more aggressive in their attitude to undocumented Windrush-generation residents’ (Gentleman, 2018), and that there was a ‘mentality’ of ‘I’m going to say no, unless you can prove me wrong’ (*ibid*, 2018). Amongst these new employees, however, the threshold for evidence was cynically high, effectively to the extent of infallibility. As a result, the ex-employee admits how ‘many experienced staff took redundancy because they disliked the *new atmosphere*’ (*ibid*: 2018, my italics), and thus there was an exodus of people with minimally good epistemic (and moral) values.<sup>91</sup> In short, the Home Office became a place that was hostile to virtuous acts and inhospitable to virtuous people.

What seems constant for both, though, is their cyclical nature. The more members promote or encourage bad epistemic values in the inner ethos, the more these values will be implemented in practice; and as these ends are implemented, the more these bad epistemic

---

<sup>91</sup> See Wendy Williams’ *Lessons Learned* inquiry into the UK Home Office’s treatment of Britain’s Black-Caribbean population for further discussion. In the report, she attributes ‘institutional ignorance and thoughtlessness’ to the Home Office (2020: 13). I think it’s worth quoting the in the main body.

values will congeal in the inner ethos. We can make good on this by enumerating a number of ways in which an actively corrupted ethos might materialise. These non-exhaustively include:

(I) *Replacing virtuous members with vicious ones* – to shore up support for its bad epistemic values, an institution might encourage those with minimally good epistemic values to ‘move on’, or it might dismiss those members entirely. In their place, it will employ individuals who exhibit the required traits and dispositions needed to execute the jointly committed values.

(II) *Valorising vicious conduct and exemplars* – an institution might praise or esteem an employee’s close-mindedness or cynicism, holding them as a paragon of success to be emulated by others.

(III) *Suppressing critics* – to ensure bad epistemic values are realised, members might suppress from view, intimidate, silence, or demote those who pose a threat to those values.

(III) *Rewarding the exercise of vices* – a tobacco company might reward its scientists who publish dishonest work that tries to discredit links between smoking and cancer (Proctor, 2011). An oil company could introduce financial incentives so that its employees work to tirelessly undermine evidence pointing to anthropogenic climate change (Oreskes and Conway, 2012).

(IV) *Establishing policies which require developing and manifesting vices* – to effectively carry out its Hostile Environment policy, the Home Office expected its employees to vastly limit immigration applications. Not only did this encourage employees to negligently rush through applications, but it also fostered a deep cynicism towards applications from British-Caribbean people.

(V) *Encouraging the development of vice-indexed skills* – if the performance of bad epistemic values helps solidify these values within the inner ethos, then an institution might be well-served by ensuring that its workforce know how to effectively plagiarise work, bullshit in front of scientific panels, hide evidence from regulatory bodies, and so on.

Depending on how it is organised, then, an institution might be more or less susceptible to developing a passively corrupted or actively corrupted ethos. However, where we find one, we are likely to find the other. Insofar as members of an institution persistently tolerate lapses

in the performative element of its ethos, despite jointly committing to good epistemic values, there is a good chance of those values being eroded in the inner ethos. However, as those values erode, it will be easier for bad epistemic values to gradually infect that institution's inner ethos. If an institution begins actively taking up bad epistemically values within its inner ethos, it can gradually replace personnel so that it excels at performing the ends of these values in its outer element. As with passive and active corruption, then, the two kinds of a corrupted epistemic ethos are likely to travel together.

### 3. The Limits of an Analogy

I spent the first section of this chapter motivating the idea that collectives and institutions can possess qualities relevant to a corruptionist critique *in their own right* and are thus vulnerable to a form of epistemic corruption. In section 2, I articulated two ways in which this might occur. The result is what I call a *corrupted epistemic ethos*, where individuals party to an institution or collective either (i) culpably fail to implement the good epistemic values in an institution's ethos and/or (ii) encourage the uptake and performance of epistemically bad values.

In this section, I want to turn my attention to Fricker's claim that ethos is the 'institutional analogue of an individual's character' (2020: 90). More specifically, I argue that the analogy she draws between individual and institutional character only goes so far. As a result, institutional ethos is only able to capture a certain range of cases in which epistemic corruption occurs at the collective or institutional level. This, I will suggest, gives us reason to be wary about the effectiveness of a purely corruptionist critique of collectives and institutions.

In section 3.1, I argue that the common knowledge required for joint commitments prevents Fricker from attributing an ethos to large-scale collectives without surrendering her non-summative commitments. In section 3.2, I contend that the uptake (and loss) of values within an ethos is fundamentally different from the habituation of values within individual character. This leads me, in section 3.3, to ultimately urge caution about the practicality of an account of epistemic corruption at the collective level.

### 3.1. Ethos and Common Knowledge

My first worry about the analogy between individual and institutional character concerns a central feature of Margaret Gilbert's model of joint commitments, which Fricker takes up into her account of institutional ethos. This is the claim that jointly committing to an ethos 'requires individual members to express willingness, under conditions of *common knowledge*' (Fricker, 2020: 96, my italics). To set up this worry, it will help to briefly consider how Gilbert understands common knowledge.<sup>92</sup>

By *common knowledge*, Gilbert roughly means the sense introduced by David Lewis (1969: 52-57), according to which there is a hierarchy or chain along which a proposition or intention is known. More precisely, if Paul knows that Lynn wants to go for a walk with him, and Lynn knows that Paul is willing to join her on the walk, then there is common knowledge between the two about going for a walk. Put simply, the willingness or expression on the part of the committing agents is 'out in the open' (Gilbert, 1989, 2001, 2023: 219). As I mentioned earlier, this expression or willingness need not be explicit. Suppose Paul and Lynn lace up their walking boots together in silence. In this case, their actions signal an implicit willingness to go for a walk together. These are instances of what Gilbert refers to as *individual common knowledge* (2006: 176). Individual common knowledge seems to be the kind at stake for institutional ethos. This is clear from Fricker's example of 'an institutional body such as a branch of government' in which 'the individual officers all endorse, or at least go along with, the set of values that comprise an ethos of truthfulness and fact-checking' (2020: 96). Common knowledge emerges in this context because the officers all go along with the ethos' stated values.

Here is where a problem emerges. For common knowledge to emerge, it must be the case that the individual officers within the branch of government '*all* endorse, or at least go along' with the values comprising the inner ethos (my italics). However, this presupposes that there are clear channels of communication within institutions such that common knowledge can freely permeate across them. Many of the institutions which comprise the fabric of our social environments, however, are not structured in this way. Often, they are characterised by heavy divisions of labour, organisational silos, and poor channels of communication – a recipe for what Carel and Kidd (2021: 481) refer to as *institutional opacity*, that is, a tendency

---

<sup>92</sup> It is worth stating that some form of common knowledge requirement is not unique to Gilbert but is widely embraced by a range of theorists working on shared intention and collective agency. See, for example, Bratman (2014), Pettit and Schweikard (2006), and Tuomela (2013). An exception is Searle (1990).

for large-scale and internally-complex institutions to become resistant to forms of assessment and understanding. Not only will it be hard for common knowledge to emerge in these institutions, but without such knowledge, those party to an institutional body cannot jointly commit to an ethos. And without an ethos, there cannot be a *corrupted epistemic ethos*.

Fricker's (2020: 102-4) own case study nicely illustrates this point. In furnishing us with the concept of ethos, she identifies an exclusively institutional epistemic vice she terms *inferential inertia*' (2020: 101).<sup>93</sup> This occurs when a hearer receives credible information from a speaker, but where there is no 'epistemic follow-through' on the hearer's behalf. Salient information is not reliably transmitted through the institutional structure. As a result, people within the institution do not receive that information, meaning they fail to draw the relevant inferences (*ibid*: 101). Fricker argues that this vice plagued the BBC during the 1970s and 80s, during which the TV entertainer Jimmy Savile committed sexual crimes at the company. Drawing on the Dame Janet Smith Report, she notes how the BBC exhibited a 'culture of separation' and 'silo mentality' during this time, meaning that that complaints and suspicions 'would not go any further in the epistemic economy of the organisation' (2020: 102).<sup>94</sup> This was only fuelled by 'unintegrated' departments and the 'lack of channels of communication' at the BBC (*ibid*: 102-3). By Fricker's own admission, the institution's 'informational states were radically unintegrated' (*ibid*: 103).

But if the BBC's informational states were radically unintegrated, how could we expect Savile's behaviour to 'be out in the open' for employees in departments far removed from his own to have known, at least at the time of its occurrence? Of course, this is not to dismiss the possibility of common knowledge amongst employees in those departments in which Savile worked. Those who worked in the music and entertainment departments, for example, would have certainly expressed an implicit willingness, under conditions of common knowledge, to go along with what was occurring when they failed to investigate complaints further. Indeed, this would have sufficed to generate the vice of inferential inertia within these departments. However, this does *not* suffice to generate the joint commitments needed to instantiate institutional ethos. For this commitment to arise, not only would all BBC employees need to express their willingness to go along with Savile's behaviour, but they would have needed common knowledge to do so. Given that the institution was so radically disjointed, though, it is difficult to see how both conditions could have been satisfied, and hence how a large, disjointed institution such as the BBC could have had an ethos *qua* institution.

---

<sup>93</sup> An 'exclusive' institutional or collective virtue or vice is one that only exists at this level, and not at the individual level. Byerly and Byerly (2016) cite 'solidarity' as an exclusively collective virtue.

<sup>94</sup> See the Dame Janet Smith Review Report (2016) for an in-depth investigation of this scandal.

To avoid this difficulty, Fricker could adopt a broader notion of common knowledge. In addition to individual common knowledge, Gilbert introduces what she calls *population common knowledge* (2006: 176), which is knowledge between people ‘considered by those involved as members of a population individuated by means of a certain general description’ (*ibid.*: 176). For our purposes here, examples might include “those who work in HR” or “those who work in finance”. Instead of requiring individual common knowledge, Fricker could say that individual *departments* – treated as specific organised populations of staff members – require common knowledge for an institution to possess an ethos. When enough departments satisfy this *qua* populations, they will be in a position to jointly commit. This move could allow Fricker to overcome the difficulties above. Assuming that there are well-designed mechanisms in place, news about Savile’s behaviour would travel from one department to another, eventually becoming common knowledge within the BBC.

Nevertheless, if Fricker were to employ this sense of common knowledge, her account would face similar difficulties. Whereas institutional ethos required all individuals to express a willingness under conditions of common knowledge beforehand, it would now just require all departments to do the same. Once more, though, the success of this depends on these departments actually coming to have common knowledge, which again is no guarantee if the institutional body is large and highly disjointed, as was the BBC. This complication aside, there remains a potentially more troubling problem for this move. Insofar as Fricker intends for ethos to be the institutional analogue of individual character, it should be attributable to a single agent, i.e., the institution in question. If she were to adopt population common knowledge, however, this would no longer be the case. Instead, it would be the case that *each* department within the institution possessed an ethos, given that the common knowledge required for the necessary joint commitment arises from these settings first. Unlike an individual agent, then, an institutional agent would come to possess multiple “characters”, all of which would need pooling to represent its character *qua* institution.<sup>95</sup> Thus, an institution’s character or ethos would ultimately be the sum of its parts, and hence be *summativist*.

The problem with ethos being a summative concept is that it would no longer be able to capture an institution’s epistemic misconduct *qua* institution. At best, any such misconduct would have to be understood as arising from the individual departments within the institution. By extension, it would only be through aggregating the sum of each department’s

---

<sup>95</sup> This might seem *prima facie* plausible since I argued in chapter 1 that our intellectual characters form in a piece-meal fashion. However, the crucial difference is that our intellectual characters still remain a single metaphysical entity, just ones whose traits can manifest in localised domains. This is not the case above, since the constituent parts of an institutional body could be said to possess their own distinct characters, which is different from local traits manifesting from a single character.

epistemic virtues or vices that Fricker would be able to explain how an institution loses or gains these qualities *qua* institution. As Reza Lahroodi (2018: 411) reminds us, though, modelling institutional virtues and vices along summativist lines means that any attributions of virtue and vice are ultimately aimed at members of the institution as private individuals, rather than the institution itself. To capture large-scale institutions, therefore, Fricker must weaken the common knowledge condition required to instantiate an ethos. But in weakening this requirement, she risks collapsing the analogy between individual and institutional character.

### 3.2. Habituating Ethos?

In the previous section, I argued that the common knowledge condition on joint commitments introduces a dilemma for Fricker's account of ethos. If she wishes to attribute an ethos to large-scale, complex institutions, she must be willing to water down this common knowledge condition. If she does this, however, then her account is no longer able to attribute an ethos or character to such institutions in their own right. Not only does this reveal that institutional ethos – and the notion of a corrupted epistemic ethos – is best suited to more homogenous institutions, but that the analogy between individual and institutional character can only go so far. In this section, I put further pressure on the analogy by drawing attention to an important feature of individual character that does not neatly translate across to institutional ethos.

I noted in the introduction to this thesis how responsibilist or character-based virtue epistemology is best understood as a form of 'regulative epistemology' (Wolterstorff, 1996). As such, the principal focus is on evaluating, assessing, and guiding the epistemic conduct of agents. The focal point of this task is an agent's intellectual *character*. A similar point presumably holds for the institutional analogue of character in ethos. In attributing an ethos to an institution or collective agent, we want to evaluate and assess how it conducts its intellectual affairs. Indeed, Fricker's motivation for developing this concept is to provide us with a means of making evaluatively rich judgements about institutions that are continuous with the way we evaluate individuals (2020: 90). Just as an individual's character allows us to judge their values and commitments, she claims, so joint commitments are 'the distinctively appropriate model to employ in elaborating the collective value-dispositions involved in forming a given ethos' (2020: 95).

The first thing we can note is that the value-commitments we hold as individuals

usually develop in our characters over a prolonged period. As Linda Zagzebski (1996: 196) might put it, they do not form ‘with the flip of a switch’. Once they have formed, moreover, they tend to go ‘all the way down’, in the sense that we cannot decide overnight to cease acting in accordance with them (Hursthouse, 1999: 12). The fact that I am generally persevering, for example, is the result of me possessing certain values and attitudes towards doing work, which have taken a great deal of time to cultivate and entrench. If I were to wake up tomorrow and begin procrastinating, I would find it very difficult. To keep the analogy between individual character and institutional ethos intact, then, we should expect the ‘collective value-dispositions’ expressed by an ethos to require time and effort to habituate, as they do in the individual case.

The problem is that the values and motives found in an institution’s ethos do not mirror this process, precisely because of Fricker’s decision to model this concept on joint commitments. A joint commitment, recall, emerges when two or more agents *commit* to some end together, under conditions of common knowledge (Gilbert, 2000: 21). If those party to an institution decide to jointly commit to some good epistemic value – open-mindedness, say – this is sufficient for the ethos to possess an epistemically good inner ethos. This seems clear from Fricker’s example, where she describes a school as having a ‘good ethos, and so the inner element is fine’ (2020: 98). What makes the school’s ethos ‘good’ is that the teachers ‘care about doing a good job’, and so jointly commit to this end (*ibid*: 98). In this way, jointly committing to a specific value or end is enough for an institution to “habituate” that value or end in its inner ethos. This is clearly far removed from the time and effort it requires for an individual to habituate a given value or motive within their character.

A similar problem arises in connection with the reverse situation. Just as an institution can incorporate a value into its ethos by jointly committing to it, so too can it theoretically abandon that value by eschewing the joint commitment. Now, we need to be careful here because Gilbert is clear that joint commitments cannot be abandoned unilaterally without some form of practical rebuke to committing members. That being said, she does recognise that joint commitments can be rescinded ‘together’ (2000: 22). So, if the teachers in Fricker’s example decided to collectively withdraw their support for the joint commitment to ‘doing a good job’, say on grounds that they all found out that the headteacher had been funnelling school money into a private bank account, they would be perfectly entitled to do so. In doing so, however, this value would disappear from the school’s inner ethos, at least until the teachers entered into a new joint commitment. Again, though, this is something that individuals cannot do; once we habituate a certain value, it is extremely hard to shake.

The fact that an institution can readily decide which values to endorse or abandon in

its ethos speaks to a more general point about individual and institutional cognition. I noted in chapter 1 how our character traits have an important cognitive component. In addition to disposing us to act and feel in certain ways, they dispose us to form certain beliefs. If I am intellectually persevering, I might believe that hard-work in my epistemic affairs is a price worth paying. If I am open-minded, I might believe that it is valuable to listen to the views of others. The kind of traits we habituate, then, can largely determine the beliefs we form, at least those arising from our characteristic value-commitments and motives. Just as I cannot wake up tomorrow and decide to be intellectually courageous, neither can I wake up and decide, ‘with the flip of a switch’, to believe in an intellectually courageous way (Zagzebski, 1996: 196).

The idea that we can voluntarily choose what to believe is known as *doxastic voluntarism*, and it is highly contested whether humans are capable of it.<sup>96</sup> Whether or not doxastic voluntarism is contestable for beliefs regarding individual character, it is very much possible for an institution if we model ethos on joint commitments. By deciding to jointly commit to a set of values, not only can an institution decide whether to incorporate those values into its inner ethos, but in doing so it can form collective beliefs about those value-commitments on the basis of this decision (Gilbert, 1989). If we accept that this is not possible for humans, then we have found another disanalogy between individual and institutional character.

Of course, I expect there to be a number of objections to the argument here. First, one might claim that the values incorporated into an institution’s ethos do develop over time because joint commitments can be flexible or gradual. This possibility seems clear from the fact that, in our everyday language, we do talk of people becoming more or less committed to certain projects or ends, that our commitments can become stronger or weaker over time, or that they can simply “fizzle out”. In each case, the commitments in question are not zero-sum matters, but rather dynamic and diachronic features of our life, that are sustained, expanded, or eroded depending on ever-changing interpersonal and social conditions. If this is correct, the argument goes, then so, too, are the joint commitments that underpin institutional ethos. As a result, the values and collective motives jointly committed to within an institution’s inner ethos could be said to evolve and decline in a gradual manner, much like those found in an individual’s character.

---

<sup>96</sup> To be more precise, what is contested is the truth of *direct* doxastic voluntarism. One person who did believe this was possible was William James in *The Will to Believe* [1896] (2010). Some contemporary defenders of direct doxastic voluntarism include Montmarquet (1986), Ginet (2001), and Steup (2008).

While I certainly agree that the everyday sense of “commitment” can be stronger or weaker, this sense of commitment is fundamentally different from the kind central to institutional ethos. As Fricker makes clear, an ethos is a plural subject because it is created by joint commitments. For Gilbert, joint commitments involve a ‘kind of commitment of the will’, where the ‘wills of two or more people create it, and two or more people are committed by it’ (2006: 134). Thus, if the teachers at the school jointly commit to do a good job, they are normatively bound by each other to ensure this happens. If one teacher reneges on this commitment, they become an adequate object of rebuke from others. However, this does not occur in the everyday sense of commitment above. If I am weakly committed to washing my car, the fact that I am not bound by anyone means that nobody (bar myself, perhaps) has the standing to rebuke me when I fail to actually wash my car. What’s important for present purposes is that jointly committing to certain values or motives in the inner ethos cannot be stronger or weaker like in the ordinary sense above, since any changes to the level of commitment must be accepted, either explicitly or implicitly, by all members. Doing this, however, creates a new joint commitment to that value or motive. Any changes to the values within an inner ethos therefore fail to replicate the organic nature with which they develop or diminish in individual character.

A second objection is that I have overlooked the importance of the outer performative element when considering how an institution’s ethos might incorporate certain values or motives. It is not sufficient for an ethos to manifest a certain value, one might argue, unless that value is reliably implemented in this outer performative aspect. For comparison, if somebody confessed to being honest but reliably failed to tell the truth, we would not likely consider them honest. Likewise, an institution which confessed to being open-minded but failed to make good on this commitment in how it acted would not seem to merit the label “open-minded”. Once we factor in this reliability component, the analogy between individual character and ethos returns.

However, this confuses things. Fricker’s (2010) decision to insert a reliability component into her account of ethos stems from her earlier decision to model institutional virtues on Zagzebski’s account of intellectual virtue. On Zagzebski’s view, intellectual virtues not only require a motive for ‘cognitive contact with reality’ but also reliable success in implementing this motive (1996: 176). It is noteworthy that Zagzebski’s view is a minority position in the virtue-epistemological literature. Whilst many theorists think that epistemic virtues will be reliable in worlds like our own, they argue that making reliability a *necessary* condition renders our actual epistemic endeavours too contingent on luck (Baehr, 2011: 135-

38; Montmarquet, 1987: 482; Riggs, 2003: 211).<sup>97</sup> For the purposes of argument, however, let us grant that there is nothing immediately problematic about a success condition on epistemic virtue. Nevertheless, the points above reveal that a success condition is only motivated by considerations of epistemic *virtue*. The reason why Fricker includes this condition in her account of ethos is because she is interested in how an institution can be said to possess epistemic virtues and vices *qua* institution. By contrast, I am simply interested in how these values or motives form in the ‘inner ethos’, and reliability does not factor into this. What determines this are the joint commitments to these values and motives. Thus, one cannot appeal to Fricker’s outer performative element of ethos to salvage the analogy.

