4

Dare to Think?

In the last two chapters, I argued that beliefs are pervasively outsourced to other people and to the environment, and that belief revision often occurs in response to changes in the cues that scaffold our beliefs. I've suggested that in light of these facts, we need to ensure the scaffolding of better beliefs. We need, that is, to manage the epistemic environment. Bad beliefs are produced by a faulty environment and better beliefs are best promoted by environmental engineering.

This kind of suggestion is certain to face opposition. Resistance will come on at least two fronts. First, changing minds by changing the environment—rather than by giving people *reasons* to change their minds—seems manipulative. I'll address that concern in Chapter 6 (where we'll see that the objection is mistaken, not because it's not manipulative to change minds without giving reasons, but because changing the environment in the relevant ways *is* giving reasons). In this chapter and the next, I'll focus on the prospects for cognition in the *absence* of environmental engineering and other kinds of scaffolding. I'll argue that individual cognition (unaided) is very much less powerful than we tend to think. Worse, if we refuse to aid the individual cognizer, others will be all too eager to fill the gap, often in ways that leave them, and us, worse off.

The view I defend in the next three chapters has a passing resemblance to a more familiar position, one I used to endorse myself. This position is also sometimes offered in response to worries about epistemic paternalism, and holds that Enlightenment ideals of intellectual autonomy are psychologically unrealistic; instead paternalistic measures must take up the slack (Conly 2013; Levy 2012; Moles 2007). This familiar response appeals to the rationality deficit view of ourselves we outlined in Chapter 1, according to which contemporary psychology has

Bad Beliefs: Why They Happen to Good People. Neil Levy, Oxford University Press. © Neil Levy 2022.
DOI: 10.1093/oso/9780192895325.003.0004

shown that rationality is a scarce resource and that for the most part we respond relatively unthinkingly. That's not a view I now endorse: one of the aims of this book is to argue that we are in fact rational animals. If epistemic engineering is justified, it's not because we're not rational enough to respond to epistemic challenges. It's because we're never sufficiently rational on our own.

Regulative Epistemology

Of course, it's easy to find examples of individual reasoning going wrong. Too easy. Piling up example after example won't go far toward showing that individual cognition isn't powerful (we can also, it should be conceded, find plenty of examples of individual reasoning going well). We need to be fair to individual reasoning: we want to focus on examples when it is conducted carefully and thoughtfully. I'll focus here on what might reasonably be taken to be individual reasoning given its best shot. I'll focus on individual reasoning in the form recommended by leading contemporary epistemologists who explicitly aim to develop practical guides to good reasoning. I'll focus, that is, on reasoning as recommended by regulative epistemology.

Traditional epistemology is concerned with the nature of knowledge and of related concepts, like justification. Traditional epistemology is in one sense rather *un*traditional: the narrowing of focus to these questions represents a departure from an earlier way of thinking about knowledge that was more explicitly regulative. Philosophers like Locke and Descartes were concerned with how we can think better—with how we can *attain* knowledge—as well as with the nature of knowledge. The tradition to which Gettier and his many intellectual offspring belong represents a newly narrowed focus for epistemology (see Ballantyne 2019 for a sketch of this history).

Regulative epistemology offers us news we can use: practical precepts and advice on how to think better, in the pursuit of knowledge. As it has been developed to date, regulative epistemology is individualistic. It is addressed to individual thinkers, offering them advice on how to gather and collate evidence, how to weigh it, how to avoid error and what warning

signs on the path to knowledge to look out for. Regulative epistemology has been developed most fully by virtue—and vice—epistemologists. While the distinction between regulative and analytic epistemology was introduced by Nicholas Wolterstorff (1996), its currency in debates today is due principally to Roberts and Woods, who subtitle their book *Intellectual Virtues* "an essay on regulative epistemology," and it's to virtue and vice epistemology that I will turn first. Virtue epistemology, in its regulative form, may have a role to play in guiding us toward better belief, but I will suggest that its role is extremely limited.

Virtue epistemology, at least in the "responsibilist" variety that will be our focus here, recenters epistemology around virtues and vices.[1] Virtues and vices are character traits (and, perhaps, ways of thinking and attitudes as well) which are, respectively, epistemically helpful and harmful. Either directly (by making us more or less responsive to evidence or to criticism, for example) or indirectly (by making us love truth or to be indifferent to it, to read widely or to be incurious, and so on) they help or harm our functioning as epistemic agents, especially the extent to which we're able to acquire knowledge. In its regulative guise, virtue epistemology might offer guidance in one or more of several broad ways.

Virtue epistemology in its regulative guide might adopt what Roberts and Wood call "the social engineering model of the regulative philosopher's role" (Roberts and Wood 2007: 29) and aim at the inculcation of virtues in us, for example, by advising us on how we should educate our children or ourselves, or more directly, by trying to shape the virtues of its readers. Alternatively or in addition, the regulative epistemologist might aim to influence our enquiry prior to (perhaps as a means toward) our becoming virtuous. She might urge us to conduct our enquiry in the way a virtuous person would conduct enquiry, or to believe what we

[1] It is this ("responsibilist," rather than reliabilist) variety of virtue epistemology that has sometimes turned regulative. Whereas virtue reliabilism thinks of virtues as reliable cognitive faculties (e.g., perception or memory), responsibilist virtue epistemology mirrors virtue ethics in thinking of the virtues principally as character traits, like generosity or humility. Both varieties of virtue epistemology were initially developed as attempts to address the nest of issues that have occupied epistemologists since Gettier upset the game board in 1963; principally, offering an analysis of knowledge. Reliabilist virtue epistemology's principal architect is Ernie Sosa (see the essays collected in Sosa 1991); Linda Zagzebski played a similar role in the development of responsibilist virtue epistemology (Zagzebski 1996).

would believe were we virtuous, or something along these lines. No matter how it goes about it, I doubt that virtue epistemology (in its regulative form) will be especially helpful to us.