Finally, Fricker might claim that her aim in introducing institutional ethos was not to create the *literal* analogue of institutional character, but to simply carve out richer conceptual space for thinking about the actual character-based qualities we do attribute to institutions. This sort of reply would echo Gilbert’s (1989: 201, 2013: 8) earlier sentiment that plural subjects and human subjects exhibit a categorically different kind of agency. Joint commitments, she says, bind the parties to a plural subject so that they act *as if* they were a single agent. Accordingly, to speak of an institution forming “collective beliefs” about the values or motives to which it has jointly committed in its inner ethos would be a category mistake. Rather than think of these beliefs as *literal* beliefs, we should instead treat them like *as if* beliefs which arise in virtue of members’ commitments to ‘emulate’ an individual agent that holds these beliefs (*ibid.* 8). If collective or “plural” beliefs are emulations of individual beliefs, and individual beliefs cannot be voluntarily chosen, then neither can collective beliefs formed about the values within an ethos.

An important reason why Gilbert makes the distinction between plural and human subjects is to avoid the suggestion that her account gives rise to collective mental states, which would seem ‘metaphysically spooky’ (Fricker, 2010: 242). By treating collective beliefs as emulations of individual beliefs, Gilbert is able to side-step this worry. However, just because collective beliefs ought to be treated like the *as if* beliefs of individuals, it does not follow that they are like individual beliefs. Whereas individuals form their beliefs in a range of ways, plural subjects can only form their collective beliefs via joint commitments. Insofar as plural subjects form collective beliefs, then, those beliefs will remain the product of joint commitments, and these require the explicit or implicit agreement of members. Given that an institution’s ethos is a plural subject, any collective beliefs formed about its values and motives would therefore remain voluntarily endorsed or rejected, depending on the joint commitment

---

<sup>97</sup> Besides Fricker, van Zyl (2015) is an outlier.

in question. In short, doxastic voluntarism is built into the plural subjectivity of institutional ethos.

### 3.3. A Plurality of Explanations

In furnishing us with the concept of ethos, Fricker intends to create the institutional analogue of individual character. As we have seen, though, her decision to model ethos on joint commitments introduces a number of complications that expose the limitations of such an analogy. While these limitations might seem somewhat trivial, they are important in the grand scheme of things. For if it is assumed that ethos is the institutional analogue of individual character, then our search for an institutional or collective-based account of epistemic corruption would be over; ethos – and specifically corrupted epistemic ethos – would fit the bill in all cases. The fact that the analogy breaks down, however, should give us reason to avoid this conclusion. In the remainder of this chapter, I will urge caution about drawing too close an analogy between institutional and individual character.

Despite emphasising the differences between individual character and institutional ethos, I would like to first draw attention to a feature shared by both. When we describe a person or institution as “cynical” or “dogmatic”, what we are doing is attempting to attribute these qualities to the relevant agent. If we take cynicism and dogmatism to be intellectual vices, then the resulting descriptions are what Tanesini (2016) has called *vice-attributions*. Vice-attributions can function in a variety of ways. In attributing an intellectual vice to an agent, we might wish to accuse or criticise them for their poor epistemic conduct. Here, the vice-attribution functions as a ‘vice-charge’ (Kidd, 2016a). Not all vice-attributions need to be vice-charges, though. Attributions of vice can serve a variety of functions. Rather than criticise somebody, for instance, we might attribute them an epistemic vice in order to *warn* or *caution* others about their intellectual conduct (Cassam, 2020). However, an overarching feature of vice-attributions is their explanatory function. In charging the BBC *qua* institution with the vice of inferential inertia, Fricker seeks to explain *how* it is that concerns about Jimmy Savile went no further in the ‘epistemic economy of the organisation’ (2020: 102). In this way, vice-explanations are typically *causal* explanations of an agent’s epistemic misconduct.

Now, vice-explanations are by no means the only way of articulating poor epistemic conduct. Sometimes, the relevant epistemic failings will manifest at the sub-personal level, such as those relating to an agent’s cognitive characteristics or implicit biases (Brownstein and Saul, 2016; Holroyd, 2020). These explanations would be *cognitive* explanations.

Alternatively, we might chalk up an agent's epistemic misconduct to various, contingent situational factors like those expressed by the situationist challenge in chapter 1 (Alfano, 2013; Fairweather and Alfano, 2017). Call these *situational* explanations. Further still, an agent might conduct their intellectual affairs poorly due to their political or ideological commitments. We can refer to these as *political-ideological* explanations (Cassam, 2019b, 2022).<sup>98</sup> Finally, we might explain an agent's epistemic conduct by reference to the social structures in which they operate. Explanations of this sort would be *socio-structural* (Haslanger, 2015).

It is worth noting that these explanations are not mutually exclusive. Indeed, many are mutually illuminating. The fact that a person comes to believe a conspiracy theory might be explained by appeal to their underlying ideological commitments, but their ardent refusal to question this ideology might be because of their dogmatism or cynicism (Cassam, 2019b). Likewise, the efficacy of racist practices and policies might be explained in structural terms, but these practices and policies can quite easily be sustained and implemented by what Lisa Tessman (2005: 54-55) refers to as the 'ordinary vices of domination'. As vice epistemologists, we should not only be alert to the many supportive relationships between explanations, but we should be willing to admit when one kind of explanation is more apt than another. For this reason, I agree with Quassim Cassam (2019a: 23-27) that vice epistemology should not be committed to explaining all instances of poor epistemic conduct in terms of vices. Indeed, we should recognise the possibility of a 'sliding scale' (*ibid.* 51). As Cassam observes:

At one end are outcomes that can only adequately be understood in structural terms. Social inequality is an excellent example. At the other extreme are outcomes that can be adequately understood in vice terms. In the middle are many outcomes that have to be understood partly in structural terms and partly by reference to epistemic vices (*ibid.* 51-52).

An insensitivity to different explanatory factors in an agent's epistemic (mis)conduct is a one-way ticket to what Cassam terms 'intellectual myopia' – an epistemic vice that is potentially inherent to vice epistemology as a discipline (*forthcoming*: 4).<sup>99</sup> It is potentially inherent insofar as vice epistemologists may be at risk of myopically focusing on vice-explanations at the expense of other explanations. I think Fricker is somewhat guilty of intellectual myopia in

---

<sup>98</sup> As Cassam (2021, *forthcoming*) notes, political-ideological explanations differ from vice-explanations insofar as they often function as *rationalising* explanations.

<sup>99</sup> In addition to intellectual myopia, Cassam considers 'political naivety' and 'overconfidence' to be two further 'disciplinary vices' that potentially blight vice epistemology (*forthcoming*). The antidote to these vices, he suggests, is a healthy dose of intellectual humility.

her attempt to explain institutional epistemic misconduct solely by reference to ethos and vices. Consider, for example, her comments on institutional racism. She says that ‘If anything is a vice in an institution, then racism is’, before noting how the idea of an institutionally racist police force furnishes us with a ‘central and prominent example of something that can be properly theorised as *institutional vice*’ (2020: 90). The purpose of this language, she emphasises, is ‘one of picking out aspects of institutions that are the *collective analogue of an individual agent’s character*’ (*ibid.*: 90, my italics). Hence her reason for modelling institutional ethos in the first place.

Fricker’s decision to explain institutional racism in vice-theoretic terms – as well as her more general attempt to draw a close analogy between institutional phenomena and individual character – demonstrates an insensitivity to other salient explanatory factors that might configure in this phenomenon. An obvious candidate seems to be the *structural* dimensions of institutional racism. Without being sensitive to this feature, any attempt to pick out aspects of an institution that are analogous with an individual’s character will not only be explanatorily impoverished; they will myopically obstruct the knowledge and understanding that structural explanations could provide. Only in tandem with a plurality of explanatory styles, then, can concepts like “vice” or “ethos” fully do justice to the epistemic misconduct of institutional or collective agents. Since this is absent from Fricker’s discussion, we should likewise not expect the concept of ethos – nor corrupted epistemic ethos – to be able to fully capture the epistemic deterioration of collective agents. Indeed, the arguments above reveal just that.

Given the sheer number of possible explanations for epistemic (mis)conduct, why would we want to draw a close analogy between the character of individuals and the character of institutions? Sure, we might wish to capture the ordinary sense in which we attribute character-based qualities to collective agents. But if our characters are shaped (and misshaped) by cultural, social, political, and ideological factors, as Robin Dillon (2012: 104) contends, then why should we not expect similar results at the institutional level? Any account of institutional character would thus need to move beyond simply offering vice-explanations and assume a ‘multi-dimensionalist’ outlook (Kidd, 2023b: 26), one that endorses and employs a plurality of kinds of explanation according to the needs of each case. The problem with such an account, though, would be that it no longer counts as purely characterological in nature, but rather would be an amalgamation of different considerations. As such, it would not be right to call it an account of institutional *character*. If it is difficult to conceptualise a purely vice-theoretic account of institutional character, then the prospects of conceptualising epistemic corruption at the collective-level seem equally difficult.

I do not find this conclusion disturbing. As Cassam (2019a: 27) points out, ‘satisfying explanations of our intellectual conduct are almost certainly going to have to be multidimensional’. I see no reason why this would not be the case for the epistemic misconduct – and deterioration – of institutional and collective agents. Adopting explanatory pluralism just means that we cannot give a clear-cut account of the epistemic corruption of collective agents. While it might be theoretically possible to offer such an account, the considerations above suggest that it will be of limited effectiveness on its own. Not only would this honour the complexities of institutional life as we find them in the world, but it would also reflect the kind of intellectual humility that Cassam (*forthcoming*: 21) recommends as the antidote to intellectual myopia. Recognising the limits of an account is sometimes the best course of action.

## 4. Conclusion

My aim in this chapter was to consider whether epistemic corruption could occur at the collective or institutional-level. Just as we saw how the epistemic corruption of individuals turns on one’s social environment facilitating the loss or development of their epistemic virtues and vices, it was my contention that individuals could epistemically corrupt the institutions and collectives to which they belong. To vindicate this, I began by motivating the idea that collectives and institutions can possess epistemic virtues, vices, and so on. I then introduced a recent attempt at conceptualising the collective analogue of individual character, namely institutional ethos. Using ethos as my basis, I articulated two ways in which a collective or institution might succumb to epistemic corruption: *passively*, when members of these bodies fail to implement good epistemic values in an ethos; and *actively* when members encourage the uptake and exercise of epistemically bad values in an ethos.

In the second half of the paper, I turned my attention to the analogy Fricker draws between individual character and institutional ethos. In doing so, I offered a number of reasons for thinking that her analogy only extends so far, and thus we should not assume that institutional ethos can provide us with a comprehensive account of how epistemic corruption operates at the collective or institutional level. Indeed, I went even further to suggest that such an account might even be of limited effectiveness, unless tied to other explanatory considerations. Given how varied institutional phenomena are, we often need to invoke different styles of explanation. For this reason, I argued that vice-explanations should not

have a monopoly on how we diagnose institutional epistemic misconduct.

My discussion also served as a cautionary tale about the methodology of vice epistemology. We should not assume that features of individual character will provide a straightforward corollary at the collective level. Just as equally, we should not assume that an account of epistemic corruption at the individual level will translate neatly across to institutions. Whilst it is theoretically possible to develop such an account, the upshot is that it is hard to implement in practice. In the next chapter, I turn my attention to the influence a different analogy has had on vice epistemology.

# 5.

## Competence in Vice Epistemology

### Introduction

So far in this thesis, I have identified four ways in which of epistemic corruption might operate. In chapter 2, I distinguished three of these varieties: the Aretaic, Phronetic, and Attitudinal conceptions. In the previous chapter, I considered whether epistemic corruption can be scaled up to collective and institutional agents. While there is certainly a sense in which this can occur, I ultimately cautioned against drawing too close an analogy between individual and collective character. Though an undoubtedly important feature of their deterioration, I concluded that sole reliance on vice-explanations cannot fully do justice to the epistemic corruption of collective and institutional agents. If we want to fully understand this phenomenon, any account must be attentive to a plurality of explanatory styles. An important lesson to come out of the previous chapter, then, was a methodological one: that relying on analogies only gets us so far in our theorising.

In this final chapter, I extend the same sentiment to a different analogy: the *skill-analogy*. This is the idea that virtue is (like) a practical skill. Throughout the history of philosophy, there has long been a close relationship between the concepts of “virtue” and “skill”. In the ancient Greek tradition, for example, Plato and the Stoics treated virtue as a skill (*Gorgias*, 449d-458c), while Aristotle used skill as a helpful aid in developing his account of virtue (*NE*, bk. II). We find a similar narrative in contemporary philosophy. On the one hand, a growing number of virtue ethicists have sought to revive the skill-analogy in their

theorising of moral virtue (Annas, 1995, 2003, 2011; Stichter, 2016, 2018, 2020).<sup>100</sup> On the other hand, skill is so central to the normativity of reliabilist virtue epistemology that David Horst (2022) simply refers to it as ‘skill epistemology’.<sup>101</sup>

The theme of skill is highly relevant to the earlier discussions of epistemic vice and the phenomenon of epistemic corruption. Corrupting an epistemic agent can involve eroding and displacing their epistemic skills and also instilling in them various vicious epistemic skills – a theme I return to later in this chapter. Losses of epistemic skills can also impair our capacity to effectively exercise various epistemic virtues, too. Moreover, a virtuous epistemic agent who works within a corrupting and corrupted institution will need to rely on certain epistemic and interpersonal skills to safely navigate and operate within that environment. So, the skill analogy is directly related to the themes of epistemic vices and corruption.

It is perhaps understandable why virtue theorists appeal to skill in their theorising. In a pre-theoretical sense, it seems that both “virtues” and “skills” are kinds or types of *excellences*. Just as one can be an excellent moral agent because of one’s benevolence towards others, so too can one be an excellent archer on account of one’s archery skills. The basic insight is that virtues and skills are normatively good or positive features of agents, at least with respect to particular domains of human activity. Therefore, we should expect to find that virtue theorists have drawn explicit connections between the former and the latter.

In this chapter, I am primarily concerned with how this analogy has shaped virtue epistemology and its treatment of epistemic vice.<sup>102</sup> If we accept that epistemic vices are “failings”, “flaws”, or “defects” of intellectual agency, and skills are excellences of agency, then it would seem natural to conclude that epistemic vices are inimical with skills, at least normatively-speaking. Unsurprisingly, this line of reasoning seems to have infiltrated much of virtue and vice epistemology. Where skill is discussed, either explicitly or implicitly, in relation to epistemic vice, there appears to be a conviction that the former is absent from the latter in some respect. Thus, Lani Watson (2020: 243) writes that ‘bad questioning is an intellectual failing found in the exercise of many of the intellectual vices’, while Miranda Fricker (2010: 244, my italics) contends that vice ‘mirrors virtue in requiring that the bad

---

<sup>100</sup> For contemporary variations of the skill-analogy in virtue ethics, see Bashour (2021), Dougherty (2019), and Swartwood (2013). See Bloomfield (2000) and Stichter (2018) for early application in virtue epistemology.

<sup>101</sup> Sosa (2007, 2015) previously referred to his brand of virtue reliabilism as *competence* virtue epistemology and the normativity underpinning it ‘performance normativity’. See, also, Greco, (2010, 2012) and Pritchard (2012).

<sup>102</sup> Given my interest in an account of epistemic corruption, it makes sense that I focus on how this normative tension has played out in discussions of epistemic vice. That said, if my arguments here are correct, then we can expect similar results in vice ethics, and presumably an account of moral corruption. However, I do not commit myself to this conclusion here.

action be done because of the bad motive or *deficit* of skill'. I will refer to this conviction as the *deficiency thesis*.

My aim in this chapter is to reject the deficiency thesis and demonstrate how skill warrants a central role in our theorising of epistemic vice, and particularly in an account of epistemic corruption. The plan is as follows: in section 1, I begin by motivating and clarifying what I mean by the deficiency thesis. As I reveal, there are two ways of understanding the claim that epistemic vices involve a lack or deficit of skill. In the rest of the chapter, I argue that both ways of interpreting the deficiency thesis are mistaken. In section 2, I take issue with the weaker interpretation by calling into question three assumptions that might lend it some credence. Then, in section 3, I turn to consider the stronger interpretation of the deficiency thesis. Not only do the theoretical resources underpinning this interpretation speak against it, but they ironically speak in favour of expanding our repertoire of epistemic vices, which I refer to as *vice-indexed skills*. After unpacking these skills, I argue that they play an important, yet overlooked, role in the epistemic corruption of agents. I bring the chapter to a close in section 4.

## 1. The Deficiency Thesis

I begin by establishing the *deficiency thesis*. This is the claim that vice is characterised by a lack or *deficiency* of skill. As it happens, there are two ways of interpreting this claim, both of which find uptake within the virtue-epistemological literature. Perhaps unsurprisingly, the form these interpretations take is determined by the relationship skill bears to epistemic *virtue*. Therefore, this section aims to motivate and clarify these two ways of understanding the deficiency thesis. I start in section 1.1 by drawing attention to a weaker interpretation endorsed by virtue responsibilists. Then, in section 1.2, I demonstrate how a stronger interpretation finds a home in reliabilist virtue epistemology. Finally, in section 1.3, I speculate on why the deficiency thesis has flourished in virtue theory.

### 1.1. Deficiency in Responsibilism

It is no secret that virtue epistemologists reserve an important role for skill in their theorising of epistemic virtue. What differs amongst them, however, is the precise role they afford to skill. And how they approach this, we shall see, determines the way in which they interpret

what I am calling the deficiency thesis. I shall start with the weaker of the two interpretations. The weak deficiency thesis holds that epistemic vices are characterised not by a deficit of skill in general, but rather by an absence of a particular class of skills, namely the *intellectual skills* associated with epistemic virtues. This interpretation is endorsed by a number of prominent responsibilists.

Jason Baehr (2016, 2020) has articulated a ‘four-dimensional’ view of epistemic virtue. As part of this, he identifies what he calls a *competence dimension* (2016: 92). On Baehr’s view, this dimension trades on the cognitive or intellectual skills associated with particular epistemic virtues. For instance, he considers taking up alternative cognitive standpoints to one’s own to be the skill associated with open-mindedness, noticing, and attending to important details as a skill proper to intellectual attentiveness, and asking thoughtful and insightful questions to be the characteristic competence of intellectual curiosity (2016). Accordingly, Baehr (*ibid.*: 93) claims that ‘an open-minded person is competent or skilled at one type of virtue-relevant activity’, that an attentive person is ‘skilled at a different type of activity, and the curious person at yet a different type’.

Against this backdrop, Baehr claims that ‘defective competence’ or ‘a mere *lack* of virtue-relevant skill can be sufficient for the possession of a vice’ (2020: 27-28). For this to be the case, though, he stipulates that agents must lack the competence or skill despite being properly motivated and having good judgement as to when to manifest the competence. Interestingly, he finds such a scenario to be a ‘psychological implausibility’ (*ibid.*: 27). If an agent is motivated by epistemic goods and has good judgement, he remarks, why would they not possess the skill or competence associated with that virtue? This might suggest a *prima facie* reason for thinking that ‘defective competence’ is not characteristic of epistemic vices. However, as this remark equally suggests, the defect in question is just chalked up to the agent’s upstream motivational defect, which then leads to defective competence downstream. This seems clear from Baehr’s contention that a ‘motivational deficiency underlies the sorts of *failure in judgement* and *competence* which in turn contribute to the possession of an intellectual vice’ (2020: 31, my emphasis). On Baehr’s view, then, epistemic vices do indeed involve a lack or deficit of the competences associated with their corresponding virtues, owing to a motivational defect. Thus, he subscribes to the weak interpretation of the deficiency thesis.<sup>103</sup>

---

<sup>103</sup> The story is not as straightforward as this. Though Baehr seems to endorse a weak interpretation of the deficiency thesis, he briefly considers the possibility that some epistemic vices might involve certain competences themselves (2020: 28). To my knowledge, Baehr is the only theorist to consider this point. I return to substantiate Baehr’s point in section 3.2 of this chapter.

A similar account is found in Lani Watson's work on questioning in relation to the appetitive epistemic virtues, pre-eminently curiosity and inquisitiveness (2015, 2018a, 2018b, 2020). According to Watson, the intellectual skill of good questioning is a necessary feature of virtuous inquisitiveness and at least constitutive of virtuous intellectual autonomy, attentiveness, and curiosity (2018a, 2018b, 2020). To possess this skill, Watson contends that one's questions must be asked *competently* – that is, elicit information because of the question – and elicit information that is *worthwhile* (2018a: 354-355). These respectively involve judging 'who, when, where, and how' to elicit information and appropriately judging 'what information to elicit' (2018a: 355).<sup>104</sup> Rather than elicit trivial or disvaluable information, that is, the good questioner extracts information that is 'significant and relevant given her aims and contexts' (2018a: 355). In this way, the skill of good questioning manifests in (and is sometimes partially constituted by) one's epistemic virtues.

Again, though, Watson claims that the correlative failing of good questioning – *bad questioning* – is 'found in the exercise of many of the intellectual vices', that it will be a feature of 'many, if not all, the intellectual vices', and is a 'feature of intellectually vicious behaviour' (2020: 243, 256). So, mirroring how the skill of good questioning involves competently eliciting worthwhile information, she describes bad questioning as an intellectual failing that 'prevents or impedes the elicitation of worthwhile information' because the questioner either 'asks the wrong thing, or in the wrong way, or at the wrong time, or place, or the wrong source' (2020: 243). Although Watson is clear that bad questioning is neither necessary nor sufficient to possess an epistemic vice, she is clear that it will often function as a 'mechanism' or 'expression' of such behaviour (2020: 256). In other words, this failing manifests in (and is sometimes partially constituted by) one's epistemic vices.

In fact, we find traces of this interpretation of the deficiency thesis in Linda Zagzebski's (1996) *Virtues of the Mind*. According to Zagzebski, epistemic virtues require 'reliable success' in carrying out the motivation for knowledge or 'cognitive contact with reality' (1996: 136-137). To be reliably successful, she is clear that the virtuous agent will have recourse to develop certain skills. As she puts it:

Effectiveness in action requires skills, and to the extent that a virtuous person is motivated to produce external consequences desirable from the point of view of the

---

<sup>104</sup> Watson emphasises that one can ask can act competently in asking a question but still fail to elicit information. This does not detract from the goodness of the question, since good questioning does not always require success in actually acquiring information (2018: 355, 2020: 240).

virtue, he would also be *motivated to acquire the skills* that are associated with such effectiveness in action (1996: 115, my italics).

To reliably succeed in the sense relevant to the intellectual virtues, an agent will be motivated to develop and exercise distinctly intellectual skills. Zagzebski does not say a great deal about the precise nature of these skills, describing them generically as ‘techniques needed for effectiveness in the pursuit of knowledge’ and as ‘sets of truth-conducive procedures’ (1996: 116, 177). Importantly, though, she claims that their effectiveness will depend on virtuous epistemic motivations: that ‘intellectually virtuous motivations lead to the development of particular *skills* suited to the acquisition of knowledge’. In turn, these skills ‘serve virtues’ by enabling virtuously motivated agents to be successful in their epistemic affairs (1996: 116, 237).

Unsurprisingly, Zagzebski contends that vicious epistemic motivations hinder the development of these skills. This is apparent from her view that epistemic vices involve a ‘lack of motivation for knowledge’ (1996: 207), and that a motivation for knowledge is what ‘leads people to follow rules and belief-forming procedures that are truth-conducive’ (1996: 167). Considering this, she claims that intellectual vices lead to ‘unreliable belief-forming procedures’ (1996: 207). Recall, though, that she is thinking of intellectual skills broadly as ‘sets of truth-conducive *procedures*’ and as ‘techniques needed for effectively pursuing knowledge’ (1996: 177). However, if intellectual vices involve a lack of motivation for knowledge, there is reason to think this motivational deficiency translates to a corresponding deficiency of the skills needed for reliable success. This seems clear from Zagzebski’s analogy of somebody with what she calls the ‘anti-skill’ of ‘playing a musical instrument very badly’ (1996: 112, fn. 21). That the person plays the instrument very badly implies that they can play, but that the corresponding quality is poor. Rather than involve intellectual skills, then, epistemic vices are understood to involve ‘anti-skills’ in this domain.

## 1.2 Deficiency in Reliabilism

In the previous section, I identified a weak interpretation of the deficiency thesis, according to which epistemic vices involve a lack or deficit of the *intellectual skills* associated with epistemic virtues. Let me now turn to a broader and stronger interpretation of the deficiency thesis. It is broader in the sense that it extends beyond a subset of skills or competences associated with epistemic virtues; instead, it takes epistemic virtues to *be* skills or competences. In turn, it is stronger because epistemic vices are thereby characterised by a

*conceptual* deficit of skill or ability. This interpretation is very much implicit within reliabilist virtue epistemology.

Virtue reliabilism, recall, is the broad set of views that takes epistemic virtues to be any reliably truth-conducive quality of an agent. While this means that traits like intellectual curiosity or attentiveness technically satisfy the requirements for a reliabilist virtue, the reality is that reliabilists tend to restrict their focus to an agent's cognitive faculties. Accordingly, the qualities of interest to them are things like good eyesight, memory, reasoning skills, etc. This, however, is a very superficial understanding of the virtue reliabilist project. As I noted in the introduction to this thesis, the reason why reliabilists tend to focus on cognitive faculties is because they are very much in the business of providing an account of knowledge.<sup>105</sup> Central to this is their shared epistemic normativity, which starts by thinking of belief-formation as a kind of cognitive performance. Just as we evaluate non-epistemic performances by reference to a person's particular competences or abilities, virtue reliabilists conceive of epistemic virtues *as* cognitive abilities or epistemic competences by which we can judge a person's epistemic performances. The resulting view is that knowledge is true belief formed *because of* cognitive ability, competence, or virtue.<sup>106</sup>

Ernest Sosa is the architect of this 'performance normativity' about knowledge (2007, 2015).<sup>107</sup> According to Sosa, an epistemic competence is a 'disposition to succeed', such that in appropriately normal conditions the exercise of the competence would manifest (or make highly likely) the success of any cognitive performance issued by it (2007: 29, 2015: 96). However, not all dispositions to succeed constitute an epistemic competence. Just as non-epistemic competencies require certain conditions to manifest, Sosa stipulates that a disposition must satisfy his SSS model to count as an epistemic competence. For starters, a competence must be an 'innermost' or *Seated* feature of one's agency, such as one's perceptual faculty. Second, an agent must be in proper *Shape* to exercise their innermost competence. If I am intoxicated, I am not in the right shape to exercise any innermost perceptual competence

---

<sup>105</sup> This is not to say that responsibilist character virtues cannot provide an account of knowledge, just that the prospects of such an account are slimmer than reliabilist ones. For discussion, see Baehr (2011) and Sosa (2015, 2021).

<sup>106</sup> Technically, this is the central thesis of *robust* virtue reliabilism, which views knowledge solely in terms of cognitive ability or competence. This contrasts with 'modest' virtue reliabilism, which stipulates an additional safety condition on knowledge. For important contributions to the former, see Broncano-Berrocal, (2017), Carter (2016), Greco (2010, 2020), Nevarro (2015), Sosa (2007, 2011, 2015), and Turri (2015). See Kelp (2013), Kallestrup and Pritchard (2014), and Pritchard (2012, 2017, 2020) for contributions to the latter.

<sup>107</sup> More recently, Sosa (2021) refers to this as 'telic' virtue epistemology, but the focus on competence remains. This is clear from his commitment to his AAA model of competence, which he uses to analyse knowledge and judgement.

because I cannot reliably see well. This prevents the realisation of what Sosa calls a ‘fuller innermost competence’. Finally, the agent must be correctly *Situated*, that is, they must be awake, not underwater etc. (2015). If an agent satisfies these conditions, they will possess a ‘complete competence’ (2015: 99-100), which will make it ‘highly likely’ that their epistemic competence disposes them to form true beliefs.

While an agent needs to be in the correct shape and situation for a competent performance, their performance must importantly manifest their pertinent skill. Here, Sosa turns to his AAA model, which we can illustrate through his example of archery. An archer’s shot is *accurate* insofar as it hits the bullseye and *adroit* if it manifests the archer’s skill or competence. Shots can be accurate without being adroit, as when I aim for the bullseye and a strong gust of wind blows the arrow into the target; likewise, they can be adroit without being accurate, such as when my shot manifests skill but strong winds blow it off-course. For Sosa, a competent performance is one that is not only accurate and adroit, but accurate *because* adroit. This is an *apt* performance (2007, 2015). In the epistemic domain, an apt belief is one that is true (accurate) because of epistemic competence (adroit). This amounts to what Sosa terms ‘animal knowledge’ (2007, 2015).<sup>108</sup>

John Greco (2010, 2020) has long championed a similar approach to Sosa.<sup>109</sup> In the opening pages of *Achieving Knowledge* (2010: 3), he explains its central thesis is that ‘knowledge is an instance of a more general normative phenomenon – that of success through ability (or success through excellences, or success through virtue)’. On his view, cognitive abilities turn on an agent’s disposition to reliably achieve true beliefs, just as non-epistemic abilities reflect one’s success at a given task. In this way, knowledge is an achievement because of an agent’s cognitive ability. Hence, Greco’s rationale for settling on the term ‘achievement-theoretic’ epistemology (2010, 2020).

An important difference between Greco’s and Sosa’s views is that the former opts to relativise cognitive abilities to certain environments, whereas the latter does not. When considering abilities, he notes, we regularly index them to parameters relevant to the activities. Thus, we might acknowledge that a professional footballer has the ability to score penalties relative to an environment in which there is a flat pitch, reasonable weather conditions, and so on. If the footballer were blindfolded or placed on a pitch with waist-high grass, this would hardly count against the footballer’s ability to reliably score penalties.

---

<sup>108</sup> In addition to animal knowledge, there is what Sosa refers to as reflective knowledge or knowledge full well, which is apt belief that is aptly endorsed by a meta-competence (2007, 2011, 2015).

<sup>109</sup> Though similar, the two diverge on a number of issues. Most notably, Sosa understands the attribution relation *because of* in terms of competence *manifesting* knowledge, whereas Greco takes ability to *explain* knowledge.

Similarly, Greco claims that we do not lose our perceptual abilities because we find ourselves in a dark room; rather, a dark room is the wrong sort of environment to epistemically evaluate one's reliable vision (2020). More formally, he states that S has a cognitive ability A just where S is 'reliably successful, when in appropriate conditions – correct shape (Sh) and situation (Si) – within that modal environment' (2020: 130). By *reliably successful*, Greco means that an agent not only has a 'disposition seated in inner seat (Se) to believe truths in an appropriate range of propositions R', but that they also achieve this 'throughout that modal space' – i.e., they are reliably disposed to form true beliefs in the close possible worlds *relative* to the range of propositions R in question (2020: 131, 133).

At the heart of the virtue-realist project, then, is a commitment to viewing epistemic virtues *as* reliably truth-conducive cognitive abilities or competences. Given their interests in providing an epistemic normativity of knowledge, it is no surprise that reliabilists say little about epistemic vices. However, what they do have to say suggests a mirror image of the discussion above. In the early days of virtue reliabilism, for example, Alvin Goldman (1992: 160-164) observed how guessing, wishful thinking, and ignoring contrary evidence are 'standardly regarded as intellectual *vices*' because they 'produce a low ratio of true beliefs'. John Greco (2000: 67) distinguishes 'cognitive abilities such as vision, memory, and reliable reasoning from *non-abilities* such as dreaming, wishful thinking, and hasty generalisation', noting that beliefs formed via these non-abilities are 'unreliable' and not epistemically justified (2009: 292). Moreover, Sosa (2021: 20) notes how an excess of luck 'aligns with *deficiency of corresponding competence*'. True beliefs formed through guesswork are not only luckily true but the product of a reliabilist vice. It is this luck, Sosa says, which 'reduces or blocks relevant credit to the agent for the success of their attempt' (*ibid.*: 20).

As these various remarks suggest, reliabilist vices are those qualities of agents which reliably block the truth or reliably produce false beliefs. Again, though, this is a superficial understanding. Insofar as reliabilists are willing to conceive of epistemic virtues on the model of skill or competence, they should presumably make the same move with regard to epistemic vices. In other words, just as a virtuous performance is one that manifests the agent's epistemic abilities or competence in arriving at their true belief, so a vicious performance would be one that either manifests the agent's epistemic *incompetence* or *inability*, either because the relevant quality fails to reliably afford them true beliefs, or because it reliably leads them to form false beliefs. Conceptually-speaking, then, epistemic vices would be characterised by a lack or deficit of competence or skill on the part of the epistemic agent. In light of their comments above, prominent reliabilists would seem to endorse such a view.

### 1.3 The Normative Story of Vice and Skill

In the previous two sections, I have shown how two interpretations of the deficiency thesis enjoy currency in the virtue-epistemological literature. The weaker interpretation, found amongst responsibilists, characterises epistemic vices in terms of a deficit of the intellectual skills associated with virtues. The stronger interpretation, implied by reliabilists, takes epistemic vices to be cognitive incompetencies or inability. Ultimately, I will argue that both interpretations are wrong. Before I get to that, however, I would like to briefly consider why it is that virtue epistemologists (indeed, virtue theorists more generally) have subscribed to some form of the deficiency thesis.

To get a grip on this question, it is worth addressing a prior consideration: why have the concepts of virtue and skill enjoyed such a long history together? A helpful place to start is by focusing on a set of views that not only takes the two concepts to be closely related but arguably takes them to be one and the same. This is the so-called *skill-analogy*, which I mentioned at the beginning of this chapter. The skill-analogy takes as its starting point the idea that a moral virtue is a kind of practical skill.<sup>110</sup> This claim can be understood in at least two ways. First, a number of theorists claim that, in almost all respects, virtues and skills are structurally analogous such that the former count as the latter (*Gorgias* 449d-458c; Stichter, 2018, 2020; Swartwood, 2013). On this way of understanding things, the skill-analogy is no analogy at all; rather, it takes virtues to actually be practical skills. The second way in which theorists tend to understand the skill-analogy is by claiming that virtue is insufficiently analogous to skill in these respects but rather is *like* a skill (Annas, 1995, 2011). Following Matt Dougherty (2019), we can refer to the former set of views as the *Virtue as Skill* thesis, and the latter as the *Virtue as like a Skill* thesis.

Matt Stichter (2018, 2020) is a proponent of the *Virtue as Skill* thesis. On his view, moral agency is a more general form of human agency. An important feature of human agency, he notes, is self-regulation. Insofar as moral agency is a form of human agency, he claims, then self-regulation will also be vital to this domain of our lives. For Stichter, successful moral self-regulation not only includes the specific ethical standards we adopt, but also the strategies that make us better or worse at implementing these standards (2018: 59). Much

---

<sup>110</sup> Although the analogy is most often associated with the Ancient Greek tradition, it also finds uptake in classical Chinese and also Roman philosophy. See, for example, Stalnaker (2010), Yao (2012) and Klein (2014).

like performances in general, that is, our moral self-regulation can be better or worse. Just as skill improves the former, so too can it improve the latter. In this way, Stichter conceives of virtues as skills for moral (and epistemic) self-regulation.

Julia Annas (1995, 2011) has defended something akin to the Virtue as like a Skill thesis. According to Annas (2011), there are two features of virtue that make it like a skill: the ‘need to learn’ and what she calls the ‘drive to aspire’ (2011: 16).<sup>111</sup> The need to learn is rather self-explanatory and uncontroversial, namely, that virtues and skills both require time and effort to acquire. To cultivate a virtue, one needs practice; to develop a skill, one needs to build up expertise through practising the relevant ability. The ‘drive to aspire’ is what brings Annas’ account into the skill-theoretical fold. On her account, moral virtue involves coming to understand what one is doing in their ethical conduct, directing oneself in accordance with that understanding, and striving to continually improve one’s level of virtue (*ibid.*: 17, 25). In this way, Annas takes the development and exercise of moral virtue to be like that of learning and refining a practical skill for living well.

At their core, both versions of the skill-analogy trade on the more general idea that living well is like another, more familiar activity – functioning well. At least within an Ancient Greek framework, this function is to flourish or achieve *eudaimonia*. Thus, the skill-analogy compares moral virtue to a practical skill in ethical activity; in exercising a virtue, that is, a person exercises a skill that enables them to flourish. From the very start, then, the skill-analogy has presumed that the skills associated with virtues are necessarily positive. Insofar as a virtue is conducive to, or constitutive of, a flourishing life, it not only takes on the form of an excellence but one that is decidedly skill-like. It is against this backdrop that the language of vice comes to derive its connotations with “defect”, “failing”, or “flaw”. Just as virtues are conducive to a flourishing life, vices encapsulate a state that detracts from this possibility. But if we are willing to understand virtues as either analogous with, or at least closely related to, skills for flourishing, then vices turn out to be normatively juxtaposed to skills. In other words, combining the two would be akin to describing the vicious as ‘skilled’ at ethical activity – an oxymoron at the least and a wholly unpalatable conclusion from the point of virtue ethics.

On the virtue-epistemological side of things, we observe a similar state of affairs. Although responsibilists do not explicitly frame epistemic virtues in terms of flourishing, they

---

<sup>111</sup> Annas (2011: 17) stresses that only *some* skills require a drive to aspire, which is why she then treats virtues like skills, not as skills.

clearly have it in mind.<sup>112</sup> For example, Roberts and Woods (2007: 158) claim that a 'love of knowledge' is not only a pervasive feature of all the epistemic virtues, but that this love 'properly resides within the concern for what Aristotle calls *eudaimonia*'. As a result, they claim that those possessing epistemic virtues will choose 'to eschew high-level but *eudaimonistically* unfitting knowledge' (2007: 178).<sup>113</sup> Moreover, Linda Zagzebski (1996: 167) remarks that the excellence of an epistemic virtue might not be rooted in an ultimate motivation for cognitive contact with reality *per se*, but rather in a more fundamental motivation such as a 'love of being in general'. If responsibilist virtues guide our conduct in a way that is conducive to a flourishing intellectual life, and responsibilist vices detract from this prospect, then we are faced with a familiar axiological clash: to say that epistemic vices involve skills becomes akin to describing somebody as 'skilled' at epistemic activity.

This is particularly evident in virtue reliabilism. In formulating his earlier account of knowledge, Sosa explicitly endorsed the sense of virtue handed down by Plato and the Stoics, both of whom were advocates of the skill-analogy. Thus, Sosa observed that (1991: 271)

[A]nything with a function – natural or artificial – does have virtues. The eye does, after all, have its virtues, and so does a knife. And if we include grasping the truth about one's environment among the proper ends of a human being, then the faculty of sight would seem in a broad sense a virtue in human beings; and if grasping the truth is an intellectual matter then that virtue is also in a straightforward sense an intellectual virtue.

Insofar as knowledge is conducive to a flourishing life, or 'among the proper ends of a human being', as Sosa puts it, and knowledge is the result of cognitive ability or epistemic competence, then reliabilist virtues are central to that pursuit. By contrast, if reliabilist vices reliably produce false beliefs or a lack of true beliefs, then they not only detract from a flourishing intellectual life, but they do so because they involve a lack of ability or competence on the part of the epistemic agent. To say otherwise would be tantamount to saying that epistemic vices involve cognitive abilities or competences. Thus, we arrive at the same axiological clash as before.

I believe the normative tension underpinning the deficiency thesis is symptomatic of a broader assumption within virtue theory, which Charlie Crerar (2018) terms the *inversion thesis*. This is an ontological and methodological assumption about the nature of virtue and

---

<sup>112</sup> Brogaard (2014a, 2014b) is an outlier here. She argues that intellectual flourishing should be the fundamental epistemic norm towards which epistemic virtues aim.

<sup>113</sup> Roberts and Woods (2007: ch. 6) take this 'love of knowledge' to be an epistemic virtue itself.

vice, according to which the two are mirror images of each other such that they are characterised by the same, if not opposing, features (2018: 764). This assumption helps to explain why virtue epistemologists endorse some form of the deficiency thesis. After all, if epistemic virtues are characterised by the presence of skills, and vice is the mirror image of virtue, then vices will be characterised by a lack or deficit thereof. This line of reasoning seems to be precisely what underpins Fricker's claim from above. As she notes, 'it is a requirement of virtue that the good conduct is performed *because of* the good motive or skill' (2010: 244). From this, she concludes that 'the same point applies to vice, which mirrors virtue in requiring that the bad action be done because of the bad motive or deficit of skill' (*ibid.*: 244). As Crerar warns, though, the picture is far messier than this; we should not expect an account of vice to fall out an account of virtue. Despite the trivial sense in which virtue and vice are contraries, the two can come apart in interesting ways (2018: 764). In the rest of this chapter, I will argue that the relationship between epistemic virtue, vice, and skill is just one of those ways.

## 2. The Skilfully Vicious I: Weak Deficiency

The discussion in the previous section helped crystalise what I am referring to as the deficiency thesis. Specifically, it elucidated two ways of understanding the claim that epistemic vices might be characterised by a lack or deficit of skill. With this clearer, I now turn my attention to challenging both interpretations. In this section, I shall focus on the weaker interpretation found amongst responsibilists, namely that epistemic vices are characterised by a lack or deficit of the *intellectual skills* associated with virtues. I argue that even this narrower, more intuitive interpretation of the deficiency thesis rests on some problematic assumptions.

In section 2.1, I introduce a problematic tendency within virtue theory to use the term "vice" as a mass noun. This leads me to a theme we first encountered in chapter 1, a presumption that epistemic vices are global traits of character. In section 2.2, I claim that the weaker interpretation of the deficiency thesis only holds if we make this assumption about the ontology of vices. Once we abandon it, though, the weak deficiency thesis loses its initial plausibility. This is further eroded in section 2.3, where I argue that the epistemic motivations underpinning vices can be normatively flexible depending on how epistemic goods configure in one's overall ends. Together, this reveals that epistemic vices can sometimes be characterised by the intellectual skills associated with epistemic virtues.

## 2.1 A Linguistic Issue

What grounds, if any, might support thinking that epistemic vices should involve a deficiency of the intellectual skills that characterise epistemic virtues? Labels of “failings” and “excellences” aside, an initial suggestion is found by reflecting on what Ian James Kidd (2016a) terms the practice of ‘vice-charging’. As I noted in the previous chapter, vice-charging is one function of vice-attributions, where the intention is to ‘charge’ or criticise another person’s intellectual conduct on account of their supposed epistemic vices. Behind vice-explanations, vice-charging is arguably the most dominant function of vice-attributions.

There are at least two ways of engaging in the practice of vice-charging. The first is by using “vice” as a *count noun*. In linguistics, a countable noun is one which exists individually and hence can be ‘counted’, so to say. As this suggests, they can be modified by a quantifier and exist as both singular and plural qualities. Thus, to charge somebody with “vice” in the count noun sense would involve pointing to particular aspects of their intellectual character for which they might be blameworthy or criticisable. For example, if I were to criticise a person for their close-mindedness, my charge “you’re being close-minded” would be directed at a particular domain of their intellectual conduct.