I won't discuss any of these different models for regulative epistemology in any detail. Whatever their value in other ways (perhaps the inculcation of the virtues is intrinsically valuable, though I'm somewhat skeptical),[2] I'll suggest that the virtues are not especially valuable as a means of regulating our epistemic conduct in the service of the acquisition of knowledge. Perhaps the virtues conduce to knowledge, but they do so to a very limited extent, and then only in appropriately structured epistemic environments. Inculcation of the virtues is helpful, if it is helpful at all, not as an *alternative* to scaffolding inquiry and structuring the environment, but only in *conjunction with* these measures.

Inculcating the Virtues

Virtue epistemology, in its regulative guise, aims to improve belief by inculcating the virtues in believers; correlatively, it explains some or many cases of bad belief by reference to a lack of intellectual virtues (or the presence of intellectual vices). Virtue epistemologists offer analyses of specific virtues and guidance on how they manifest, with the aim of allowing us to identify them in others and develop them in ourselves. The intellectual virtues are the cognitive parallels of the ethical virtues. Just as we ought to strive to be and to act out of virtues like generosity, courage and honesty, so we ought to strive to be open-minded, epistemically humble, charitable and so on, and conduct our thinking in a way that manifests these virtues.

Virtue theories, in their ethical and their epistemic guises, face the charge that they fail to be action- (or conduct-) guiding. These high-minded phrases are at best difficult to pin down. It's difficult to give

[2] Elsewhere, Mark Alfano and I have expressed doubts that possession of the epistemic virtues conduces to the generation of knowledge. We argue that much of our most significant knowledge is generated collectively, by agents who individually manifest epistemic vices, rather than virtues (Levy & Alfano 2019).

anything like a precise content to the virtues, at least in a way that's principled and doesn't beg crucial questions. Famously, Aristotle identified the virtues as a mean between two extremes: courage is the mean between cowardice, on the one hand, and foolhardiness, on the other. In doing so, he highlighted the problem: identifying the virtues requires us to engage in interpretative activity and interpretations are easy to contest and difficult to defend. If I'm accused of cowardice in not standing up to that bully, I may retort that that would have been foolhardy, not brave. How can we settle this dispute in a principled way? Virtue theorists often appeal to practical wisdom at this point; that is, to a kind of intellectual capacity to judge cases well. At best, that's little use to those of us who apparently lack practical wisdom. How are *we* to proceed, now? At worst, it might be nothing more than an appeal to the inherent superiority of an interlocutor to put an end to a dispute.

Virtue epistemology faces the same set of problems: the intellectual virtues, too, are at best hard to pin down in a principled way. Virtue epistemologists (and virtue ethicists) have a reasonable response to this worry, however. They may dismiss the wish for anything more precise to guide action and intellectual conduct as a fantasy, insisting that doing and thinking the right thing really just *is* a matter of difficult judgment. They'll point out, moreover, that they don't simply appeal to practical wisdom to end debates. They offer detailed arguments and analyses of the virtues, arguments, and analyses that can bring those of us who currently lack practical wisdom a little closer to understanding the virtues and a little nearer to practical wisdom ourselves.

I'm agnostic on the extent to which virtue epistemology and its analyses are genuinely helpful. If virtue epistemology can help, however, it's not by substituting for apt deference to others and socially distributed cognition; instead, it's by playing a (small) role in helping us to do these things better. Virtue epistemologists may be on to something, but in its current guise their (explicit) recommendations are far too individualistic. Explicitly at least, they appear to aim to bring us each to inculcate the virtues in ourselves and then, guided by our intellectual excellences, to tackle hard problems largely on our own. *That* strategy must fail; or so I'll argue (in fact, many of the examples of intellectual excellence they

themselves cite involve agents whose epistemic success is due mainly to apt deference, and not—much, at any rate—to the deployment of the virtues).

I'll work through a case study of sorts: I'll focus on one virtue epistemologist's attempt to delineate the virtue of open-mindedness from a kind of intellectual "flaccidity," on the one hand, and from dogmatism, on the other, and to provide concrete guidance for intellectual inquiry.[3] I'll suggest that the enterprise is misguided as a central focus of an explanation of bad belief and a remedy for it. Even if we succeed in delineating open-mindedness and inculcating it in ourselves, we won't have the tools we need we for better beliefs. For that, we need other people and effective epistemic institutions much more than we need virtues.