Count nouns contrast with *mass nouns*. A mass noun does not follow an indefinite article such as “a” or “the” and instead corresponds to properties or qualities in a more general sense. Examples include “looking for *evidence*”, “lack of *diligence*”, or “paragon of *virtue*”. Just as we can charge people with the count noun sense of “vice”, so too can we charge them with the mass noun sense. What differs this time, though, is the scope of our criticism. On this kind of vice-charge, one’s criticism is not confined to a particular aspect or domain of a person’s intellectual life; rather, it is intended as a generalised (negative) evaluation of that person’s intellectual conduct *simpliciter*. In accusing somebody of close-mindedness, we would now be chastising them for something that characterises them *en masse*.

I raise this linguistic distinction because virtue epistemologists – and virtue theorists more generally – have tended to invoke the mass noun sense of vice in their theorising, by which I mean they have used “vice” as a means of characterising a person’s *general* intellectual conduct. For example, Alessandra Tanesini (2021: 21) states that ‘virtues contribute to being a good epistemic agent’, before claiming that ‘vices are characteristic of agents whose *intellectual life does not go well*’ (my emphasis). One way of interpreting this claim is to say that a person who has intellectual vices conducts their epistemic conduct in a worse manner than

somebody who lacks those vices. I think this is relatively uncontroversial. However, even if vices make a person's epistemic conduct worse in this respect, it does not follow that those vices thereby shape one's cognitive life *wholesale*. In a similar vein, José Medina (2013: 30-31, my emphasis) conceives of epistemic vices as 'attitudinal structures that *permeate one's entire cognitive life*'. While some epistemic vices might be capable of this; very much will depend on the agent's character and life.

Indeed, we find this sense of vice in Kidd's paper on vice-charging (2016: 186), when he notes how 'there is a lot of *vice* out there and we often judge people to evince a vicious character *or* epistemic pathology', where 'a lot of vice' is intended to capture a person's general epistemic conduct. More generally, Jason Baehr (2011: 18) observes that 'an intellectually virtuous *person* is one who thinks, reasons, judges, interprets, evaluates ... in an intellectually appropriate or rational way, while an intellectually vicious person is one who is *deficient or defective* in this regard'. Again, the 'intellectually vicious person' is generically described in a way intended to capture their deficient epistemic conduct.<sup>114</sup>

In fact, this conflation is arguably clearer in discussions of "virtue". Take what Robert Merrihew Adams (2006: 32) has called 'capital V virtue', which he says, 'normally has no plural and does not take an indefinite article'. This is precisely what a mass noun is, and his use of "virtue" is expressive of this. In keeping with the linguistic story above, Adam describes capital V virtue as the 'holistic property of having a good moral character' (*ibid*: 32), which is an over-arching feature of the agent in possession of this kind of virtue. By extension, we can assume that "capital V vice" would take a similar form and be the *holistic* property of having a bad moral character. Such a character would be the result of "vice" permeating one's life.

Linda Zagzebski (1996: 94, 105) adds to this impression when she claims that a judge is 'closer to being a *virtuous person* with compassion, and that 'he *cannot be a virtuous person* without compassion'. Here, it is as though the quality of the judge's wider character cannot be detached from the individual (countable) virtue of compassion. Likewise, Roberts and Woods (2007: 39, 69) frequently identify a person's intellectual virtues with being a 'virtuous agent', part of a 'virtuous personality', and the 'virtuous life' more generally, as if one's individual virtues can be agglomerated into one. If we treat an agent's epistemic vices as a generalised or *mass* feature of their intellectual conduct, then it would make perfect sense to think that they would lack the characteristic competence or skill associated with epistemic virtues. There would be no room at the inn for that virtue, as it were.

---

<sup>114</sup> For similar confluents, see Crerar (2018: 757-758), Watson (2020: 244); Hursthouse (1999), and Wright et. al. (2020).

## 2.2 Global vs Local Vice

Of course, there is a crucial assumption running alongside the mass noun sense of epistemic vice above: that epistemic vices are *global* properties of agents. As I noted in chapter 1, an intellectual character trait is global in scope if it extends across every domain of one's epistemic life. To say that a person possesses the vice of close-mindedness would be to say that their close-mindedness permeates all their epistemic activity; likewise, somebody who is open-minded would be said to exhibit open-mindedness throughout their intellectual conduct. In Guy Axtell's words, such a trait would be 'robustly held such that it resists undermining, and as one so settled or habitual that the agent will manifest it not just in a few situations that invite it, but in many' (2010: 74).

Notice, though, that this presumption of global vice offers responsibilists an intuitive rationale on which to premise the weaker interpretation of the deficiency thesis under consideration. For if they are willing to accept that a person's epistemic virtues are "global" or holistic in the sense above, then not only would it follow that they think, feel, and act in virtuous ways across the board; as virtuous agents, they would also develop and exercise the intellectual skills or characteristic competences associated with that virtue across their entire cognitive affairs. If this is how we should treat epistemic virtues, then it makes sense to conceive of epistemic vices as "global" properties. Not only would a person in possession of an epistemic vice lack the contrary virtue, then, but they would also lack the corresponding competence or skill associated with that virtue. Thus, we arrive at the weaker interpretation of the deficiency thesis.

As I argued in chapter 1, however, this is a questionable theoretical commitment to hold. At best, it is increasingly acknowledged that global virtue and vice are extremely difficult to cultivate; at worst, it is empirically-unfounded. Even if we accept that some of the empirical evidence from situationist social psychology has failed to replicate (Alfano, 2018), this does not re-light the flame of global virtue and vice by any means. Given the interest that the initial situationist literature sparked, virtue theorists cannot rest content that further experiments and evidence will not come to light undermining the connection between a person's conduct and situational influences. If my ontology of character from chapter 1 is correct, in fact, then we should not expect people to develop global virtues or vices at all; as the product of summary attitudes, our character traits – and *ex hypothesi* our epistemic virtues and vices – will not only come in degrees but their strength will usually depend on our

interests.<sup>115</sup> Just because a person is intellectually persevering and flexible when it comes to solving their 1000-piece puzzle, it would not follow that they exhibit this same trait to fix a difficult fault with their car. In short, our virtues and vices are likely to be “local” in scope.<sup>116</sup>

Furthermore, if we are willing to accept that our vices are local qualities, why should we think that an agent could not develop or exercise the competence or skill associated with a particular epistemic virtue whilst also possessing its contrary vice? For example, suppose a person possesses the local virtue of open-mindedness. If the skill or competence associated with this open-mindedness is taking up alternative cognitive standpoints to one’s own (Baehr, 2016), then not only would it be possible for such a person to develop and exercise this skill in particular domains; they could also possess it while also possessing the local vice of close-mindedness in other domains of their intellectual life. In such a scenario, the agent does not “forget” how to exercise this competence when in a domain of their life that they are locally close-minded. More likely is that they still retain this skill, it is just not exercised in this domain. This, however, is very different from saying that they *lack* the skill of taking up alternative cognitive standpoints to one’s own because of their close-mindedness.

What’s more, the presumption of global virtue and vice seems to carry an additional cost. On the one hand, if an epistemic virtue or vice is in fact a global feature of a person’s intellectual conduct, then there is a sense in which it is too difficult to be “virtuous” or “vicious” in the relevant sense. For if these qualities require levels of cross-situational consistency that have been brought into question by empirical data, then we will struggle to manifest a given virtue or vice in this robust way. Indeed, we need only reflect on the fact that our intellectual exemplars invariably have their epistemic blind spots – areas in which their characteristic virtues are lacking. On the other hand, if we conceive of virtues and vices as global properties of agents, then we are seemingly able to explain a person’s *general* intellectual conduct by reference to those qualities. By cultivating the virtue of open-mindedness, for example, it would follow that the agent manifests this virtue across their entire epistemic affairs. But this seems to now conflict with the plausible consideration above, that even our intellectual exemplars let slip their characteristic virtues. There would thus appear to be a tension inherent to the idea of global virtue and vice.

Perhaps most importantly, if a key rationale for the weak deficiency thesis rests on a

---

<sup>115</sup> Note, this might not extend to so-called ‘professional’ virtues and vices. See Cassam (2017).

<sup>116</sup> When discussing close-mindedness, for example, Heather Battaly (2018a: 23, my italics) adds the following qualification: ‘In short, Paul is close-minded, *at least when it comes to this issue*’. That is, Paul’s close-mindedness is not – and should not – be taken as a generalised feature of his intellectual conduct, but rather a localised feature of his character.

presumption of global virtue and vice, then this interpretation begins to lose its initial motivation. After all, it is against the backdrop of epistemic *virtues* and intellectual skills that this interpretation is premised. But if our epistemic virtues do not possess the kind of global robustness assumed above, then it is unclear why we should uncritically correlate virtues with intellectual skills. Possessing the former might correspond with the latter, but on this basis alone we should not presume that one kind of excellence (virtue) will automatically translate to another kind (skill). Indeed, Watson seems to recognise this when she remarks that the skill of good questioning is difficult to acquire. If this is the case, we should pause before correlating the skill with epistemic virtues. Together, these points cast initial doubt on the weak interpretation of the deficiency thesis.

## 2.3 The Flexibility of Motivations

At this point, there is a final line of defence supporting the weaker interpretation of the deficiency thesis. One might argue that we are justified in correlating epistemic virtues with intellectual skills precisely because they involve a final or intrinsically valuable motivation for epistemic goods. If we think that intellectual skills are an aid in acquiring, transmitting, and retaining a range of these goods, then surely being motivated by them for their own sake will entail the development of these skills. This point seems to underscore Roberts and Woods' (2007: 141) claim that 'where intrinsic motivation is low, performance quality will also be low', as well as Tanesini's (2021: 28) contention that 'the development of one's intellectual skills consists, at least partly, in the *love* or desire to acquire knowledge and understanding in the relevant domain'. Since epistemic vices lack this motivational structure, however, it is difficult to see how they might be compatible with the kind of intellectual skills associated with virtues.

This is an intuitive way of grounding the weaker interpretation of the deficiency thesis. However, there are good reasons to be cautious about it. For argument's sake, suppose a "love" or intrinsically good motivation towards epistemic goods does result in the development of intellectual skills. Nevertheless, it seems that such a motivational requirement is hardly necessary.<sup>117</sup> The reality is that many people are not motivated by epistemic goods

---

<sup>117</sup> I am willing to accept that virtue epistemologists might disagree that intrinsically good epistemic motivations are *necessary* for the development of good quality intellectual skills. Even so, the claims above leave the impression such a motivational state goes a very long way to accounting for these skills.

*for their own sake*. Often, their interests in epistemic goods are wholly derivative of their practical concerns in life: a journalist motivated to improve their questioning skills because they want to break an unfolding story; a scientist moved to consider alternative views and evidence in order to strengthen their grant application; and a graduate student intent on solving logical puzzles out of a desire to publish their paper. Strictly speaking, the agents in each case fail to exhibit a “love” or intrinsic motivation for epistemic goods, but it is unclear why this fact about them would not lead them to develop the relevant intellectual skills.<sup>118</sup> While virtuous epistemic motivations might be capable of leading agents to develop the relevant intellectual skills, they are by no means necessary. More often, it seems, this role will fall to one’s epistemic *traits*.

I worry that a pre-occupation with ultimate epistemic motivations has led virtue epistemologists to often malign another set of epistemic motivations present in virtually all intellectual traits: *proximate motivations*. These are what simply motivate an intellectually curious person to ask questions, an open-minded person to consider alternative views, and so on. Just as equally, they help to explain why an intellectually arrogant person might be motivated to feel superior to their intellectual peers, think they are in fact superior, and act in ways that shield this sense of their self (Tanesini, 2021: 107-108). What is interesting for my purposes is that these epistemic motivations are normatively flexible, which is to say that somebody can be proximately motivated to ask good questions *because* they are ultimately motivated by a desire for epistemic goods. Alternatively, they might be proximately motivated to ask good questions because they are ultimately motivated to exert intellectual dominance over their peers. How we put our proximate epistemic motivations to work will likely depend on how, and to what extent, they configure in our overall ends.

With this in mind, consider the case of ‘intellectual stinginess’. As Roberts and Woods (2007: 293) understand it, this is a disposition to ‘reserve for oneself what epistemic goods one has acquired and to be disinclined to share them with others’. On their analysis, this involves being in a ‘sort of bondage’ to these epistemic goods, ‘subject to anxiety and distress’ that they could be lost (2007: 287). For this reason, they take intellectual stinginess to be an epistemic vice that disposes one to prevent others from obtaining ‘prized’ epistemic goods. Now, imagine an academic who clings dearly to their research out of a worry that it could be lost to a certain colleague or ‘competitor’. Suppose this intellectual stinginess is the product of a harsh, competitive (and, we might add, epistemically corrupting) environment, which

---

<sup>118</sup> Charlie Crerar (2020) offers a similar assessment when criticising the ‘performance’ argument offered in favour of what he calls the ‘Strong Motivational Requirement’ on epistemic virtue.

highly prizes prolific research output, punishes the lack thereof, and also defines individual self-worth in terms of epistemic productivity. According to the weak interpretation of the deficiency thesis, our academic will lack the characteristic competence or skill characteristic of the epistemic virtue contrary to stinginess, namely intellectual generosity. Again, Roberts and Woods (2007: 287) describe this as a disposition to give valuable epistemic goods to others. More precisely, then, our academic will be said to lack the competence or skill of identifying when and how often to share those epistemic goods with others.

However, this would only be the case if our academic's intellectual stinginess is in fact robustly global. If it is not, it seems unclear why their stinginess would extend across their entire epistemic affairs. Indeed, it seems possible that our academic might be intellectually generous about the latest shopping sales or trends, letting their friends or neighbours know about the kinds of deals up for grabs. Here, we could either attribute to them the "local" virtue of intellectual generosity or think of their behaviour as instantiating *acts* of intellectual generosity.<sup>119</sup> Regardless of how we carve up their epistemic conduct, the academic is going to have recourse to develop an exercise the competence or skill associated with intellectual generosity. If they possess the local virtue, they will likely possess its characteristic competence or skill whilst also possessing the "local" vice of intellectual stinginess in their academic affairs; if they manifest acts of intellectual generosity, they will gradually acquire an ability to know how and when to gift their friends and neighbours the relevant epistemic goods.

These possibilities allow us to importantly appreciate how the academic might deploy the skill or competence associated with intellectual generosity to retain their prized epistemic goods. Rather than share their epistemic goods at the most opportune time, the academic could share them when it is inconvenient for their colleague: when they are ill, busy marking exams or preparing lectures, and so on. Alternatively, they might "gift" them valuable information that is partially or wholly tangential to their prized research.<sup>120</sup> If they engage in this activity long enough, not only will they develop the skill or competence to identify when to share some valuable epistemic goods with their colleague, but it seems in their best interest to develop this skill. Failing this, after all, would likely result in them losing their prized

---

<sup>119</sup> This need not entail that the academic possesses the *virtue* of intellectual generosity, not even the *trait*; rather, that they act in this way reflects how such an intellectually generous person *would* act. This bears parallels to Zagzebski's notion of an 'act of intellectual virtue' (1996).

<sup>120</sup> In his memoir, *The Double Helix*, James Watson recalls how Maurice Wilkens employed a tactic like this when hoarding his research on DNA from Linus Pauling. When Pauling requested research, Wilkens either downplayed its importance, said that he wished to 'look more closely at the data', or offered Pauling just enough data to buy time until he, alongside Watson and Crick, could crack the structure themselves (1968: 29).

epistemic goods. Here, the proximate motivation characteristic of virtuous intellectual generosity – being disposed to share valuable epistemic goods with others – provides our stingy academic with a means to achieve their ultimate end in preventing their colleague from acquiring particular epistemic goods. In this way, their epistemic vice is *parasitic* on the intellectual skill or competence associated with its contrary virtue, precisely because that skill feeds into the ultimate end of their vice.<sup>121</sup>

A similar case emerges in connection with those we recognise as “gossips”, those who desire to acquire ‘juicy’ information from their interlocutors (Robinson and Alfano, 2017; Kotsonis, 2021). In addition to a lack of epistemic restraint (Manson, 2011), being a gossip often involves a great deal of what Heather Battaly calls ‘epistemic self-indulgence’ (2010). This is a disposition to desire to consume and enjoy inappropriate information from others, to undertake this activity at inappropriate times, or to take pleasure in this sort of behaviour. Given this, suppose that our gossip is epistemically self-indulgent primarily when it comes to office news, such that they seek out the latest gossip about their colleagues’ romantic affairs, promotions, and so on. Why should we think that this epistemic vice would result in the intellectual failing of bad questioning? If the self-indulgent gossip is to acquire ‘juicy’ information, it seems that they will be cognisant of how best to acquire it from their interlocutors. For this reason, being a bad questioner is not an option for them.<sup>122</sup> Instead, they are likely to satisfy Watson’s two conditions on the skill of good questioning.

First, they will usually ask their questions *competently*. That is, they will tend to correctly judge *whom* to direct their questions (the person with the information); they will tend to correctly judge *when* to ask their question (away from other colleagues); they will also tend to correctly judge *where* to ask their questions (in private and away from relevant parties, rather than in front of others), and finally they are likely to correctly judge *how* they ask their questions: they will aim to drive the conversation to the ‘juicy information’ and ask questions that elicit it. But is ‘juicy’ information *worthwhile information*? Even if we think the ‘juicy’ information they seek is “epistemic rubbish”, the fact that it is ‘relevant and significant’ given their self-indulgent ends is enough to render it ‘worthwhile information’, at least according to Watson’s stipulations (2018, 2020). If this is the case, though, then a vice like epistemic

---

<sup>121</sup> This is further attested to by the existence of false modesty or humility: those who act as if they are modest or humble, when in fact they are lazy or arrogant (Hazlett, 2021). One way of appearing modest or humble is by developing and exercising a competence or skill that tracks such behaviour. For discussion, see Driver (1989, 2001) and Bommarito (2018).

<sup>122</sup> Of course, this is not to say that gossips will unfailingly be good questioners. They will certainly be candidates for bad questioning but given their ends, they have good reasons to develop this intellectual skill.

self-indulgence is unlikely to result in the intellectual failing of bad questioning. On the contrary, those so disposed will often be prime candidates to develop and exercise the intellectual *skill* of good questioning. Rather than take epistemic vices to be incompatible with the skills associated with epistemic virtues, as the weak deficiency thesis proscribes, we should first consider the extent to which the skill or competence characteristic of a given virtue might, itself, feed into the ends of a vice. Where this is a possibility, we should expect vices and skills to run together.

### 3. The Skilfully Vicious II: Strong Deficiency

My aim in the previous section was to challenge the weaker interpretation of the deficiency thesis. In doing so, I considered and rejected three grounds for thinking that epistemic vices might involve a lack or deficit of the intellectual skills associated with epistemic virtues. This revealed an important insight, namely, that epistemic vices can be parasitic on the intellectual skills that characterise epistemic virtues.

In this section, I turn my attention to the strong interpretation of the deficiency thesis, according to which epistemic vices are characterised by a deficit of skill *conceptually-speaking*. Just as epistemic virtues are viewed as cognitive abilities or competences of agents, that is, epistemic vices are taken to be inabilities or incompetencies. I argue that this claim is also mistaken. Instead of calling into question assumptions that might support this interpretation, however, I will show how the very same normativity underpinning this interpretation provides grounds for rejecting it. To help set this up, I start in section 3.1 by unpacking the role that cleverness plays in our practical reasoning. In section 3.2, I use insights about cleverness to develop an account of what I call *vice-indexed* skills, which I argue count as reliabilist epistemic vices. Then, in section 3.3, I end by noting how these skills are often a gateway to epistemic corruption.

#### 3.1 Cleverness and the Practical Syllogism

In chapter 2, I drew attention to what I called the Phronetic conception of epistemic corruption. An important feature of this conception is that it turns on conditions that erode an agent's ability to exercise the Aristotelian notion of *phronesis*, or practical wisdom. Though I did not mention it there, practical wisdom is subsumed within the broader faculty of what

Aristotle calls *cleverness* (*NE*, bk. VI, 1144a13, bk. VII, 1152a5). In this section, I will unpack the role Aristotle affords cleverness in our conduct. In doing so, I will explain why this faculty might serve as the basis for thinking that a person who possesses vices might have recourse to develop and exercise various skills.

Before I turn directly to cleverness, it will be worth briefly reconsidering the role practical wisdom plays in virtue. In my discussion of the Phronetic conception in chapter 2, I noted that Aristotle traces practical wisdom back to his separation of the soul into the rational and appetitive parts, which he further divides into the scientific (theoretical) and calculative (practical) parts (*NE*, VI, 1138b-1139a22: 138). For Aristotle, this distinction serves as the basis for two different kinds of excellence, or virtue: excellences of the appetitive part of the soul, or character virtues, and excellences in the rational part, or intellectual virtues. Though he stresses that these two virtues are fundamentally different – only the former concern pleasure and pain (*EE*, II.4.1221b-1222a) and the latter are taught as opposed to habituated – Aristotle seems to treat them as necessary and sufficient for each other.<sup>123</sup> Thus, Aristotle concludes that ‘we cannot really be good without practical wisdom, or practically wise without virtue of character’ (*NE*, bk. VI, 13, 1144b). In other words, we cannot possess character virtue without practical wisdom, which is an intellectual virtue; nor, though, can we possess this intellectual virtue without virtues of character.

Despite this symbiotic relationship, practical wisdom and character virtue clearly occupy different roles. This is borne out in Aristotle’s contention that virtue ‘makes the aim right, and practical wisdom the thing towards it’ (*NE*, VI, 1144a). Put differently, character virtue ensures that we make the correct choices about our ends (aim), whilst practical wisdom sees to it that we consistently identify the ‘correct’ means of arriving at those ends (*NE*, VI, 1142b11-b30). These two roles, one concerned with our ends and the other with the means, roughly correspond to the major and minor premises in Aristotle’s model of practical reasoning, what he calls the *practical syllogism*. Here is a classic example from the *Nicomachean Ethics* (VII.3.1147a):

MAJOR PREMISE: Everything sweet must be tasted.