Open-Mindedness as an Epistemic Virtue

The vulnerability of open-mindedness to excess is a part of folk wisdom. We all recognize that there's such a thing as being *too* open-minded. "Be open-minded, but not so open-minded your brains fall out," as the quip has it. We don't think that the virtuous agent will give just as much weight to the health advice of their physician on the one hand and their cousin who posts memes about vaccines on Facebook on the other. They shouldn't be equally open-minded about each.

Perhaps we shouldn't be open-minded at all about those vaccine memes; perhaps we should just dismiss them more or less unthinkingly. Kripke (2011) argues that dogmatism—a paradigmatic epistemic vice—is warranted in the face of certain claims. He gives the example of astrology: "I once read part of a piece by a reasonably well-known person defending astrology […] I was not in a position to refute specific claims but assumed that this piece was of no value" (Kripke 2011: 49). Rather

[3] Admittedly, Cassam, my stalking horse, is a self-described vice theorist, not a virtue theorist. Nevertheless, he seems an apt theorist to discuss under the heading of virtue epistemology. Vice epistemology follows the example of virtue theory in explaining bad belief as the upshot of the character traits and dispositions of agents (Cassam adds "attitudes" to the mental states that constitute vices, but he thinks that virtues too are sometimes constituted by attitudes). Moreover, the epistemic vices are the mirror images of the epistemic virtues: they imply one another and each can be analyzed in terms of an extreme lack of the other.

than being open-minded with regard to every claim we encounter, he suggests, we should "delineate cases when the dogmatic attitude is justified" (Kripke 2011: 49).

As Cassam points out, dogmatism looks very much like a manifestation of the "archetypical epistemic vice": closed-mindedness (Cassam 2018: 39). Kripke advocates refusing to consider evidence and arguments against conclusions he takes to be settled. Cassam argues that such dogmatism is epistemically vicious, and advocates serious engagement, even under conditions like these. Moreover, he suggests, dogmatism threatens knowledge. Kripke had argued that dogmatism can protect knowledge: when we know that an argument leads to a false conclusion, but can't see how to rebut it, we do best to refuse to engage (see Fantl 2018; Levy 2006 for related arguments). Cassam argues that on the contrary, dogmatism undermines knowledge, on the grounds that knowledge requires justified confidence in a belief. If we don't have a right to our confidence, because we maintain our belief dogmatically, we don't genuinely *know* the proposition we believe.

Of course, Cassam accepts that the virtuous agent should be slow to abandon their justified convictions in the face of arguments they can't immediately see how to refute. We should never be dogmatic, he maintains, but we often ought to be appropriately firm in our opinions (open-mindedness is the mean between intellectual flaccidity and dogmatism). Kuhn (1970) influentially argued that scientists are and should be much less open-minded than we tend to think. Scientists are very reluctant to abandon their "paradigms;" the set of findings, exemplars of good scientific practice, methods, and assumptions that for them constitute good science. In the face of an observation that conflicts with a well-established theory, they'll usually reject the observation, rather than the theory (as we saw in the previous chapter). That looks like dogmatism: rejecting a claim or an observation as false, simply on the grounds that it conflicts with what we expected to see.

Cassam accepts that it's entirely appropriate for the scientist to assume that the anomaly can be accommodated by the paradigm (or, alternatively, that it's the upshot of measurement error), even if she can't immediately see how. That's not dogmatism; that's appropriate "firmness or tenacity" (113). What distinguishes firmness from dogmatism is that

whereas the dogmatic scientist would shrug her shoulders and move on in the face of the anomalous, the firm scientist will work to accommodate it or to show that it's spurious. The virtuous scientist is not dogmatic, because she's willing "to acknowledge fundamental flaws in established tools and beliefs, and abandon those tools and beliefs" (113).

Cassam thereby offers us a genuine alternative to Kripke, not a mere relabeling of what the latter calls "dogmatism" as "firmness." Kripke argues that in the face of an argument for a view we know to be false, sometimes we should just shrug and move on. Cassam advocates engagement with these arguments. If we refuse to engage, he maintains, we lose our right to confidence in our beliefs and thereby lose knowledge. One should be confident in one's beliefs, and that requires confidence in our capacities to tackle spurious arguments. I'll argue that insofar as Cassam urges us—ordinary agents, who lack any special expertise in the domain of the argument—to tackle these arguments on our own, he's wrong: engaging with them risks knowledge to a far greater extent than does dogmatism. Reason, argument, evidence, and practical wisdom unaided leave even the fully virtuous person vulnerable. Without heavy-duty social and environmental scaffolding, even virtuous agents can't reliably acquire knowledge about difficult and complex issues. The right response to spurious arguments is often to shrug and move on, relying on others to tackle them for us.

The dispute between views like Cassam's and those more like Kripke's turns in important part on an empirical question. Kripke asserts, and Cassam denies, that it is sometimes very difficult to discover where spurious arguments go wrong; difficult enough that we are more likely to lose knowledge by confronting a spurious argument than by dogmatically refusing engagement. Kripke maintains he couldn't see where the argument in favor of astrology went wrong and implies he was unlikely to be able to identify any flaws quickly had he persisted. Cassam thinks this shows a vicious lack of intellectual self-trust on Kripke's part. *Of course* he could identify flaws: the more dubious a view, the easier it is to debunk. When we're confronted with dubious claims—like our cousin's anti-vaxxer rants on Facebook—we must respond seriously, or we risk

our knowledge. We must identify the flaws, do our own research and (when technical expertise is necessary) consult the experts.

I'll argue that Cassam is wrong on the empirical question: in fact, we're at much greater risk of losing knowledge from "doing our own research" than we are from dogmatism. It *is* true that we are often able to rebut spurious claims, but that's not by probing them for ourselves: it's by apt deference. The intellectual virtues play only the smallest of roles in any of this. I'll also argue that Cassam is wrong in his characterization of the behavior of scientists. Scientists often are (appropriately) untroubled by the anomalous. They don't work to accommodate it. They shrug and go on, dogmatically. If that's how scientists (our paradigms of epistemic excellence) should behave, it's even more the case for the ordinary person. We, too, should often shrug and be dogmatic.

Climate Change Skepticism, Holocaust Denial, and Other Fantasies

One of Cassam's principal aims, in developing vice epistemology, is to enable us to understand the origins and persistence of conspiracy theories and the like.[4] He argues that epistemic vice is an important factor in explaining why people accept these theories. If he's right, then we should find that the epistemic virtues are protective against them: the virtues enable us to see through these theories. Moreover, he's committed to a further claim: there won't just be a correlation between possession of the epistemic virtues and the rejection of conspiracy theories and the like. There'll be a direct causal connection: agents will reject such theories *because* they're virtuous and as a result of *deploying* these virtues in

[4] David Coady (2007) and Charles Pigden (2007) both urge that we drop the pejorative use of "conspiracy theories," because conspiracies are all too often real and we risk unjustifiably stigmatizing conspiratorial explanations by this usage. I have some sympathy for this view. I use the term "conspiracy theory" here somewhat tentatively, to pick out that subset of theories that postulate a conspiratorial explanation of events where that explanation runs counter to the explanation offered by duly constituted epistemic authorities (see Levy 2007 for discussion). Given our pervasive dependence on testimony, rejecting such explanations is usually irrational, even if the explanation is in fact false.

assessing these theories.[5] Cassam is committed to thinking that the epistemic virtues will make a significant difference in our capacity to assess such theories accurately, because they'll enable us to see through them.

We're each exposed to a massive amount of misinformation; many of us daily. Our cousin may post his anti-vaxx memes on our Facebook wall; we may see "plandemic" fliers on the street and the news may report the false claims of certain politicians. If Cassam is right, at least some of this material requires a serious response from each of us: we are able to retain knowledge *that COVID-19 is not a hoax* only if we each confront the claims of those who maintain that it is. How hard can it be?

Many of these claims venture into areas of specialist knowledge. I don't have specialist training in immunology or in climate science. That puts me at an immediate disadvantage when it comes to assessing these claims, especially when they stem (directly or indirectly) from those who do possess specialist training in these areas. I take myself to be reasonably sophisticated and reasonably knowledgeable about climate science, which I've been following for more than a decade. But very often, when I come across sophisticated denialism I find myself in the same position that Kripke took himself to be in reading that article on astrology: I know it's wrong, but I have no real idea how.

Here's an example I happened to come across recently. Visiting the website for Springer publishing (a reputable publisher, for whom I edited a journal for a decade), I was presented with an advert for a book called *The Rise and Fall of the Carbon Dioxide Theory of Climate Change*. The author is one Rex Fleming, who has a PhD in atmospheric science from the University of Michigan and is a fellow of the American Association

[5] Mark Alfano and colleagues have found that the virtue of epistemic humility correlates with lower rates of acceptance of conspiracy theories and lower rates of acceptance of fake news (Meyer & Alfano forthcoming). While this is evidence in favor of Cassam's view, it's correlational: it doesn't show that agents reject conspiracy theories due to the deployment of epistemic virtues. As a matter of fact, I'm skeptical of the causal claim. Meyer and Alfano used a sub-scale from a previously validated epistemic humility scale (Alfano et al. 2017) to measure the virtue. I suspect that this sub-scale is transparent to respondents. The items (sample items: *I don't take people seriously if they're very different from me*—reverse scored, of course—and *I appreciate being corrected when I make a mistake*) are clearly designed to probe dispositions that we, as good liberals, are *supposed* to value. People respond accordingly. Both rejection of conspiracy theorizing and the scale tap into the same underlying dispositions: dispositions to respond in the way that left-leaning, educated Western individuals value. I suspect sociological factors explain the correlation, not the virtues.

for the Advancement of Science. Fleming has had a decades' long career as a scientist and has published work on climate science and on modeling in reputable journals, including the *Journal of the Atmospheric Sciences* (impact factor 3.194) and *Environmental Earth Sciences* (impact factor 2.18). Fleming is also a climate change skeptic: he argues that CO_2 has no impact on global warming.

If Cassam is right, we—you as much as I—now have an obligation to confront Fleming's claims seriously, if we're to retain the knowledge *that climate change is real and CO_2 is a very major contributor to it*. We can't just shrug our shoulders and move on. According to his own website, Fleming's book establishes:

> the failure of the Schwarzschild radiation integrations to maintain the CO_2 longwave radiation intensity achieved in the surface warming by H_2O and CO_2. The resultant Planck radiation intensity is severely depleted in the upper atmosphere. The result is the CO_2 molecules merely pass their remaining small residual heat to space un-impeded.