MINOR PREMISE: This piece of food is sweet.

CONCLUSION: This piece of food must be tasted.

---

<sup>123</sup> References to the *Eudemian Ethics* (*EE*) are to the translation by Inwood and Woolf (2012).

As shown above, a practical syllogism is meant to depict a deductively valid argument: the major premise picks out a ‘universal’ or general principle, while the minor premise specifies ‘particulars’ of a given object or situation. The conjunction of these two premises then leads to a valid conclusion. Aristotle believed that this form of reasoning not only helped to explain why people hold certain beliefs, but why they might act in certain ways.<sup>124</sup> In practical reasoning, that is, the major premise picks out one’s end and the minor premise guides that person to act in ways that entail it. Inserting character virtue and phronesis back into the picture, we can see how the former makes the practical aim ‘right’ by determining the good end (major premise), at the same time as the latter procedurally guides one to this end in an effective way (minor premise).

Now, it is worth emphasising that the practical syllogism is not meant to be something that the virtuous agent formally follows; rather, it serves as a reconstruction of the reasons taken for their actions (McDowell, 1979: 22). In reality, the virtuous person will usually ‘grasp’ or perceive the multiple particulars of a situation that are salient to them and do so in a spontaneous manner, which will then help them determine how they ought to act in a given situation. While virtuous benevolence might generally dictate that one ought to help others in need, for example, practical wisdom will determine *how* to achieve this in particular circumstances. Thus, benevolence might require me to stop and help somebody on the street on one occasion but require me to dial 999 on another. In this way, practical wisdom is something that is dynamic and informed by virtue. This allows us to better appreciate how the erosion of epistemic *phronesis* under the Phronetic conception of epistemic corruption prevents agents from ‘perceiving’ that education is the correct place to manifest their intellectual virtues.

My interest in introducing the practical syllogism is not with practical wisdom, however, but rather with the overarching faculty of which it is a part. This is *cleverness* or *deinotes*. Whereas practical wisdom guides one’s conduct towards the right principle, cleverness is something that enables us more generally to ‘tend towards the mark we have set before ourselves, and to hit it’ (VI, 1144a18). In other words, it allows us to set our respective ends and guides the course of action we take to achieve those ends, regardless of whether they are ‘correct’. Indeed, cleverness is indifferent to whether the ends in question are good or bad. As Aristotle puts it, if ‘the aim is noble, then the cleverness is praiseworthy; if it is bad, then it is villainy. This is why both practically wise and villainous people are called

---

<sup>124</sup> Note that *akratic* people might also be capable of the kind of reasoning found in the practical syllogism yet fail to act accordingly.

clever' (VI.12.1144a).<sup>125</sup> For this reason, only virtuous people can be practically wise, in that only they can select the right means that facilitate virtuous ends.

That 'villainous' or bad people can be considered clever is the first stage in realising that the strong deficiency thesis is mistaken. For starters, it is often claimed that virtues are not like skills because the latter can also be used for bad ends, whereas the former cannot (Wallace, 1978; Putnam, 1997; Stalnaker, 2010).<sup>126</sup> Daniel Putnam (1997: 303) makes this objection vivid: excellence in bricklaying is a good skill, but 'excellent bricklaying can be done for ends which are destructive to human flourishing (concentration camp)'.<sup>127</sup> The point here is that a skilled person can employ the faculty of cleverness in their practical reasoning to put the knowledge of their craft to bad ends. Indeed, this was something Plato recognised long ago in his *Hippias Minor*.<sup>128</sup> In this dialogue, Socrates and Hippias of Elis discuss the difference between somebody who performs an activity badly in an involuntary way and somebody who voluntarily performs it badly. To their surprise, they conclude that only those who have mastered the activity in question are in a position to voluntarily perform it badly. The reason for surprise is that it follows from this that the person most capable of committing injustice is the one who excels at justice. As Socrates observes, 'it belongs to a good man to do injustice voluntarily' (*HM*, 376a).

The more general point we can extract from the *Hippias Minor* is that deliberately performing an activity badly can only be achieved by those who, themselves, excel at the activity. But those who are capable of deliberately performing bad ends will often be the most skilled. Thus, if one wants to achieve some bad end, one will have recourse to develop the skills required to do so. Here is where cleverness enters the picture. Just as practical wisdom enables one to consistently select the correct means of action, cleverness enables somebody in possession of a vice to identify the particular means of ensuring that they perform their bad end well. To do this effectively, we can now see that this means developing the skills related to the activities constitutive of those bad ends.

If this is right, cleverness offers us a partial explanation as to why our stingy academic and office-gossip from the previous section were cognisant of how the competences usually

---

<sup>125</sup> Aristotle can be read as describing the Sophists in this way. See, *Politics*, bk. VIII (2014), translation due to Rackham.

<sup>126</sup> Annas (2015) is a noteworthy outlier. She considers the prospects of vice within the skill-analogy, drawing attention to the role of cleverness. To my knowledge, Annas is the only person to consider this.

<sup>127</sup> David Roochnik (1990) interprets Socrates as making a similar point, namely that the good practitioner of a skill, unlike the virtuous person, is best-suited to use their skill for bad (*Laches*, 195c, *Republic*, 333e).

<sup>128</sup> References to the *Hippias Minor* are to the translation by R. E. Allen (1996). For an insightful discussion of this dialogue, see Callard (2022).

associated with epistemic virtues could feed into their practical ends. In the former case, the stingy academic cunningly deploys the skill associated with intellectual generosity – the skill of knowing when and how to share valuable epistemic goods – to retain their prized epistemic goods and ensure that their colleague does not easily acquire them. In the latter case, the gossip recognises that the skill of good questioning will afford them the lewd information they seek. What we see is that cleverness can help vicious agents excel at being epistemically bad.

### 3.2 Vice-Indexed Skills as Reliabilist Vices

In the previous section, I argued that cleverness can figure prominently in the exercise of an agent's epistemic vices. More specifically, it allows a person in possession of an epistemic vice to effectively identify the means of achieving their vicious end. What is often central to their success, I claimed, is the development and exercise of certain skills. Throughout this thesis, I have referred to these skills as *vice-indexed* skills. In this section, I will demonstrate how these skills can form the basis for thinking about reliabilist epistemic vices, and hence how their existence undermines the strong deficiency thesis.

On the strong deficiency thesis, recall, epistemic vices are conceived as cognitive inabilities or incompetencies. This interpretation is the logical conclusion of virtue reliabilists' decision to conceive of epistemic virtues as cognitive abilities or competences of agents. What makes these virtues such is that they are reliably-truth conducive properties of agents. But insofar as reliabilist vices are those qualities which reliably produce false beliefs or a lack of true beliefs, and truth-conduciveness is what constitutes a cognitive ability or competence, then it follows that reliabilist vices involve a lack of ability or competence.

I noted in section 1 earlier that Sosa is arguably the most prominent advocate of virtue reliabilism in contemporary epistemology. Ironically, I think his 'performance' or 'telic' normativity of epistemic virtue provides the roots for developing the notion of a vice-indexed skill. For Sosa, recall, an epistemic competence is a disposition to reliably form true beliefs when one attempts under suitable conditions.<sup>129</sup> What is interesting for my purposes is that epistemic competences are a subset of competences more broadly. As Sosa understands them, competences are simply dispositions to reliably succeed at performing a given task or aim

---

<sup>129</sup> Not all dispositions are competences: glass is disposed to reliably break upon impact, but we would not say that glass has a competence to break. See Vetter (2019), and Sosa (2007, 2015) for relevant discussion. I thank Quassim Cassam for raising this point.

(2015: 96, 2021: 45). Within his framework, ‘telic credit’ accrues to an agent just if the success of their performance is attributable to their competence. In this way, manifesting competence is what decides the normative status of one’s performance. Indeed, Sosa (2021: 21, fn. 7) remarks that ‘telic excellence allows awful aims, as in the “perfect” murder’ because all that goes into the normative evaluation of the murder is whether it is performed through skill or competence. Thus, an assassin who murders their target because of their skill not only performs competently but is a *good* assassin.

I think something similar can be said about epistemic vices. To see how, let me reflect on some remarks recently made by Jason Baehr (2020). Baehr endorses a weak interpretation of the deficiency thesis, evident from his contention that ‘a mere lack of virtue-relevant skill can be sufficient for the possession of a vice’ (*ibid.*: 28). Despite this, he speculates that there could be some epistemic vices ‘such that, to possess them (or to possess them fully), one must possess certain skills characteristic of these vices’ (*ibid.*: 28). Though brief, he notes skills associated with ‘self-deception’, ‘keeping counterevidence at bay’, and ‘misrepresenting opposing viewpoints’ – the latter two he associates with intellectual dishonesty and close-mindedness, respectively (*ibid.*: 28). In doing so, he suggests that with enough training or practice, one can become “better” – or more competent – at these mental operations’ (*ibid.*: 28). Thus, somebody who is skilled at misrepresenting other peoples’ views, he thinks, is more close-minded than somebody who is not as skilled in this activity.

Not only do I take this to be entirely right, but I think we can make sense of it by appealing to Sosa’s performance or telic normativity. Again, let us take a competence to be a disposition to reliably succeed at performing a given task or aim. If those who engage in self-deception, keeping counter-evidence at bay, or misrepresenting opposing views reliably succeed at these tasks through exercising these ‘characteristic competences’, then there is not only a sense in which they are ‘competent’ at these respective tasks; there is a further, telic sense in which they are *good* at doing so. Much like the good assassin, then, these agents are good at being epistemically bad.<sup>130</sup> This much reveals that these *vice-indexed skills* or characteristic competences are entirely compatible with the framework Sosa employs to model reliabilist virtues. But if this is the case, then notice that we can carve out conceptual space to view these kinds of skills or competences *as* reliabilist vices. Before I make this claim,

---

<sup>130</sup> A similar point might hold in the moral realm: imagine a torturer who excels at torturing his victims. There is something more morally disvaluable about a skilled torturer than one who achieves the aims of his torture through lucky means. If so, these competencies might instantiate the sort of ‘bads’ that are constitutive of ill-will, at least according to those accounts which take ill-will to involve the presence of bads as opposed to the absence of goods. For discussion, see Kagan (2014) and Woodard (2022).

though, let me clarify several things.

Within the vice-epistemological literature, the notion of a ‘reliabilist vice’ has generally been correlated with the production of bad epistemic effects, or what Battaly (2016) calls ‘effects-vices’. Quassim Cassam (2016, 2019a) has articulated a broadly ‘reliabilist’ account along these lines. On his obstructivist account, epistemic vices are character traits, attitudes, or ways of thinking that systematically obstruct the acquisition, retention, and transmission of knowledge.<sup>131</sup> Like Cassam’s view, the proposal I am suggesting would take reliabilist vices to be qualities that reliably produce bad epistemic effects, or systematically obstruct knowledge. For instance, competently misrepresenting other peoples’ viewpoints is a reliable means of acquiring false beliefs about others, whilst being competent in self-deception reliably obstructs one’s ability to acquire self-knowledge (Cassam, 2014). As this already suggests, though, my proposal does not restrict reliabilist vices to features of intellectual character. Unlike Cassam’s view, that is, the reliabilist vices I have in mind not only extend beyond character but are also attuned to the very same normative framework that underpins virtue reliabilism. What I am suggesting is that we can conceive of reliabilist vices simply as competences or skills that reliably produce bad epistemic effects.

An upshot to framing these vices within a performance or telic framework is that it allows us to track their epistemic disvalue by appealing to their status as competences. Just as Sosa derives the value of reliabilist virtues from their status as competences to reliably form true beliefs, so reliabilist vices would derive their epistemic disvalue from their status as competences which frustrate this. Indeed, we can employ Sosa’s AAA model to illustrate this. At a first approximation, this might involve assessing the following criteria: (i) whether an agent performs the bad epistemic activity or practice *accurately* – does somebody who attempts to keep counter-evidence at bay, for example, succeed at doing so? (ii) Is the bad epistemic practice or activity undertaken *adroitly* – to what extent does the agent keeping counter-evidence at bay perform this task deliberately through their competence or skill, as opposed to luck? (iii) Finally, is the resulting performance *apt* – does the agent succeed at keeping counter-evidence at bay *because of* their competence, and hence because of their reliabilist vice? Far from being characterised by a lack or deficit of skill or competence, then, we can see that reliabilist vices can plausibly be modelled on the very same basis as reliabilist virtues – that of competence. If this much is true, we should reject the strong interpretation of the deficiency thesis.

---

<sup>131</sup> Cassam develops his ‘obstructivist’ view from Julia Driver’s (2000, 2001, 2003) consequentialist account of moral and epistemic virtue.

To bring the discussion to a close, I want to reflect on a consensus that has emerged in much of the virtue-reliabilist literature. This is the claim the responsibilist character virtues are likely to play a supportive or ‘auxiliary’ role in an account of knowledge (Sosa, 2015, Greco, 2010). If my contention in this section is correct, then the inverse conclusion seems to be true. For notice that the kind of activities that make for reliabilist vices, namely keeping counter-evidence at bay, self-deception, and misrepresenting others’ views, extend from responsibilist character vices – intellectual dishonesty and closed-mindedness, etc. Thus, reliabilist vices will often amount to competences to reliably succeed at performing the activities or ends characteristic of responsibilist vices. What this would suggest, then, is that reliabilist vices are auxiliary to responsibilist vices. Where we find responsibilist vices, that is, we are likely to find reliabilist vices not too far behind. In the remainder of this chapter, I will demonstrate this possibility by returning to my account of epistemic corruption.

### 3.3 The Competence to Corrupt

In this chapter, I have sought to cast doubt on the deficiency thesis – the claim that epistemic vices are characterised by a lack or deficit of skill in some respect. Across the previous two sections, I focused my attention on the two respective ways of interpreting this claim. In both cases, my basic conclusion was the same: that epistemic vices can be characterised by skill or competence. I hope, therefore, that the arguments in this chapter have demonstrated the important role skill can play in vice epistemology. In this final section, I will add to this impression by highlighting the role that vice-indexed skills can play in an account of epistemic corruption.

We can begin approaching this task by casting our minds back to what I referred to in chapter 3 as active epistemic corruption. Amongst the many ways in which this might materialise, one particular means stands out for the present discussion. This is that one’s social environment can encourage and facilitate the development of certain skills that attach to epistemic vices, or what we now know as vice-indexed skills. If my contention above is correct, that these skills are often auxiliary to epistemic vices, then encouraging the former is likely to give rise to the latter. To illustrate this, consider the following case.

Samantha is interning at a reputable company. Knowing that this company is Samantha’s dream workplace, her boss exploits her willingness to take orders by tasking her with hiding evidence of the company’s previous malpractice. At first, Samantha refuses but her boss quickly reminds her that the internship is highly competitive, and that there are

many qualified candidates with whom he could replace her. Under threat of losing the internship, then, she begrudgingly obliges. As Samantha engages in this activity more and more, she devises increasingly elaborate ways of hiding and destroying evidence about the company's history. For this, she receives high praise from her boss. Over time, she develops her obfuscatory abilities so well that she becomes considered the 'go-to' person in the office, and other teams begin approaching her to carry out their dirty work. Eventually, this status becomes integral to her self-conception whilst interning. In light of this excellent obfuscation, Samantha is offered her dream job at the company, where it is strongly and repeatedly implied that these skills will continue to be highly relevant.<sup>132</sup>

In this case, Samantha's workplace encourages her to develop and exercise an array of skills to obfuscate evidence effectively. Now, there are two ways of proceeding from here, both of which lead us to the same conclusion. On the one hand, we might think that being skilled at obfuscating is not an epistemic vice, but rather a 'mechanism' or 'expression' of intellectual dishonesty at the least, and 'epistemic malevolence' at the worst – that is, an 'opposition to another's share of epistemic goods' and their 'epistemic well-being' (Baehr, 2010: 204). Insofar as Samantha's workplace encourages her to develop and deploy these skills – finding out when, where, and how best to obfuscate the company's dirty deeds – then her workplace equally encourages her to develop epistemically malevolent dispositions. In this case, Samantha is subjected to active epistemic corruption because her social environment encourages the development and exercise of epistemic vices, and this occurs because that environment encourages her to develop a particular vice-indexed skill.

On the other hand, the discussion from the previous section presents us with a different possibility. By encouraging Samantha to develop these obfuscatory skills, her workplace simply encourages the development of a reliabilist epistemic vice, understood again as a competence to reliably produce false beliefs or a lack of true beliefs. After all, being skilled at obfuscating evidence means that Samantha routinely misses out on the chance to form true beliefs about the company's prior conduct. However, if people are destroying evidence of institutional malpractice, we deserve to know what sort of malpractice was involved, especially if it could affect our lives. Thus, Samantha's competence in destroying this evidence deprives others from forming true beliefs about such malpractice. In both scenarios, though, we arrive at the same conclusion: encouraging the development of a vice-indexed skill

---

<sup>132</sup> Here, we might distinguish skills for (i) *doing* epistemically bad things and (ii) skills for *getting away* with those deeds. Somebody in a similar situation to Samantha might develop the latter skills because they lack the power to get away with their epistemic misdeeds. Contrast this with a powerful executive, who can openly and brazenly do bad things since nobody can stop them.

encourages one or more epistemic vices. This is one way in which reliabilist vices can be auxiliary to responsibilist character vices. We can represent this relationship as follows:

Vice-Indexed Skill/Reliabilist Vice → Responsibilist Vice.

In other cases, though, the encouragement of a responsibilist character vice will likely lead to the development of a vice-indexed skill/reliabilist vice. One case I have in mind is what Quassim Cassam (2019a) calls ‘epistemic insouciance’. This is an attitude of indifference to the truth. In his discussion of this vice, Cassam draws an explicit connection between insouciance and the activity of bullshitting (Frankfurt, 2005). As he observes, bullshit is not only the ‘primary product of epistemic insouciance’ but ‘epistemic insouciance *leads* to bullshit’ (2019: 80-81, fn. 6, my italics). Put differently, if one wants to be epistemically insouciant, one will have recourse to engage in the activity of bullshitting.<sup>133</sup>

Now imagine a scenario in which epistemic insouciance is widely encouraged. Suppose, for example, that one had the misfortune of working in the West Wing during Donald Trump’s presidency. The Washington Post’s Fact Checker has Trump on record making 30,573 false or misleading claims during his presidency – an average of twenty-one untruths each day (Kessler et. al., 2021). As this statistic reveals, a disregard for the truth was rife in the Trump administration. It is not too difficult to imagine that something like epistemic insouciance was widely encouraged in this environment. Insofar as bullshit is the ‘primary product’ of epistemic insouciance, though, this bullshitting would have equally been encouraged. Thus, the more insouciant one becomes about the truth, the more one will engage in bullshitting. And the more one bullshits, the more competent one is likely to become at this activity.

This is not to say, of course, that one must be a competent bullshitter to count as viciously insouciant. I very much acknowledge that one can be epistemically insouciant and be bad at bullshitting. However, this possibility just underscores the distinction we made in the above case. On the one hand, we can take an agent’s bullshitting skills to function as expressions or mechanisms by which they manifest the attitude-vice of epistemic insouciance. In which case, this environment has encouraged a vice-indexed skill. Alternatively, we can say that this corrupting environment encourages a reliabilist epistemic vice, namely the skill of bullshitting. A competent bullshitter, much like the skilled obfuscator, is going to reliably

---

<sup>133</sup> This is plausibly due to bullshitting being a ‘high-fidelity’ activity, in Alfano’s (2013: 31-32) terms. In order to effectively disregard the truth through bullshitting, that is, one will have to engage in this activity fairly regularly and this will often involve practicing or refining one’s bullshitting abilities. We can contrast this with a ‘low-fidelity’ activity such as riding a bicycle or swimming, which requires far less exercise to remain competent in.

produce bad epistemic effects. Specifically, they are likely to transmit a host of false beliefs to others. By encouraging the development of epistemic vices, then, actively corrupting environments can facilitate the development, entrenchment, and enrichment of vice-indexed skills or reliabilist epistemic vices. Thus:

Responsibilist Vice  $\rightarrow$  Vice-Indexed Skill/Reliabilist Vice.

I suspect the reality is far less clear-cut than these suggestions make out. It seems that in cases whereby the encouragement of vice-indexed skills leads to the development of epistemic vices, the maintenance and refinement of those skills will partially turn on the vicious dispositions encouraged. For instance, as Samantha becomes more competent or skilled at obfuscating evidence, not only will she likely develop epistemically malevolent or dishonest dispositions, but those dispositions will simultaneously nurture and sustain the development of her obfuscatory skills. A similar process is likely to occur in cases whereby the encouragement of epistemic vices leads to vice-indexed skills. Consider our skilled bullshitter again. To refine and hone their bullshitting skills, it seems plausible that the Trump official will continue to be disposed to epistemic insouciance. But the longer they are so disposed, the more engrained this vice becomes in their intellectual character. And the more engrained this attitude is in their character, the more reason they will have to deploy and exercise their bullshitting skills. What we are left with is a mutually sustaining relationship between epistemic vices and vice-indexed skills:

Responsibilist Vice  $\rightleftharpoons$  Vice-Indexed Skill/Reliabilist Vice.<sup>134</sup>

While the initial focus in each case was on the causal relationship between the respective qualities, we can begin to see that actively corrupting conditions can give rise to a symbiotic process. As conditions encourage epistemic vices, the development and exercise of those vices facilitates the development of vice-indexed skills; and as one develops and exercises these skills, the more entrenched epistemic vices become in one's character. Through this process, the relationship between vice-indexed skills/reliabilist vices and responsibilist epistemic vices can become a *constitutive* one. Part of what it means to manifest the latter, that is, can be borne out in exercising the skill or competence involved in the former. In a sense, the two sets of qualities would be inseparable.