To retain our knowledge, we're going to have to rebut this claim, show it's irrelevant or that it's gobbledygook. If we can't, we may lose our right to confidence.

Can you do this? I have to confess that I can't (not in the way that Cassam recommends: by deploying my intellectual capacities to grapple with the first-order evidence). I'm at an immediate disadvantage when it comes to the task: I have little idea what the words I just quoted actually *mean*. I'm tempted to dismiss them as gobbledygook (a commentator on the website *Skeptical Science* dismisses Fleming as a purveyor of half-baked physics he learned on Wikipedia), but his track record of publication in reputable scientific journals and his other credentials makes me suspect that something more sophisticated than the mere repetition of ill-understood verbiage is going on. I'm going to have to do better.

How should I go about it? I can turn to Google and search for discussions of Fleming's book (that's how I came across the quote from *Skeptical Science*), but most of this discussion is beyond me. Perhaps Wikipedia would serve me better: there is, I see, an entry for "Schwarzschild's equation for radiative transfer," but even the Wikipedia

entry is very heavy going for me. It's math-heavy, and I'm extremely limited in that area. Even coming to understand what Fleming is claiming—rather than getting on with assessing his claims—is likely to be the work of many hours for me. Without the math, I doubt I'd *ever* be in a good epistemic position to genuinely assess and rebut Fleming. Even with it, I'd likely require hundreds of hours of immersion in the technical literature to rebut the claims of someone with a PhD in the relevant area and a decent track record of publication on technical questions like mathematical modeling. Since much of that literature is currently beyond me, I'll have to start small: maybe with a high-school math textbook (and I'd better get going on the syllabuses in chemistry and physics too). I'm going to have to acquire genuine expertise to rebut Fleming's claims (at least, that is, to rebut them in the kind of way Cassam wants: by identifying for myself the flaws in the arguments).

How much genuine expertise am I going to need? I don't know enough to know how much I don't know or how much ignorance I need to remedy. Would a good undergraduate degree in climate science suffice? Maybe, although the fact that Fleming publishes in the area makes me think that may not be enough. Of course, Fleming is far from the only sophisticated skeptic; the famed 97 percent consensus on climate change is, if anything, an underestimate of the proportion of those with relevant expertise who accept the mainstream view on climate change, but that still leaves room for hundreds of dissenters. If the issues are difficult enough that the experts sometimes get them wrong, what chance do *I* have, working without institutional support and colleagues and acquiring the needed skills along the way?

Let me illustrate the difficulty with a different example, from an area in which the issues are surely less complex. In *Vices of the Mind*, Cassam discusses implicit bias at some length. He concludes his survey with the observation that such biases are malleable and respond to efforts on the part of motivated individuals to change them (173). I happen to possess some degree of genuine expertise on this topic, on which I've published multiple papers (see Levy 2017, 2016, 2015, 2014a, 2014b). I've read many papers on the very topic Cassam here considers: the extent to which our biases respond to efforts at self-cultivation. Yet I remain unsure whether Cassam is right in claiming that they're malleable in the

way he suggests. The literature on implicit biases is large, but it's tiny compared to the multi-disciplinary literature on climate change. The range of expertise required for it is correlatively small. If I haven't been able to answer the implicit bias question for myself, I despair at my capacity to rebut sophisticated climate skeptics.

It's worth adding that the multi-disciplinary nature of climate science (like many other areas of contemporary science) entails that many actual climate scientists may lack the skills to rebut the sophisticated skeptic. Only a minority of climate scientists work on issues like radiative transfer; the rest work elsewhere, on different issues. Some (in my ignorance I have no idea how many) of those who work on different issues will have a working knowledge of radiative transfer, and perhaps that'll enable them to rebut Fleming or rapidly to acquire the capacity to rebut him. Few will waste their time in this way. Few climate scientists want to take time off from their research to answer the arguments of a crank: if those they trust tell them the points raised have been dealt with, they'll shrug and move on. Surely that's all we can reasonably ask of them. If I'm right that the majority of climate scientists would have to take a significant amount of time away from their research to rebut Fleming for themselves (how much time? In some cases, mere hours; in others, weeks or months), then we can't expect each of them to invest this time. They'd have no time for their own research if they had to do this with regard to every crank they encountered.

It may be important that *someone* rebuts the claim (that will depend on its novelty). The great majority of climate scientists will outsource the job, and defer to whoever does it. For the most part, this deference will itself be dogmatic: they won't search for rebuttals. Rather, they'll move on, confident that if the claims are worth engaging with, someone well-placed to do so will take on the task. They'll deal with challenges they take to be worth taking seriously in their own areas instead. They know if the challenge that Fleming raises is genuinely troubling for the field, the news will reach them (it would reach all of us soon enough). In the meantime, dogmatism is justified epistemically and pragmatically.

Climate science is highly unusual, and perhaps taking it as our case study unfairly tilts the odds against Cassam. After all, climate science is highly specialized and highly technical. Many readers, like me, will face

a high barrier to entry into discussions into the area due to its reliance on mathematical modeling and very specialist knowledge. But it's surely not out of place to focus on the single most significant case of bad belief in the contemporary world: if virtue epistemology struggles to handle this case, what's it good for? Nevertheless, perhaps it has a purchase in less technical areas. Cassam's own principal example in *Vices of the Mind* is Holocaust denial. Perhaps on a topic like this, virtue epistemology can make significant inroads, even if it struggles with climate science.

It is surely true that the barriers to entry into discussions of history are lower (at least for likely readers of this book) than into climate science. In academic history, claims are conveyed largely in natural language, not the technical languages of mathematics and statistics, and we're all experts at parsing natural language. We should, however, take care to refrain from the philistinism of those who regard only the sciences as genuinely worthwhile intellectual pursuits. History isn't just story-telling. It has its own tools and techniques and its own experts. We're not experts in history or in adjudicating historical claims merely in virtue of being expert language users. To see through the lies of sophisticated Holocaust deniers like David Irving (on whom Cassam focuses), we need field-specific expertise, where the field is not simply "history" but more precisely twentieth-century German history (or more precisely still, World War Two history or even Holocaust history). Perhaps we can acquire that expertise more rapidly than we can expertise in radiative transfer, but it would nevertheless be a very significant investment of time and energy for any of us.

An incident that occurred just a few months after *Vices of the Mind* was published illustrates just how demanding and specialized historical expertise really is. In May 2019, Naomi Wolf gave a series of interviews to promote her new book *Outrages* (Wolf 2019). This book argues that the persecution of gay men in Britain increased dramatically in the second half of the nineteenth century. Wolf pointed to the occurrence of the term "death recorded" in the records of the trials of men accused of "sodomy." She argued that this indicated an increase in the use of capital punishment by the courts and a correspondingly harsher attitude. She suffered the on-air embarrassment of having her error pointed out by Matthew Sweet, a radio host and cultural historian. As he pointed out,

"death recorded" was used for a nominal death sentence: when the accused was found guilty of a capital offence but not executed.

Wolf's public embarrassment was widely taken to indicate that she hadn't done her research diligently and hadn't used fact checkers. While I'm in no position to assess the quality of her research, prima facie she had every reason to be confident in her claims. The book was based on her 2015 University of Oxford doctoral thesis: it had, therefore, been supervised by an expert (in nineteenth-century literature) and had been examined by independent experts. She had also employed a genuinely expert fact-checker to verify her interpretation of the law. Dame Helena Kennedy, a prominent human rights lawyer, had interpreted "death recorded" in the same way Wolf had (H. Kennedy 2019). Wolf's DPhil and her book were both checked by multiple experts, but that wasn't enough to allow her to avoid embarrassing error. What's needed, in science and the humanities alike, is *specific* expertise: expertise not just on that period, but on that practice, in that place.

Our genuine expertise as natural language users may lead us to miss this fact, and to think we can wade in for ourselves, in the humanities if not in the sciences. But we're liable to be mistaken, and embarrassments like Wolf's may easily result. Wolf took care to have her claims checked; others get themselves into trouble because they approach history casually, certain that their native intelligence and common sense equips them to adjudicate historical claims. 2019 seems to have been a treasure trove of such incidents: in that year the journalist Cokie Roberts took it upon herself to charge historians writing about abortion with distorting the record (Wulf 2019). The historians she targeted had claimed that advertisements for abortion services were plentiful in nineteenth-century newspapers. Roberts denied this, on the grounds she couldn't find any. What was missing wasn't the ads, which were indeed plentiful, but the capacity to identify them by the euphemisms they used. Being a good journalist didn't equip Roberts to call out the historians.

If Roberts' embarrassment indicates that historical expertise is needed to adjudicate historical claims, Wolf's indicates just how specific that expertise needs to be. It's not enough to possess expertise in the period or in the law: one needs expertise in the law of that period. Expertise is *brittle* (Kilov forthcoming): an expert in a particular domain is often

unable to transfer their skills to another, intuitively similar domain. Kilov provides multiple illustrations. Surgical skills not only fail to transfer across surgical procedures: they don't even correlate well with performance on tasks designed to mimic them. Performance on surgical simulations is predicted by practice on the simulator, not the possession of skill at the procedure simulated. Professional-level skill in hitting a baseball doesn't equip batters to hit a softball thrown by an expert. Not only do skills not transfer across intuitively similar domains, they also degrade rapidly in response to changes within the domain of expertise: expert bridge players and accountants perform badly in response to arbitrary changes within their domains; the excellent memory for board configurations and moves of expert chess players vanishes when the configurations or moves are not meaningful to them. Expertise in a very specific domain may provide someone with the *confidence* they'll perform well in an adjacent area, but they may nevertheless lack the *competence*.