---

<sup>134</sup> To avoid confusion, I use the  $\rightleftharpoons$  symbol not to represent a bi-conditional but rather to denote its scientific connotations, where a reaction is reversible and consequently has the potential to reach a state of equilibrium.

## 4. Conclusion

The previous three chapters in this thesis were devoted to a sustained analysis of the concept of epistemic corruption. In this chapter, I stepped back to consider a more general feature of vice epistemology, namely the way in which theorists understand the relationship between epistemic vices and skills. My reason for exploring this relationship was because of a long tradition within virtue theory that takes virtues and skills to be closely related, if not the same thing. This is the so-called skill-analogy, according to which virtues are (like) practical skills. In the first part of this chapter, I demonstrated how this analogy has, to differing degrees, influenced the ways in which virtue epistemologists understand the relationship between epistemic virtues and skills. This, I argued, has come to shape the way that vice epistemologists conceive of the relationship between epistemic vices and skills. Whereas epistemic virtues are characterised by their presence, epistemic vices are taken to involve a lack or deficit. I referred to this claim as the deficiency thesis.

In the remainder of the chapter, I argued that both the weak and strong interpretations of this claim are mistaken. On the one hand, I claimed that reliabilist character vices can be characterised by the intellectual skills that are present in epistemic virtues. On the other hand, I demonstrated that the conceptual resources of virtue reliabilism allow us to model reliabilist epistemic vices *on* competences. The upshot of my arguments here is twofold: first, that virtue and vice epistemologists are mistaken to endorse either interpretation of the deficiency thesis; and second, that there is an important role for skill or competence to play in vice epistemology going forward. One such role is within an account of epistemic corruption. Thus, I ended this chapter by considering the way in which our social environments can not only encourage the development of vice-indexed skills or reliabilist vices, but how these qualities can contribute to the epistemic corruption of agents. This completes our account of epistemic corruption.



## Concluding Remarks

My aim in this thesis was to begin filling a theoretical lacuna surrounding how our social environments can damage or corrupt our intellectual character. If there are three themes that run throughout this process, they are *character*, *corruption*, and *competence*. Together, these ingredients formed the basis of my account of epistemic corruption. I started by critically discussing the notion of character. Specifically, I argued that a prominent account of character faces a dilemma I termed the *ontological problem*. To overcome this dilemma, I sketched an ontology of intellectual character which takes our intellectual traits to be grounded in prior summary attitudes towards epistemic practices and objects.

After sketching this ontology, I proceeded to develop my account of epistemic corruption. This began in chapter 2, where I distinguished three ways of understanding the claim that our social environments can degrade our intellectual character. I referred to these as the Aretaic, Phronetic, and Attitudinal conceptions of epistemic corruption. Having distinguished these varieties of epistemic corruption, I then turned my attention in chapter 3 to the ways in which they might play out in practice. A question I left unanswered at this point was how it is that epistemically corrupting conditions might come into existence. This question was the focus of chapter 4. There, I drew on a recent attempt at conceptualising collective character, in an effort to scale up epistemic corruption to the institutions that form our social environments. Two insights came from this discussion: (i) an account of epistemic corruption can be extended to collectives and institutions in theory, but (ii) such an account must concede that character-based concepts are only so effective on their own. As such, offering a corruptionist critique of collectives and institutions will often be difficult in practice.

In the final chapter of the thesis, I took a step back from epistemic corruption and investigated the relationship between epistemic vices and skills. In particular, I took aim at what I called the *deficiency thesis*, which is an assumption that epistemic vices will be characterised by some form of absence or deficit of skills. What we found is that epistemic vices can sometimes be parasitic on the intellectual skills associated with epistemic virtues. Perhaps more striking was the conclusion that the very theoretical resources underpinning one way of interpreting the deficiency thesis actually lend support for conceiving of reliabilist epistemic vices as competences. Far from being characterised by a deficit of skill, then, I concluded that vice epistemologists would do well to pay closer attention to the role that skill

and competence can play in our accounts of epistemic vice. I demonstrated this prospect by showing how these competences can configure in my account of epistemic corruption.

While the over-arching aim of this thesis was to develop and articulate an account of epistemic corruption, the general discussion here also provides us with at least four directions for future research. The first of these is a contention explicit within Kidd's (2020) work on epistemic corruption, namely that epistemic vices 'damage' intellectual character. Though I did not explicitly address this 'characterological damage' here, it will be well worth our time as vice epistemologists to better understand what constitutes this damage. One possibility is handed down to us from ancient Greek tradition, which is that (epistemic) vices prevent us from flourishing, epistemically-speaking.<sup>135</sup> Of course, though, some might think that flourishing is an ethically-loaded concept, one that is not strictly apt for epistemology. Indeed, if the damage inflicted to character by epistemic vices is taken to involve a failure to flourish, then in what sense is intellectual character distinct from moral character? This brings to the fore what Julia Driver (2003) has called the 'conflation problem', or the problem of clearly distinguishing moral and epistemic qualities from each other. Insofar as vice epistemologists are interested in the nature and significance of epistemic vices, then addressing this problem will help them make sense of how those vices damage character in a distinctly epistemic way.

The second area of future research turns on the prospects of epistemic corruption at the collective-level. Although I expressed some doubts about this in chapter 4, my concerns there applied strictly to a corruptionist critique that models collective character on institutional ethos (Fricker, 2020). As such, it remains to be seen whether other accounts of collective character in the literature would fare any better at capturing the epistemic deterioration of institutions and collectives. One potential alternative might be de Rooij and de Bruin's (2022) account, according to which a collective possesses a character insofar as it functions as an agent. Though I cannot unpack the account here, a noticeable difference is its rejection of the joint commitment model of collective agency that underpins institutional ethos. Given that this model is responsible for several of the problems facing ethos, perhaps their account offers a more viable alternative for modelling the epistemic corruption of collective agents. Only time will tell if this is the case. Regardless, such an account will have to make important room for other, non-characterological considerations in explicating this phenomenon.

---

<sup>135</sup> Similar thoughts are found in other traditions, albeit in different forms. For the Buddha, the epistemic vices referred to as *avijja* – literally 'ignorance', but including dogmatism, resistance to truth, and wilful ignorance – are ultimately bad insofar as they entrap us within *samsara*, the cycle of *kamma* and rebirth. See, for instance, *Majjhima Nikaya* 9 in Nyanatiloka (2004).

The third area of research falls out of the discussion from chapter 5. One of the main conclusions to come out of this chapter was that vice epistemologists have been unwarranted in neglecting skill from their theorising of epistemic vice. This led me to identify what I called *vice-indexed skills*. If these skills can be a causal, if not constitutive, feature of epistemic vices, as I argued, then their existence opens up a new area of interest for vice epistemologists. In addition to studying the nature and epistemological significance of epistemic vices, it would also be incumbent on them to address some of the following questions: do all epistemic vices have vice-indexed skills? If not, which vices lack these skills? Why might they lack such skills? Does the possession or development of these skills contribute to the blame or criticism an agent garners in virtue of possessing an epistemic vice? How might we ameliorate the development of vice-indexed skills? To the extent that vice epistemologists are interested in the nature of epistemic vices, they should also be interested in the nature of these skills.

The final area of research is related to the previous one. I claimed that vice-indexed skills might not just be an important feature of epistemic vices; I further suggested that vice epistemologists have the theoretical resources to model reliabilist epistemic vices *on* these skills. This is an exciting prospect, but it generates a potential challenge. This is because skills or competences are said to involve practical knowledge or *knowledge-how* (Ryle, 1949). This is the kind of knowledge that is said to underpin our abilities to ride bicycles or drive cars. Just as these activities frequently require us to acquire practical knowledge, a similar thing can be said about those who develop their bullshitting skills or who become highly competent in obfuscating evidence. In order to become skilled or competent in these activities, that is, the agents are going to have to acquire a deal of knowledge-how, from knowing how to hide evidence in particular ways, to knowing how to recognise which audiences are likely to see through one's bullshit. However, if it turns out that these vice-indexed skills involve knowledge-how, and these skills are often partially constitutive of character vices – or indeed *count* as reliabilist vices – then in what sense do they obstruct knowledge or produce bad epistemic effects? Anyone wishing to model reliabilist vices on vice-indexed skills, then, would do well to grapple with this question, and indeed, the relationship between epistemic vices and knowledge-how more generally. By addressing these research areas, vice epistemology will undoubtedly become an ever richer discipline going forward.



# Bibliography

- Adams, R. M (2006) *A Theory of Virtue: Excellence in Being for the Good*, Oxford, Clarendon.
- Alfano, M (2012) 'Expanding the Situationist Challenge to Responsibilist Virtue Epistemology', *The Philosophical Quarterly*, 62: 247, pp. 223-249.
- (2013a) *Character as Moral Fiction*, Cambridge: Cambridge University Press.
  - (2013b) 'The Most Agreeable of All Vices: Nietzsche as Virtue Epistemologists', *British Journal for the History of Philosophy*, 21:4, pp. 767-790.
  - (2018) 'A Plague on Both Your Houses: Virtue Theory After Situationism and Repligate', *Teoria*, 38: 2, pp. 115-122.
- Alfano, M and Fairweather, A (eds.) (2017) *Epistemic Situationism*, Oxford: Oxford University Press.
- Alfano, M. R. and Robinson, B (2017) 'Gossip as a Burdened Virtue', *Ethical Theory and Moral Practice*, 20:3, pp. 473-487.
- Allport, G. W and Odbert, H. S (1936) 'Trait-names: A Psycho-lexical Study', *Psychological Monographs: General and Applied*, 47:1, pp. 1-176
- Annas, J (1995) 'Virtue as a Skill', *International Journal of Philosophical Studies*, 3:2, pp. 227-243.
- (2003) 'The Structure of Virtue', in M. dePaul and L. T. Zagzebski (eds.) *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Oxford: Oxford University Press, pp. 15-33.
  - (2011) *Intelligent Virtue*, Oxford: Oxford University Press.
  - (2015) 'Virtue, Skill, and Vice', *Etica e Politica*, 17:2, pp. 94-106.
- Anscombe, G. E. M (1958) 'Modern moral philosophy', *Philosophy*, 33:124, pp. 1-19.
- Aristotle (1998) *Nicomachean Ethics*, J. L. Ackrill and J. O. Urmson (trans.), Oxford: Oxford University Press.
- (2012) *Eudemean Ethics*, B. Inwood and R. Woolf (trans.), Cambridge: Cambridge University Press.

- (2014) *Politics*, H. Rackham (trans.), Massachusetts, Harvard University Press.

Armstrong, D. M (1980) *The Nature of Mind and Other Essays*, Queensland: University of Queensland Press.

Astell, M (2002) *A Serious Proposal to the Ladies. Parts I and II* [1694], P. Springborg (ed.), Ontario: Broadview Literary Text.

Axtell, G (2010) 'Agency Ascriptions in Ethics and Epistemology: Or, Navigating Intersections, Narrow and Broad', *Metaphilosophy*, 41:1-2, pp. 73-94.

Bacon, F (1859) *Novum Organum* [1620], Rev. A Johnson (trans.), Bell and Daldy: London.

Badwhar, N. K (1996) 'The Limited Unity of the Virtues', *Noûs*, 30:3, pp. 306-329.

Baehr, (2010) 'Epistemic Malevolence', *Metaphilosophy*, 41: 1-2, pp. 189-213.

- (2011) *The Inquiring Mind: An Essay on Intellectual Virtues and Virtue Epistemology* Oxford: Oxford University Press.
- (2013) 'Educating for Intellectual Virtues: From Theory to Practice', *Journal of Philosophy of Education*, 47:2, pp. 248-262.
- (2016) 'The Four Dimensions of an Intellectual Virtue' in C. Mi, M. Slote, and E. Sosa (eds.) *Moral and Intellectual Virtues in Western and Chinese Philosophy*, New York: Routledge, pp. 86-98.
- (2020) 'The Structure of an Intellectual Vice' in I. J. Kidd, Q, Cassam, and H. Battaly (eds.) *Vice Epistemology*, New York: Routledge, pp. 21-36.
- (2021) *Deep in Thought: A Practical Guide to Teaching for Intellectual Virtues*, Massachusetts, Harvard University Press.

Baker, H K., Kumar, S., and Goyal, N. (2021) 'Personality Traits and Investor Sentiment', *Review of Behavioural Finance*, 13: 4, pp. 354-369.

Banaji, M. R., and Heiphetz, L (2010) 'Attitudes', in S. T. Fiske, D. T. Gilbert, and G. Lindzey (eds), *Handbook of Social Psychology* (5th ed.), New Jersey: John Wiley and Sons, pp. 353-393.

Baron-Cohen, S (2011) *The Science of Evil: On Empathy and the Origins of Cruelty*, New York: Basic Books.

Bashour, B (2021) 'How Did She Get So Good? On Virtue and Skill', *Ethical Theory and Moral Practice*, 24:2, pp. 563-575.

Batson, C (2011) *Altruism in Humans*, New York: Oxford University Press.

Battaly, H (2010) 'Epistemic Self-Indulgence', *Metaphilosophy*, 41:1-2, pp. 214-234.

- (2013) 'Detecting Epistemic Vice in Higher Education Policy: Epistemic Insensibility in the Seven Solutions and the REF', *Journal of Philosophy of Education*, 47: 2, pp. 263-280.
- (2014) 'Varieties of Epistemic Vice' in J. Matheson and R. Vitz (eds.) *The Ethics of Belief*, Oxford: Oxford University Press, pp. 51-76.
- (2016) 'Epistemic Virtue and Vice: Reliabilism, Responsibilism, and Personalism' in C. Mi, M. Slote, and E. Sosa (eds.) *Moral and Intellectual Virtues in Western and Chinese Philosophy*, pp. 99-120.
- (2017) 'Testimonial Injustice, Epistemic Vice, and Vice Epistemology' in I. J. Kidd, J. Medina, and G. Pohlhaus Jr (eds.) *The Routledge Handbook of Epistemic Injustice*, New York: Routledge, pp. 223-231.
- (2018) 'Can Close-mindedness be an Intellectual Virtue?', *Royal Institute of Philosophy Supplements*, 84, pp. 23-45.
- (2018b) 'A Third Kind of Intellectual Virtue: Personalism', in H. Battaly (ed.) *The Routledge Handbook of Virtue Epistemology*, New York: Routledge, pp. 115-126.
- (2019) 'Vice Epistemology has a Responsibility Problem', *Philosophical Issues*, 29:1, pp. 24-36.

Baumeister, R. F, Smart, L., and Boden, J. M (1996) 'Relation of Threatened Egotism to Violence and Aggression: The Dark Side of High Self-Esteem', *Psychological Review*, 103:1, pp. 5-33.

Berg, M. and Seeber, B (2016) *The Slow Professor: Challenging the Culture of Speed in the Academy*. Toronto: University of Toronto Press.

Biddle, J. B., Kidd, I. J. and Leuschner, A (2017) 'Epistemic Corruption and Manufactured Doubt: The Case of Climate Science', *Public Affairs Quarterly* 31.3, pp. 165-187.

Bloomfield, P (2000) 'Virtue Epistemology and the Epistemology of Virtue', *Philosophy and Phenomenological Research*, 60:1, pp. 23-43.

Bloom, A. D (1988) *The Closing of the American Mind: how higher education has failed democracy and impoverished the souls of today's students*, New York: Simon and Schuster.

Bloomfield, P (2013) 'Error Theory and the Concept of Morality', *Metaphilosophy*, 44:4, pp. 451- 469.

Bommarito, N (2018) *Inner Virtue*, New York: Oxford University Press.

Brady, M (2014) *Emotional Insight: The Epistemic Role of Emotional Experience*, Oxford: Oxford University Press.

Bratman, M. E (2014) *Shared Agency: A Planning Theory of Acting Together*, Oxford: Oxford University Press.

Brogaard, B (2014a) 'Intellectual Flourishing as the Fundamental Epistemic Norm', in C. Littlejohn and J. Turri (eds.) *Epistemic Norms: New Essays on Action, Belief, and Assertion*, Oxford: Oxford University Press, pp. 11-31.

- (2014b) 'Towards a Eudaimonistic Virtue Epistemology', in A. Fairweather (ed.) *Virtue Epistemology Naturalised*, London: Synthese Library, pp. 83-102.

Broncano-Berrocal, F (2017) 'A Robust 'Enough' Virtue Epistemology', *Synthese*, 194, pp. 2147-2174.

Byerly, T. R and Byerly, M (2016) 'Collective Virtue', *Journal of Value Inquiry*, 50:1, pp. 33-50.

Cacioppo, J. T., and Petty, R. E. (1982) 'The Need for Cognition', *Journal of Personality and Social Psychology*, 42, pp. 116-131.

Cacioppo, J. T., Petty, R. E., Feinstein, J. A., and Jarvis, W. B. G. (1996) 'Dispositional Differences in Cognitive Motivation: The Life and Times of Individuals Varying in Need for Cognition', *Psychological Bulletin*, 119:2, pp. 197-253.

Callard, A (2022) 'Being Good at Being Bad: Plato's *Hippias Minor*', in D. Ebrey and R. Kraut (eds.) *The Cambridge Companion to Plato*, Cambridge, Cambridge University Press, pp. 146-172.

Callcut, D (2009) *Reading Bernard Williams*, Oxford: Routledge.

Card, C (1996) *The Unnatural Lottery: Character and Moral Luck*, Philadelphia: Temple University Press.

Carel, H and Kidd, I. J (2021) 'Institutional Opacity, Epistemic Vulnerability, and Institutional Testimonial Justice', *International Journal of Philosophical Studies*, 29:4, pp. 473-496.

Carment, D., Miles, C. and Cervin, V (1965) 'Persuasiveness and Persuasibility as Related to Intelligence and Extraversion', *British Journal of Social and Clinical Psychology*, 4, pp. 1-7.

Carter, J. A (2016) 'Robust Virtue Epistemology as Anti-Luck Epistemology: A New Solution', *Pacific Philosophical Quarterly*, 97: 1, pp. 140-155.

Casey, L (2023) *Baroness Casey Review: An independent review into the standards of behaviour and internal culture of the Metropolitan Police Service*, London: Metropolitan Police Service.

Cassam, Q (2014) *Self-Knowledge for Humans*, Oxford: Oxford University Press.

- (2016) 'Vice Epistemology', *The Monist*, 99:2, pp. 159-180.
- (2017) 'Diagnostic Error, Overconfidence, and Self-Knowledge', *Palgrave Communications*, 3:1, pp.1-8.
- (2019a) *Vices of the Mind: From the Intellectual to the Political*, Oxford: Oxford University Press.
- (2019b) *Conspiracy Theories*, Oxford: Oxford University Press.
- (2020a) 'The Metaphysical Foundations of Vice Epistemology', in I. J. Kidd, Q. Cassam, and H. Battaly (eds.) *Vice Epistemology: Theory and Practice*, New York: Routledge, pp. 37- 52
- (2021) 'Epistemic Vices, Ideologies, and False Consciousness', in M. J. Hannon and J. de Ridder (ed.) *The Routledge Handbook of Political Epistemology*, pp. 301-311.
- (2022) *Extremism: A Philosophical Analysis*, Oxford: Oxford University Press.
- (forthcoming) 'Some Vices of Vice Epistemology', *Metaphilosophy*.

Cheng-Hung, T (2014) 'Xunzi and Virtue Epistemology', *Universitas*, 41:3, pp. 121-142.

Clarkson, J. J., Tormala, Z. L., and Rucker, D. D. (2008) 'A New Look at the Consequences of Attitude Certainty: The Amplification Hypothesis', *Journal of Personality and Social Psychology*, 95:4, pp. 810-825.

- Code, L. (1984) 'Toward a "Responsibilist" Epistemology', in *Philosophy and Phenomenological Research*, 45:1, pp. 29-59.
- Confucius (2003) *Analects: With Selections from Traditional Commentaries*, E. Slingerland (trans.), Indianapolis, IN: Hackett.
- Cooper, D. E. (2008) 'Teaching and Truthfulness', *Studies in Philosophy and Education* 27, pp. 79-87.
- (2010) 'Education and the Experience of Beauty', in W. Keping (ed.) *International Yearbook of Aesthetics: Diversity and Universality in Aesthetics*, Beijing: Institute for Transcultural Studies, pp. 62-80.
  - (2018) *Animals and Misanthropy*, New York: Taylor and Francis.
- Cordell, S (2017) 'Group Virtues: No Great Leap Forward with Collectivism', *Res Publica*, 23:1, pp. 43-59.
- Corneanu, S (2011) *Regimens of the Mind: Boyle, Locke, and the Early Modern Cultural Anima Tradition*, Chicago, IL: University of Chicago Press.
- Costa jr., P. T and McCrae, R (1992) 'Four Ways Five Factors are Basic', *Personality and Individual Differences*, 13:6, pp. 653-665.
- Crerar, C (2018) 'Motivational Approaches to Intellectual Vice', *Australasian Journal of Philosophy*, 96:4, pp. 753-766.
- (2020) 'Must We Love Epistemic Goods?', *The Philosophical Quarterly*, 71:4, pp. 886-907.
- Darley, J. M and Batson, C. D (1973) 'From Jerusalem to Jericho: A Study of Situational and Dispositional Variables in helping Behaviour', *Journal of Personality and Social Psychology*, 27:1, pp. 100-108.
- Daukas, N. V (2006) 'Epistemic Trust and Social Location', *Episteme*, 3:1-2, pp. 109-124.
- (2018) 'Feminist Virtue Epistemology' in H. Battaly (ed.) *The Routledge Handbook of Virtue Epistemology*, pp. 379-392.
- Davies, B (2014) *Thomas Aquinas' Summa Theologiae: A Guide and Commentary*, Oxford: Oxford University Press.

de Raad, B., and Perugini, M. (eds.) (2002) *Big Five Assessment*, Göttingen: Hogrefe and Huber Publishers.

de Ridder, J (2022) 'Three Models for Collective Intellectual Virtues', in M. Alfano, C. Klein, and J. de Ridder (eds.) *Social Virtue Epistemology*, New York: Routledge, pp. 367-385.

de Rooij, B and de Bruin, B (2022) 'Real Life Collective Epistemic Virtue and Vice', in M. Alfano, C. Klein, and J. de Ridder (eds.) *Social Virtue Epistemology*, New York: Routledge, pp. 396-414.