Of course, expertise can transfer to some degree.[6] The genuine expert often has an edge over a novice, even when there is some degree of mismatch between the domain of expertise and the domain in which it is applied. A historian of, say, modern Europe is surely better placed than I am to adjudicate between David Irving and Richard J. Evans, whose book *Telling Lies about Hitler* Cassam praises as an antidote to Irving. The historian's expertise surely gives her an edge when it comes to evaluating the arguments and evidence each presents. Often, however, expertise fails to transfer: it's far from obvious that my philosophical expertise qualifies me to adjudicate the Evans/Irving dispute any better than an engineer, a plumber or a tax accountant could. Like Fleming, Irving possesses genuine expertise: his 1964 book on the German V-weapons program continues to be well-regarded. This expertise gives him a capacity to fool the naïve—when they attempt to adjudicate claims for

[6] Episodes like those involving Naomi Wolf and Cokie Roberts are examples of what Nathan Ballantyne (2019) calls "epistemic trespassing," where someone with genuine expertise in one field takes themselves to have sufficient expertise to engage seriously with another. Ballantyne has a somewhat more optimistic take on the transfer of expertise than I do. However, his examples of successful transfer involve relatively simple tasks in an expert domain, and the bar for success was set low: expertise transferred just in case the expert performed better than the novice (rather than performed well). Even by this undemanding standard, expertise is surprisingly brittle.

themselves—and ensures that his historical fantasies require genuine and specialized expertise to rebut.

Of course, Cassam recognizes that we often can't be expected to acquire specialist expertise for ourselves, at least when that expertise is scientific (it's much less clear that he recognizes how demanding it would be to acquire historical expertise). 9/11 conspiracy theories often turn on claims about the melting point of steel beams and the like; Cassam recognizes it would be "unreasonable" (117) to expect ordinary people to acquire the knowledge of physics or engineering required to rebut these theories. Instead, he advocates that we consult the experts and work out which of the competing views is correct. While we may not (always) be able to do our own research, in the sense of grappling directly with the first-order evidence, we may deploy the virtues to choose between competing experts.

Cassam is surely right that something along these lines is and ought to be how we should adjudicate disputes like that between Irving and Evans. But insofar as his prescription requires directly and virtuously adjudicating the *second-order evidence*—the evidence that bears not on the truth of the Holocaust, but rather on who is more reliable on this question—it's still too demanding and too individualistic. We face the same risks of losing knowledge by engaging at this level as we do at the first, and dogmatism remains a better strategy. Dogmatism, here, involves the proper scaffolding of inquiry: relatively unquestioning deference to authoritative sources *because they're authoritative* and not because we've assessed their degree of expertise ourselves. It's because they have the right credentials—primarily because they represent the expert consensus view or are endorsed by duly constituted epistemic authorities—that we should defer, not because we've virtuously probed their track records or their citation indexes, let alone because we've evaluated their arguments.

Cassam holds that when we read Evans' *Telling Lies about Hitler*, we'll see Irving's deceptions "brilliantly exposed" (114). But why should we believe Evans' claims? We're in no position to verify his claims about inconsistency with the historical record for ourselves. The same sort of problem arises with other sources we might check. Cassam advises us to turn to Wikipedia and to Google. As he notes, we'll soon learn that the

British High Court ruled against Irving when he brought suit against the historian Deborah Lipstadt for calling him a Holocaust denier, and that he is widely regarded as discredited. But why should we believe the Wikipedia article, or (if we do believe it) that the High Court reached the right verdict? More googling will turn up all too many other sites that support Irving, laud him as a hero, accuse the historical establishment of a conspiracy against him and cite evidence to back up all these claims. How are we to adjudicate any of this?

As a matter of fact, most of us will trust Evans against Irving, and give little weight to Irving's many supporters (some of whom may turn out to have PhDs, alas). We may take ourselves to be convinced by Evans' arguments, judged on their own merits; alternatively, we may take Evans' word for it, but do so because we've diligently assessed his degree of expertise and judged it greater than Irving's. Perhaps, but there's room for a great deal of self-deception here. Why do *I* find Evans' argument more plausible than Irving's? Is it really because Evans' arguments are better and I'm well placed to recognize this? Or am I swayed by him, rather than Irving, because I'm disposed to accept the consensus view? I strongly suspect that's an important element, for me and for Cassam. No doubt Evans' arguments matter, but I bet I'm more receptive to them because I know them to reflect the consensus view, and I suspect the same is true of Cassam. We accept the claims made by Evans and the judgment of the British High Court in very important part *because these are authoritative sources*, and not because they're claims we're in a good position to assess on their merits, or even because we're in a good position to assess how good they are as sources (*that*, too, is a specialist topic, one on which neither he nor I possesses sufficient expertise for confident and well-justified judgment).

Let's conclude this section by turning to the behavior of working scientists, our paradigm of epistemic success. Cassam, recall, denied Kuhn's suggestion that the responsible scientist was a dogmatist who set aside anomalies rather than attempt to explain them or explain them away. In fact ethnographic and historical studies of scientists indicate that Kuhn was right. The scientist often just shrugs in the face of the anomalous, setting anomalies to one side and trusting in the march of science to accommodate them (or to conclude that they must reflect error, since

they aren't replicated). Consider the response of Darwin and his supporters to the work of Lord Kelvin on the age of the Earth and the Sun. Kelvin's genuinely expert estimates of the age of each were vastly too short for the evolutionary account of the origin of life to be plausible. Darwin and his supporters recognized this fact, and recognized that they were unable to refute Kelvin. But rather than following Cassam's advice, and standing ready "to acknowledge fundamental flaws in established tools and beliefs, and abandon those tools and beliefs" (113) in the face of anomalies, they simply set them aside and continued with their work. They relied on the march of science and the work of others, in other fields, to vindicate them, which it duly did later (C. Lewis 2002).[7]