Descartes, R (2012) *Rules for the Direction of the Mind* [1701], J. Cottingham, R. Stoothoff, and Dugald Murdoch (eds.) *The Philosophical Writings of Descartes*, Cambridge: Cambridge University Press.

DeYoung, R. D (2009) *Glittering Vices: A New Look at the Seven Deadly Sins and Their Remedies*, Grand Rapids, Michigan: Brazos Press.

Dillon, R (2012) 'Critical Character Theory: Toward a Feminist Perspective on 'Vice' (and Virtue)' in S. L. Crasnow and A. M. Superson (eds.) *Out from the Shadows: Analytical Feminist Contributions to Traditional Philosophy*, Oxford: Oxford University Press, pp. 83-114.

Doris, J (1998) 'Persons, Situations, and Virtue Ethics', *Noûs*, 32, pp. 504, 530.

- (2002) *Lack of Character: Personality and Moral Behaviour*, Cambridge: Cambridge University Press.
- (2010) 'Heated Agreement: Lack of Character as being for the Good', *Philosophical Studies*, 148, pp. 135-146.
- (2021) *Character Trouble: Undisciplined Essays on Moral Agency and Personality*, Oxford: Oxford University Press.

Dougherty, M. R (2019) 'The Importance of Roles in the Skill Analogy', *Journal of Ethics and Social Philosophy*, 17:1, pp. 75-102.

Driver, J (1989) 'The Virtues of Ignorance', *The Journal of Philosophy*, 86:7, pp. 373-384.

- (2000) 'Moral and Epistemic Virtue' in G. Axtell (ed.) *Knowledge, Belief, and Character*, Oxford, Rowman and Littlefield, pp. 123-134.
- (2001) *Uneasy Virtue*, Cambridge: Cambridge University Press.

- (2003) 'The Conflation of Moral and Intellectual Virtue', *Metaphilosophy*, 34:3, pp. 367-383.

Duncker, K (1945) 'On Problem-Solving', *Psychological Monographs*, 58: 5, pp. i-113.

Durand, R, Newby, R., Tant, K., and Trepongkaruna, S (2013) 'Overconfidence, Overreaction, and Personality', *Review of Behavioural Finance*, 5:2, pp. 104-113.

Eagly, A. H., and Chaiken, S (1993) *The Psychology of Attitudes*, Texas: Harcourt Brace Jovanovich College Publishers.

Ebersole, C., Atherton, O., Belanger, A., Skulborstad, H., Allen, J., Banks, J., Baranski, E., Bernstein, M., Bonfiglio, D., Boucher, L., Brown, E., Budiman, N., Cairo, A., Capaldi, C., Chartier, C., Chung, J., Cicero, D., Coleman, J., Conway, J., Davis, W., Devos, T., Fletcher, German, K., Grahe, J., Hermann, A., Hicks, J., Honeycutt, N., Humphrey, B., Janus, M., Johnson, D., Joy-Gaba, J., Juzeler, H., Keres, A., Kinney, D., Kirshenbaum, J., Klein, J., Klein, R., Lucas, R., Lustgraaf, C., Martin, D., Menon, M., Metzger, M., Moloney, J., Morse, P., Prislín, R., Razza, T., Re, D., Rule, N., Sacco, D., Sauerberger, K., Shrider, E., Shultz, M., Siemsen, C., Sobocko, K., Sternglanz, R. W., Summerville, A., Tskhay, K., van Allen, Z., Vaughn, L., Walker, R., Weinberg, A., Wilson, J., Wirth, J., Wortman, J., Nosek, B. (2016). Many Labs 3: Evaluating participant pool quality across the academic semester via replication', *Journal of Experimental Social Psychology*, 67, pp. 68-82.

Eisenberg, N. (2000) 'Emotion, Regulation, and Moral Development', *Annual Review of Psychology*, 51, pp. 665-97.

Fairweather, A (2001) 'Epistemic Motivation' in A. Fairweather and L. Zagzebski (eds.) *Virtue Epistemology: Essays in Epistemic Virtue and Responsibility*, Oxford: Oxford University Press.

Fazio, R. H., and Olson, M. A. (2007) 'Attitudes: Foundations, Functions and Consequences, in M. A. Hogg and J. Cooper (eds), *The SAGE Handbook of Social Psychology*, London: SAGE, pp. 139-160.

Fleeson, W and Jayawickreme, E (2015) 'Whole Trait Theory', *Journal of Research in Personality*, 56, pp. 82-92.

Fleisher, W (2017) 'Virtuous Distinctions', *Synthese* 194: 8, pp. 2973-3003.

Flores, C and Woodard, E (2023) 'Epistemic Norms of Evidence-Gathering', *Philosophical Studies*, 180: 9, pp. 2547-2571.

Foot, P (1958) 'Moral Arguments', *Mind*, 67: 268, pp. 502-513.

- (2002) *Virtues and Vices and Other Essays in Moral Philosophy* [1978], Oxford: Oxford University Press.

-

Forstenzer, J (2018) 'The Teaching Excellence Framework: Epistemic Insensibility and the Question of Purpose', *Journal of Philosophy of Education*, 52:3, pp. 548-574.

Frankfurt, H. G (2005) *On Bullshit*, Princeton: Princeton University Press.

Fricker, M (2007) *Epistemic Injustice: Power and the Ethics of Knowing*, Oxford: Oxford University Press.

- (2010) 'Can there be Institutional Virtues', in T. S. Gendler and J. Hawthorne (eds.) *Oxford Studies in Epistemology*, Oxford: Oxford University Press, pp. 235-252.
- (2013) 'Epistemic Justice as a Condition of Political Freedom?', *Synthese*, 190: 7, pp. 1317-1332.
- (2020) 'Institutional Epistemic Vice: The Case of Inferential Inertia', in J. Kidd, Q. Cassam, and H. Battaly (eds.) *Vice Epistemology*, New York: Routledge, pp. 89-107.

Friedman, J (2020) 'The Epistemic and the Zetetic', *The Philosophical Review*, 129:4, pp. 501-536.

Gardiner, S. M. (2011) *A Perfect Moral Storm: The Ethical Tragedy of Climate Change*, Oxford: Oxford University Press.

Geen, R (2001) *Human Aggression*, 2<sup>nd</sup> edition, Buckingham: Open University Press.

Gelfert, A (2013) 'Climate Scepticism, Epistemic Dissonance, and the Ethics of Uncertainty', *Philosophy and Public Issues – Filosofia E Questioni Pubbliche*, 3:1, pp. 167-208.

Gelles, D., Kitroeff, N., Nicas, J., and Ruiz, R. R (2019) 'Boeing Was 'Go, Go, Go' to Beat Airbus with the 737 Max', *The New York Times*, 23<sup>rd</sup> March, accessible at: <https://www.nytimes.com/2019/03/23/business/boeing-737-max-crash.html>.

Gentleman, A (2018) 'Whistleblowers contradict No 10 over destroyed Windrush landing cards', *The Guardian*, Wednesday 18<sup>th</sup> April, accessible at: <https://www.theguardian.com/uk-news/2018/apr/18/whistleblowers-contradict-no-10-over-destroyed-windrush-landing-cards>.

Geertz, C (1973) *The Interpretation of Culture: Selected Essays*, New York: Basic Books.

Gilbert, M (1989) *On Social Facts*, New York: Routledge.

- (2000) *Sociality and Responsibility: New Essays in Plural Subject Theory*, Maryland: Rowman and Littlefield.
- (2004) 'Collective Epistemology', *Episteme*, 1: 2, pp. 95-107.
- (2006) *A Theory of Political Obligation: Membership, Commitment, and the Bonds of Society*, Oxford University Press.
- (2013) *Joint Commitment: How We Make the Social World*, Oxford: University Press.
- (2023) *Life in Groups: How We Think, Feel, and Act Together*, Oxford: Oxford University Press.

Ginet, C (2001) 'Deciding to Believe', in M. Steup (ed.) *Knowledge, Truth, and Duty*, Oxford: Oxford University Press, pp. 63-76.

Golberg, L. R (1993) 'The Structure of Phenotypic Personality Traits', *The American Psychologist*, 48:1, pp. 26-34.

Goldman, A (1979) 'What is Justified Belief?' in G. S. Pappas (ed.) *Justification and Knowledge* Dordrecht: Reidel.

- (1992) 'Epistemic Folkways and Scientific Epistemology' in A. Goldman (ed.) *Liaisons*, Cambridge: MA: MIT Press, pp. 155-175.

Grasswick, H (2017) 'Epistemic Injustice in Science', in I. J. Kidd, J. Medina, and G. Pohlhaus Jr. (eds.) *The Routledge Handbook of Epistemic Injustice*, pp. 313-323.

Greco, J (1993) 'Virtues and Vices of Vice Epistemology', *Canadian Journal of Philosophy*, 23:3, pp. 413-432.

- (2000) 'Virtue, Skepticism, and Context', in G. Axtell (ed.) *Knowledge, Belief, and Character*, Oxford, Rowman and Littlefield, pp. 55-71.
- (2009) 'Knowledge and Success from Ability', *Philosophical Studies*, 142:1, pp. 17-26.

- (2010) *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*, Cambridge: Cambridge University Press.
- (2020) 'Knowledge-Producing Abilities' in C. Kelp and J. Greco (eds.) *Virtue-Theoretic Epistemology: New Methods and Approaches*, pp. 124-146.

Haidt, J and Haslam, N (2016) 'Campuses are places for open minds – not where debate is closed down', *The Guardian*, Sunday 10<sup>th</sup> April, accessible at: <https://www.theguardian.com/commentisfree/2016/apr/10/students-censorship-safe-places-platforming-free-speech>.

Hare, R. M (1952) *The Language of Morals*, Oxford: Oxford University Press.

Hare, W (1987) 'Open-mindedness in Moral Education: Three Contemporary Approaches', *Journal of Moral Education*, 16:2, pp. 99-107.

Harman, G (1999) 'Moral Philosophy Meets Social Psychology: Virtue Ethics and the Fundamental Attribution Error', *Proceedings of the Aristotelian Society*, 99, pp. 313-331.

- (2000) 'The Nonexistence of Character Traits', *Proceedings of the Aristotelian Society*, 100: 1, pp. 223-226.
- (2009) 'Scepticism about Character Traits', *The Journal of Ethics*, 13:2/3, pp. 235-242.

Haslanger, S (2015) 'Social Structure, Narrative, and Explanation', *Canadian Journal of Philosophy*, 45, 1-15.

Hazlett, A (2021) 'False Intellectual Humility', in M. Alfano, M. P. Lynch, and A. Tanesini (eds.) *The Routledge Handbook of Philosophy of Humility*, New York: Routledge, pp. 313-324.

Hollis, M (2002) *The Philosophy of Social Science: an introduction*, Cambridge: Cambridge University Press.

Horst, D (2022) 'Is Epistemic Competence a Skill?', *Australasian Journal of Philosophy*, 100:3, pp.509-523.

House Committee on Transportation and Infrastructure, (2020) The Boeing 737 MAX Aircraft: Costs, Consequences, and Lessons from its Design, Development, and Certification – Preliminary Investigative Findings. Accessible at: [https://www.govinfo.gov/content/pkg/GOVPUB-Y4\\_T68\\_2](https://www.govinfo.gov/content/pkg/GOVPUB-Y4_T68_2)

Hume, D (1960) *A Treatise of Human Nature* [1739], L. A. Selby-Bigge (ed.), Oxford: Clarendon Press.

- (1975) *An Enquiry Concerning the Principles of Moral* [1751], P. H. Nidditch (ed.), Oxford: Clarendon Press.
- (1975) *An Enquiry Concerning Human Understanding* [1748], P. H. Nidditch (ed.), Oxford: Clarendon Press.

Hurka, T (2001) *Virtue, Vice, and Value*, Oxford: Oxford University Press.

Hursthouse, R (1999) *On Virtue Ethics*, Oxford: Oxford University Press.

Isen, A. M and Levin, P. F (1973) 'Effect of Feeling Good on Helping: Cookies and Kindness', *Journal of Personality and Social Psychology*, 21: 3, pp. 384-388.

Isen, A. M., Daubman, K. A., & Nowicki, G. P (1987). 'Positive Affect Facilitates Creative Problem Solving', *Journal of Personality and Social Psychology*, 52: 6, pp. 1122-1131.

James, W (2010) *The Will to Believe* [1896], Auckland, The Floating Press.

Jeffrey, A and Beary, A (2023) 'Against Reductivist Character Realism', *Philosophical Psychology*, 36: 1, pp. 186-213.

Joëls, M., Pu, Z., Wiegert, O., Oitzl, M. S., and Krugers, H. J. (2006) 'Learning Under Stress: How Does it Work?', *Trends in Cognitive Sciences*, 10, pp. 152-158.

John, O. P. and Srivastava, S (1999). 'The Big Five Trait Taxonomy: History, Measurement, and Theoretical Perspectives', in O. P. John, R. W. Robins, and L. A. Pervin (eds.) *Handbook of personality: Theory and Research*, New York: Guilford Press, pp. 102-138.

Johnson, C (2020) 'Teaching to the Test: How Schools Discourage *phronesis*', in I. J. Kidd, H. Battaly, and Q. Cassam (eds.) *Vice Epistemology*, New York: Routledge, pp. 225-238.

Kagan, S (2014) 'An Introduction to Ill-Being', in M. Timmons (ed.) *Oxford Studies in Normative Ethics*, vol. 4, pp. 261-288.

Kallestrup, J and Pritchard, D (2014) 'Virtue Epistemology and Epistemic Twin Earth', *European Journal of Philosophy*, 22, pp. 335-357.

Kelp, C (2013) 'Knowledge: The Safe-Apt View', *Australasian Journal of Philosophy*, 91:2, pp. 265-278.

Kessler, G., Rizzo, S., and Kelly, M (2021) 'Trump's false or misleading claims total 30,573 over 4 years', *The Washington Post*, 24<sup>th</sup> January, accessible at: <https://www.washingtonpost.com/politics/2021/01/24/trumps-false-or-misleading-claims-total-30573-over-four-years/>.

Kidd, I. J (2016a) 'Charging Other with Epistemic Vice', *The Monist* 99.3, pp. 181-197.

- (2016b) 'Intellectual Humility, Confidence, and Argumentation', *Topoi*, 35, pp. 395-402.
- (2017) 'Capital Epistemic Vices', *Social Epistemology Review and Reply Collective*, 6:8, pp. 11-16.
- (2018a) 'Deep Epistemic Vices', *Journal of Philosophical Research*, 43, pp. 43-67.
- (2018b) 'Confucianism, Curiosity, and Moral Self-cultivation', in I. Inan, L. Watson, and D. Whitcomb, S. Yigit (eds.) *The Moral Psychology of Curiosity*, New York: Rowman and Littlefield, pp. 97-116.
- (2019) 'Epistemic Corruption and Education', *Episteme*, 16:2, pp. 220-235.
- (2020) 'Epistemic Corruption and Social Oppression', in I. J. Kidd, Q. Cassam, and H. Battaly (eds.) *Vice Epistemology*, New York: Routledge, pp. 69-86.
- (2021a) 'Epistemic Corruption and Political Institutions', in M. J. Hannon and J. de Ridder (eds.) *The Routledge Handbook to Political Epistemology*, New York, Routledge, pp. 347-358.
- (2021b) 'A Case for an Historical Vice Epistemology' in L. Candiotto (ed.) *The social Dimensions of the Ethics of Knowledge: Intellectual Virtues and Intellectual Vices in Epistemic Practices*, a special issue of *Humana Mente*, 14:39, pp. 69-86.
- (2021c) 'Epistemic Corruption and the Research Impact Agenda' (co-authored with J. Chubb and J. Forstenzer), *Theory and Research in Education*, 19.2, pp. 148-167.
- (2022) 'Character, Corruption, and 'Cultures of Speed' in the Academy', in A. Mahon (ed.) *The Promise of the University; Reclaiming Humanity, Humility, and Hope*, Dordrecht: Springer, pp. 17-28
- (2022) 'From Vice Epistemology to Critical Character Epistemology' in M. Alfano, C. Klein, and J. de Ridder (eds.) *Social Virtue Epistemology*, New York: Routledge, pp. 84-102.

- (2023a) 'Corrupted Temporalities, 'Cultures of Speed', and the Possibility of Collegiality', in G. Dunne (ed.) *Reimagining Epistemic Injustice*, special issue of *Educational Philosophy and Theory*, 55.3, pp. 330-342.
- (2023b) 'Multidimensionalism, Resistance, and the Demographic Problem', *European Journal of Analytic Philosophy*, 19:1, pp: 5-30.

Klein, J (2014) 'Of Archery and Virtue: Ancient and Modern Conceptions of Value', *Philosopher's Imprint*, 14:19, pp. 1-16.

Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams, R. B., Jr., Alper, S., Aveyard, M., Axt, J., Bahník, Š., Berkics, M., Bernstein, M., Bialobrzeska, O., Bocian, K., Brandt, M., Cantarero, K., Cemalcilar, Z., Cicero, D., Chandler, J., Chatard, A., Chen, E., Cheong, W., Coen, S., Collisson, B., Conway, J., Corker, K., Curran, P., Cushman, F., Dalla Rose, A., Davis, W., Devos, T., Dogulu, C., Dunham, Y., Eller, A., Finck, C., Friedman, M., Giessner, S., Gnams, T., Gómez, Á., Graham, J., Grahe, J., Green, E., Haigh, M., Haines, E., Heffernan, M., Hicks, J., Houdek, P., Huntsinger, J., Ijzerman, H., Inbar, Y., Kende, A., Innes-Ker, Å., Jiménez-Leal, W., Kappes, H., Karabati, S., Keller, V., Kervyn, N., Krueger, L., Lakens, D., Lazarević, L., Levitan, C., Lins, S., John, M.-S., Mallett, R., Milfont, T., Morris, W., Myachykov, A., Neave, N., Nichols, A., O'Donnell, S., Orosz, G., Pérez-Sánchez, R., Petrovic, B., Pilati, R., Pollmann, M., Salomon, E., Schmidt, K., Sekerdej, M., Smith, M., Smith-Castro, V., Sobkow, A., Stouten, J., Street, C., Traczyk, J., Torres, D., Theriault, J., Ujhelyi, A., van Aert, R., van Assen, M., Vaughn, L., Vázquez, A., Verniers, C., Verschoor, M., Vranka, M., Wichman, A., Williams, L., Young, L., Zelenski, J., Nosek, B. (2017). Many Labs 2: Investigating Variation in Replicability Across Sample and Setting. *Advances in Methods and Practices in Psychology Science*, 1:4, pp. 443-490.

Klein, R., Ratliff, K., Vianello, M., Adams, R., Bahník, Š, Bernstein, M., Bocian, K., Brandt, M., Brooks, B., Brumagh, C., Cemalcilar, Z., Chandler, J., Cheong, W., Davis, W., Devos, T., Eisner, M., Frankowska, N., Furrow, D., Galliani, E., Hasselman, F., Hicks, J., Hovermale, J., Hunt, S. J., Huntsinger, J., IJzerman, H., John, M.-S., Joy-Gaba, J., Kappes, H., Krueger, L., Kurtz, J., Levitan, C., Mallett, R., Morris, W., Nelson, A., Nier, J., Packard, G., Pilati, R., Rutchick, A., Schmidt, K., Skorinko, J., Smith, R., Steiner, T., Storbeck, J., van Swol, L., Thompson, D., van 't Veer, A., Vaughn, L., Vranka, M., Wichman, A., Woodzicka, J., & Nosek, B (2014) 'Investigating Variation In Replicability: A "Many Labs" Replication Project', *Social Psychology*, 45:3, pp. 142-52.

Korsgaard, C (1996) *The Sources of Normativity*, Cambridge: Cambridge University Press.

Kotsonis, A (2021a) 'The Curious Case of the Excellent Gossiper', *Philosophia*, <https://doi.org/10.1007/s11406-021-00444-1>.

- (2021b) 'The Platonic Conception of Intellectual Virtues: Its Significance for Virtue Epistemology', *Synthese*, 198: 3, pp. 2045-2060.
- (2022) 'The Curious Case of the Excellent Gossiper', *Philosophia*, 50: 3, pp. 1207-1222.

Kotsonis, A (2021) 'The Platonic Conception of Intellectual Virtues: Its Significance for Virtue Epistemology', *Synthese*, 198: 3, pp. 2045-2060.

Kotzee, B (2014) *Education and the Growth of Knowledge: perspectives from social and virtue epistemology*, Massachusetts: John Wiley and Sons Inc.

Kouchaki, M., Gino, F., and Jami, A (2014) 'The Burden of Guilt: Heavy Backpacks, Light Snacks, and Enhanced Morality', *Journal of Experimental Psychology*, 143: 1, pp. 414-424.

Krahé, B. (2001) *The Social Psychology of Aggression*, Philadelphia: Taylor and Francis.

Kunda, Z (1990) 'The Case for Motivated Reasoning', *Psychological Bulletin*, 108:3, pp. 480-498.

Kvanvig, J (1992) *The Intellectual Virtues and the Life of the Mind*, Maryland: Rowman and Littlefield.

Lackey, J (2016) 'What is Justified Group Belief?', *The Philosophical Review*, 125:3, pp. 341-396.

Lahroodi, R (2007) 'Collective Epistemic Virtues', *Social Epistemology*, 21:3, pp. 281-297.

- (2018) 'Collective Virtue Epistemology', in H. Battaly (ed.) *The Routledge Handbook of Virtue Epistemology*, New York: Routledge, pp.

Latané, B and Darley, J (1970) *The Unresponsive Bystander: Why Doesn't he Help me?* New York: Appleton-Century-Crofts.

Latané, B and Rodin, J (1969) 'A Lady in Distress: Inhibiting Effects of Friends and Strangers on Bystander Intervention', *Journal of Experimental Social Psychology*, 5:2, pp. 189-202.

Lewis, D. K (1969) *Convention: A Philosophical study*, Massachusetts, Wiley-Blackwell.

List, C and Pettit, P (2011) *Group Agency: The Possibility, Design, and Status of Corporate Agents*, Oxford: Oxford University Press.

Locke, J (1996) *Some Thoughts Concerning Education and Of the Conduct of the Understanding* [1706], R. W. Grant and N. Tarcov (eds.), Indianapolis, IN: Hackett

- (2008) *An Essay Concerning Human Understanding* [1690], P. Phemister (ed.) Oxford: Oxford University Press.

Lounsbury, J. W, Saudargas, R. A, and Gibson, L. W (2004) 'An Investigation of Personality Traits in Relation to Intention to Withdraw from College', *Journal of College Student Development*, 45:5, pp.517-534.

Machery, E. (2016) 'De-Freuding Implicit Attitudes, in M. S. Brownstein and J. M. Saul (eds) *Implicit Bias and Philosophy: Metaphysics and Epistemology* (Vol. 1), Oxford: Oxford University Press, pp. 104–129.

MacIntyre, A. C. (2013) *After Virtue: A Study in Moral Theory*, 3<sup>rd</sup> ed., Indiana: University of Notre Dame Press.

Maio, G. R and Haddock, G (2018) *The Psychology of Attitudes and Attitude Change* (3<sup>rd</sup> ed.), London: SAGE.

Malcom, N (2001) *Ludwig Wittgenstein: A Memoir, with Wittgenstein's Letters to Malcolm*, 2<sup>nd</sup> ed., Oxford: Clarendon Press.

Manson, N (2012) 'Epistemic Restraint and the Vice of Curiosity', *Philosophy*, 87:2, pp. 239-259.

Marshall, G. N., Wortman, C. B., Kusulas, J. W., Hervig, L. K., & Vickers, R. R. Jr., (1992) 'Distinguishing Optimism from Pessimism: Relations to Fundamental Dimensions of Mood and Personality', *Journal of Personality and Social Psychology*, 62, pp. 1067–1074.

Masala, A and Webber, J (eds.) (2016) *From Personality to Virtue: Essays on the Philosophy of Character*, Oxford: Oxford University Press.

Matthews, T (2022) 'Deepfakes, Intellectual Cynics, and the Cultivation of Digital Sensibility', *Royal Institute of Philosophy Supplements*, 92, pp. 67-85.

- (2023) 'Deepfakes, Fake Barns, and Knowledge from Videos', *Synthese*, 201:3, <https://doi-org.nottingham.idm.oclc.org/10.1007/s11229-022-04033-x>.
- McAdams, D. P., and Pals, J. L. (2006) 'A New Big Five: Fundamental Principles for an Integrative Science of Personality', *American Psychologist*, 61:3, pp. 204-217.
- McCrae, R. R. (2009) 'The Five-Factor Model of Personality Traits: Consensus and Controversy', in P. J. Corr and G. Matthews (eds.) *The Cambridge Handbook of Personality Psychology*, Cambridge: Cambridge University Press, pp. 148-161.
- McDowell, J. (1979) 'Virtue and Reason', *The Monist*, 62: 3, pp. 331-350.
- (1998) *Mind, Value, and Reality*, Massachusetts: Harvard University Press.
- McKinnon, C. (1989) 'Ways of Wrong-Doing, the Vices, and Cruelty', *The Journal of Value Inquiry*, 23:4, pp. 319-335.
- Medina, J. (2013) *The Epistemology of Resistance; Gender and Racial Oppression, Epistemic Injustice, and Resistant Imaginations*, Oxford: Oxford University Press.
- Mehl, M. R., Gosling, S. D. And Pennebaker, J. W. (2006) 'Personality in its Natural Habitat: Manifestations and Implicit Folk Theories of Personality in Daily Life', *Journal of Personality and Social Psychology*, 90:5, pp. 862-877.
- Meir, B. (2018) *Painkiller: An Empire of Deceit and the Origin of America's Opioid Epidemic*, New York: Random House.
- Merritt, M., Doris, J., and Harman, G. (2010) 'Character', in J. Doris (ed.) *The Moral Psychology Handbook*, Oxford: Oxford University Press, pp. 355-401.
- Mi, C. (2015) 'What is Knowledge? When Confucius Meets Sosa', *Dao: a Journal of Comparative Philosophy*, 14:3, pp. 355-367.
- (2017) 'Reflective Knowledge: Confucius and Virtue Epistemology', *Comparative Philosophy*, 8:2, pp. 30-45.
- Milgram, S. (1974) *Obedience to Authority: An Experimental View*, New York: Harper and Row.
- Miller, C. B. (2013) *Moral Character: An Empirical Theory*, Oxford: Oxford University Press.

- (2014) *Character and Moral Psychology*, Oxford: Oxford University Press.
- (2017) 'Character and Situationism: New Directions', *Ethical Theory and Moral Practice*, 20: 3, pp. 459-471.
- (2018) 'Virtue as a Trait', in N. Snow (ed.) *Oxford Handbook of Virtue*, New York: Oxford University Press, pp. 9-34.

Millar, B and Richardson, L (2023) 'Grief, Smell and the Olfactory Air of a Person', *Pacific Philosophical Quarterly*, DOI: 10.1111/papq.12443.

Mischel, W (1968) *Personality and Assessment*, New York: Routledge.

Montmarquet, J (1986) 'The Voluntariness of Belief', *Analysis*, 46:1, pp. 49-53.

- (1987) 'Epistemic Virtue', *Mind*, 96, pp. 482-497.
- (1993) *Epistemic Virtue and Doxastic Responsibility*, Lanham, MD: Rowman and Littlefield.
- (2000) 'An Internalist Conception of Virtue' in G. Axtell (ed.) *Knowledge, Belief, and Character*, Oxford, Rowman and Littlefield, pp. 135-147.

Monypenny, A (2021) 'Between Vulnerability and Resilience: A Contextualist Picture of Protective Epistemic Character Traits', *Journal of Philosophy of Education*, 55:2, pp. 358-370.

Nagel, T (1979) *Mortal Questions*, Cambridge: Cambridge University Press.

Navarro, J (2015) 'No Achievement Beyond Intention: A New Defence of Robust Virtue Epistemology', *Synthese*, 192: 10, pp. 3339-3369.

Nyanatiloka (2004) *Buddhist Dictionary: Manual of Buddhist Terms and Doctrines*, Kandy, Sri Lanka: Buddhist Publication Society.

Olin, L. and Doris, J (2014) 'Vicious Minds', *Philosophical Studies*, 168:3, pp. 665-692.

Oreskes, N and Conway, E. M (2012) *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*, London: Bloomsbury.

Pereboom, D (2014) *Free Will, Agency, and Meaning in Life*, Oxford: Oxford University Press.

- (2021) *Wrongdoing and the Moral Emotions*, Oxford: Oxford University Press.

Pettit, P and Schweikard, D (2006) 'Joint Actions and Group Agents', *Philosophy of the Social Sciences*, 36:1, pp. 18-39.

Picone, P. M., De Massis, A., Tang, T., Piccolo, R. F. (2021) 'The Psychological Foundations of Management in Family Firms: Values, Biases, and Heuristics' *Family Business Review* 34: 1, pp. 12-32.

Piefke, M and Glienke, K (2017) 'The Effects of Stress on Prospective Memory: A Systematic Review', *Psychology and Neuroscience*, 10: 3, pp. 345-362.

Plato (1966) *Republic*, I. A. Richards (ed. & trans.), Cambridge: Cambridge University Press.

- (1979) *Gorgias*, T. Irwin (trans.), Oxford: Clarendon Press
- (1996) *Hippias Minor*, R. Allen (trans.), *The Dialogues of Plato, Volume 3: Ion, Hippias Minor, Laches, and Protagoras*, Yale: Yale University Press.
- (1998) *Symposium*, R. Waterfield (trans.), Oxford: Oxford University Press.
- (2002) *Phaedrus*, R. Waterfield (trans.), Oxford: Oxford University Press.

Pritchard, D (2012) 'Anti-Luck Virtue Epistemology', *Journal of Philosophy*, 109: 3), pp. 247-279.

- (2017) 'Anti-Risk Virtue Epistemology and Negative Epistemic Dependence', *Synthese*, 197: 7, pp. 2879-2894.
- (2020) 'Anti-Risk Virtue Epistemology' in J. Greco and C. Kelp (eds.) *Virtue Theoretic Epistemology: New Methods and Approaches*, pp. 203-224.

Proctor, R. N (2011) *Golden Holocaust: Origins of the Cigarette Catastrophe and the Case for Abolition*, Berkeley, University of California Press.

Putnam, D (1997) 'The Intellectual Bias of Virtue Ethics', *Philosophy*, 72:280, pp. 303-311.

Quinton, A (1976) 'The Presidential Address: Social objects', *Proceedings of the Aristotelian Society*, 76:1, pp. 1-28.

Reuter, K (2019) 'Dual Character Concepts', *Philosophy Compass*, 14:1, accessible at: <https://doi.org/10.1111/phc3.12557>.

Riggs, W. D (2003) 'Understanding "Virtue" and the Virtue of Understanding', in M. dePaul and L. T. Zagzebski (eds.) *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Oxford: Oxford University Press, pp. 203-226.

- (2010) 'Open-Mindedness', *Metaphilosophy*, 41:1-2, pp. 172-188.
- Roberts, R and Woods, J (2007) *Intellectual Virtues: An Essay in Regulative Epistemology*, Oxford: Clarendon Press.
- Roberts, R. C and West, R (2020) 'The Virtue of Honesty: A Conceptual Exploration', in C. B. Miller and R. West (eds.) *Integrity, Honesty, and Truth Seeking*, Oxford: Oxford University Press, pp. 97-126.
- Roochnik, D (1990) *Of Art and Wisdom: Plato's Understanding of Techne*, Pennsylvania, Penn State University Press.
- Ross, L and Nisbett, R. E (2011) *The Person and the Situation: Perspectives of Social psychology*, London: Pinter and Martin.
- Ryle, G (1949) *The Concept of Mind*, Illinois: University of Chicago Press.
- Scheffler, S (1987) 'Morality Through Thick and Thin: A Critical Notice of *Ethics and the Limits of Philosophy*', *The Philosophical Review*, 96: 3, pp. 411-434.
- Schwabe, L and Wolf, O. T (2010) 'Learning under stress impairs memory formation', *Neurobiology of Learning and Memory*, 93: 2, pp.183-188.
- Searle, J (1990) 'Collective Intentions and Actions, in P. R. C. Jerry Morgan & M. Pollack (eds.), *Intentions in communication*, Massachusetts: MIT Press, pp. 401-415.
- Segerstrom, S. C, Castañeda, J. O, and Spencer, T. E (2003) 'Optimism Effects on Cellular Immunity: Testing the Affective and Persistence Models', *Personality and Individual Differences*, 35:7, pp.1615-1624.
- Sharpe, J. P., Martin, N. R., & Roth, K. A. (2011). 'Optimism and the Big Five Factors of Personality: Beyond Neuroticism and Extraversion', *Personality and Individual Differences*, 51:8, pp. 946-951.
- Sher, G (2005) *In Praise of Blame*, New York: Oxford University Press.
- Sherman, N (1991) *The Fabric of Character*, Oxford: Clarendon Press.
- Singh, Y., Adil, M. & Haque, S.M.I (2022) 'Personality Traits and Behaviour Biases: The Moderating Role of Risk-Tolerance', *Quality and Quantity*, 57, pp. 3549-3573.
- Smith, A (2008) 'Control, Responsibility, and Moral Assessment', *Philosophical Studies*, 138: 3, pp. 367-392.

Smith, Dame J (2016) *The Dame Janet Smith Review Report: An Independent Review into the BBC's Culture and Practices during the Jimmy Savile and Stuart Hall Years*, London: BBC.

Snow, N (2013) 'Notes Towards an Empirical Psychology of Virtue: Exploring the Personality Scaffolding of Virtue' in J. Peters (ed.) *Aristotelian Ethics in Contemporary Perspective*, New York: Routledge, pp. 130-144.

Sosa, E (1980) 'The Raft and the Pyramid' in *Midwest Studies in Philosophy*, 5:1, pp. 3-26

- (1991) *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge: Cambridge University Press.
- (2007) *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Oxford: Oxford University Press.
- (2009) *Reflective Knowledge*, Oxford: Oxford University Press.
- (2011) *Knowing Full Well*, Princeton: Princeton University Press.
- (2015) *Judgement and Agency*, Oxford: Oxford University Press
- (2021) *Epistemic Explanation: A Theory of Telic Normativity, and What it Explains*, Oxford: Oxford University Press.

Stalnaker, A (2010) 'Virtue as Mastery in Early Confucianism', *The Journal of Religious Ethics*, 38:3, pp. 404-428.

Statman, D (ed.) (1997) *Virtue Ethics*, Edinburgh: Edinburgh University Press.

Steup, M (2008) 'Doxastic Freedom', *Synthese*, 161: 3, pp. 375-392.

Stichter, M (2016) 'Practical Skills and Practical Wisdom in Virtue', *Australasian Journal of Philosophy*, 94:3, pp. 435-448.

- (2018) *The Skilfulness of Virtue: Improving our Moral and Epistemic Lives*, Cambridge: Cambridge University Press.
- (2020) 'Virtue as Skill: Self-regulation and Social Psychology', in E. Fridland and C. Pavese (eds.) *The Routledge Handbook of Philosophy of Skill and Expertise*, New York: Routledge, pp. 487-501.

Strawson, P (1963) *Freedom and Resentment*, London: Oxford University Press.

Stuewig, J., Tangney, J. P., Heigel, C., Harty, L., & McCloskey, L (2010) 'Shaming, Blaming, and Maiming: Functional Links Among the Moral Emotions, Externalization of Blame, and Aggression', *Journal of Research in Personality*, 44:1, 91-102.

Swartwood, J (2013) 'Wisdom as an Expert Skill', *Ethical Theory and Moral Practice*, 16: 3, pp. 511–528.

Sylvan, K (2020) 'Responsibilism within Reason' in C. Kelp and J. Greco (eds.) *Virtue-Theoretic Epistemology: New Methods and Approaches*, pp. 225-256.

Tanesini, A (2016) 'Calm Down Dear: Intellectual Arrogance, Silencing, and Ignorance', *Aristotelian Society Supplementary Volume*, 90:1, pp. 71-92.

- (2018a) 'Epistemic Vice and Motivation', *Metaphilosophy*, 49: 3, pp. 350-367.
- (2018b) 'Eloquent Silences', in C. R. Johnson (ed.) *Voicing Dissent*, New York: Routledge, pp. 109-128.
- (2021) *The Mismeasure of the Self: A Study in Vice Epistemology*, Oxford: Oxford University Press.

Tessman, L (2005) *Burdened Virtues: Virtue Ethics for Liberatory Struggles*, Oxford: Oxford University Press.

Toner, C (2014) 'The Full Unity of the Virtues', *Journal of Ethics*, 18:3, pp. 207-227.

Tuomela, R (2013) *Social Ontology*, New York: Oxford University Press.

Turri, J (2015) 'From Virtue Epistemology to Abilism', in C. B. Miller, R. Michael Furr, A. Knobel, and W. Fleeson (eds.) *Character: New Directions from Philosophy, Psychology, and Theology*, Oxford: Oxford University Press, pp. 315-330.

Vaccarezza, M. S (2017) 'The Unity of the Virtues Reconsidered: Competing Accounts in Philosophy and Positive Psychology', *Review of Philosophy and Psychology*, 8:3, pp. 637-651.

van Zyl, L (2015) 'Against Radical Pluralism', in M. Alfano (ed.) *Current Controversies in Virtue Theory*, Abingdon: Routledge, pp. 22-32.

Väyrynen, P (2021) 'Thick Ethical Concepts' in E. Zalta (ed.) *Stanford Encyclopaedia of Philosophy* (Spring 2021 Edition), accessible at: <https://plato.stanford.edu/archives/spr2021/entries/thick-ethical-concepts>.

Vernon, P. E (1963) *Personality Assessment: A Critical Survey*, New York: Routledge.

Vetter, B (2019) 'Are Abilities Dispositions?', *Synthese*, 196:1, pp.201-220.

- Vitz, R (2009) 'Doxastic Virtues in Hume's Epistemology', *Hume Studies*, 35:1-2, pp. 211- 229.
- Vogel, S and Schwabe, L (2016) 'Stress in the Zoo: Tracking the Impact of Stress on Memory Formation over Time', *Psychoneuroendocrinology*, 71, pp. 64-72.
- Wallace, J (1978) *Virtues and Vices*, New York: Cornell University Press.
- Watson, J. D. (1968) *The Double Helix: A Personal Account of the Discovery of the Structure of DNA*, London, Weidenfeld and Nicolson.
- Watson, L (2015) 'What is Inquisitiveness?', *American Philosophical Quarterly*, 52:3, pp. 273-288.
- (2018a) 'Education for Good Questioning: A Tool for Intellectual Virtues Education', *Acta Analytica*, 33:5, pp. 353-370.
  - (2018b) 'Curiosity and Inquisitiveness', in H. Battaly (ed.) *The Routledge Handbook of Virtue Epistemology*, New York: Routledge, pp. 155-166.
  - (2020) 'Vices of Questioning in Public Discourse', in I. J. Kidd, Q. Cassam, and H. Battaly (eds.) *Vice Epistemology*, New York: Routledge, pp. 239-258.
- Watt, S. E., Maio, G. R., Haddock, G., and Johnson, B. T (2008) 'Attitude Functions in Persuasion: Matching, Involvement, Self-affirmation, and Hierarchy', in W. D. Crano and R. Prislin (eds.) *Attitudes and Attitude Change*, New York: Psychology Press, pp. 189-211.
- Weber, M (1922) *Economy and Society*, G. Roth and C. Wittich (eds.), Berkeley: University of California Press.
- Wilkowski, B and Robinson, M (2010) 'The Anatomy of Anger: An Integrative Cognitive Model of Trait Anger and Reactive Aggression' *Journal of Personality*, 78, pp. 9-38.
- Williams, B (1981) *Moral Luck: Philosophical Papers, 1973-1980*, Cambridge: Cambridge University Press.
- (1985) *Ethics and the Limits of Philosophy*, Cambridge: Harvard University Press.
  - (2002) *Truth and Truthfulness: An Essay in Genealogy*, Princeton: Princeton University Press.
- Williams, D. W (1992) 'Dispositional Optimism, Neuroticism, and Extraversion', *Personality and Individual Differences*, 13:4, pp. 475-477.

Williams, W (2020), *Windrush Lessons Learned Review*: independent review by Wendy Williams, House of Commons 93, accessible at: [https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment\\_data/file/874022/6.5577\\_HO\\_Windrush\\_Lessons\\_Learned\\_Review\\_WEB\\_v2.pdf](https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/874022/6.5577_HO_Windrush_Lessons_Learned_Review_WEB_v2.pdf).

Wollstonecraft, M (1995) *A Vindication of the Rights of Men with A Vindication of the Rights of Women and Hints* [1792], S. Tomaselli (ed.), Cambridge: Cambridge University Press.

Wolterstorff, N (1996) *John Locke and the Ethics of Belief*, Cambridge: Cambridge University Press.

Woodard, C (2022) 'The Value and Significance of Ill-Being', *Midwest Studies in Philosophy*, <https://doi.org/10.5840/msp202212130>.

Wright, J. C, Warren, M. T, and Snow, N. E (2020) *Understanding Virtue: Theory and Measurement*, Oxford: Oxford University Press.

Yao, X (2012) 'The Way, Virtue, and Practical Skill in the Analects', *Journal of Chinese Philosophy*, 39:1, pp. 26-43.

Yolton, S (1998) *John Locke: A Descriptive Bibliography*, Bristol: Thoemmes.

Zagzebski, L. T (1996) *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge: Cambridge University Press.

Zunick, P. V., Teeny, J. D., and Fazio, R. H. (2017) 'Are Some Attitudes More Self-defining than Others? Assessing Self-related Attitude Functions and their Consequences', *Personality and Social Psychology Bulletin*, 43:8, pp. 1136-1149.