Darwin and his followers were right to set aside the problem, because their research program was so successful in its own domain. This was true even though they recognized that Kelvin's work produced genuine evidence conflicting with their program. Of course, Cassam is right in thinking this strategy has its risks: sometimes the anomalous really is an indication that the research program is flawed. The dispute over the cause of stomach ulcers is an excellent recent example. The primary cause was long believed to be stomach acid, often linked to stress. Doctors were sufficiently convinced of the truth of this theory that some supported fines for those who advocated a rival, bacterial, theory (Zollman 2010). But the bacterial theory was in fact true and the dissidents were eventually vindicated. This little tale doesn't have the moral that Cassam might hope to draw from it, though. It doesn't show that we shouldn't be dogmatic in the face of anomalies. Anomalies are cheap and plentiful. The scientist can't abandon her research whenever she hears of one; that would mean abandoning her research forever. If she is to hang on to her knowledge, she'd better be able to respond by shrugging

[7] In *The Knowledge Machine* (2020), Michael Strevens recounts some of the history of the dispute between Kelvin and Darwin's defenders (Strevens focuses on Kelvin's direct evidence concerning the age of the Earth; Kelvin also argued that the Sun was too young for evolution to be plausible). This episode in the history of science is one of many that Strevens details in which both sides dogmatically (though reasonably) stuck to their guns in the face of evidence they couldn't satisfactorily explain. As Strevens puts it, "Science is driven onward by arguments between people who have made up their minds and want to convert or at least to confute their rivals. Opinion that runs hot-blooded ahead of established fact is the life force of scientific inquiry" (79).

her shoulders and setting aside the many contrarian views she hears expressed every day.

If these reflections are on the right track, virtue epistemology goes wrong to the extent to which it suggests that we can and must secure knowledge by individual cognition. Of course, individual cognition *is* powerful and we ought to deploy it, to the extent we can, in our own fields; perhaps the virtues are an important element in doing so. No matter how virtuously we conduct enquiry, however, we can't rely on it alone reliably to sort out fact from fiction and from fake; even in our own fields, we remain heavily dependent on scaffolding of all kinds for epistemic success. I see no reason to think that the argument won't generalize well beyond virtue epistemology. We non-experts can't hope to rebut the climate skeptics or the Holocaust deniers for ourselves; not without becoming experts (and that requires a great deal of time and effort). Even genuine experts often can't rebut frauds and cranks, if they lack field- and topic-specific specialized knowledge, and those who possess the precise knowledge they need themselves owe their epistemic success to a great deal of scaffolding. By ourselves, none of us are all that epistemically impressive.

Dissent in a Time of COVID

Let me finish this chapter by discussing an apparent counterexample. As I write, the world is in the grip of the COVID-19 pandemic. Governments in most countries claim that their responses are "led by the science" (Peck 2020). Yet very many people, scientists or not, have felt qualified to question that response. Suddenly, everyone is apparently an epidemiologist. Of course, the mere fact that people with no prior expertise feel able to make confident pronouncements at odds with those of (apparent) experts is by itself no surprise. There's nothing new in that. But there may be good grounds for thinking that the pandemic warrants a less deferential response than other areas in which scientific advice guides policy.

Right now, it's not obvious (to me) what the right response to the pandemic is. Most governments have opted for strict lockdowns and strong social distancing measures to reduce transmission, with the aim either

of "flattening the curve"—ensuring that intensive care units are not overwhelmed—or actually eliminating the virus from the population. No one denies that these measures will save the lives of people who would otherwise have died from the virus. Some, however, worry that the economic and social costs of the measures might be greater than the benefits. These are not costs to be weighed *against* the costs to health and well-being: they are *also* such costs. The recession that has resulted from the shutdown of much of the economy across large parts of the world will itself be deadly. The recession that followed the 2008 financial crash is estimated to have led to at least 10,000 extra suicides in Europe and North America (Reeves et al. 2014) and more than a quarter of a million extra cancer-related deaths in OECD countries (Maruthappu et al. 2016). Feelings of isolation linked to the lockdown imposed in many countries will also take a toll on mental health (S. K. Brooks et al. 2020). The economic impact, and therefore (in all probability) the impact on mortality and morbidity of the COVID recession is likely to be much greater this time round. On this kind of basis, it has been suggested that the current response might be more costly than can be justified (Ioannidis 2020). Instead, critics argue, we should collect more data before we settle on a response.

Of course, governments can't wait for more data to come in when they confront a crisis. In this kind of situation, it's widely held they should err on the side of caution, which has been interpreted as entailing that they should lock down hard now, in advance of the evidence. Again, it's not clear to me right now that's right, because it is not clear to me which side is the side of caution. Assessing that issue requires an assessment not just of the health effects of the pandemic, but also of the shutdown. The modeling on which government policy has been based makes assumptions (about infection fatality rates, for instance) which are evolving as we learn more, and has paid little attention to the health effects of the shutdown.

It may be that the actual, highly restrictive, response is partly due to what we might call the goalkeeper's fallacy. There's evidence that penalty kicks aimed down the middle of the goal are less likely to be saved than those aimed to the left or to the right (Chiappori et al. 2002). Part of the reason for this is that goalkeepers usually dive to the left or right to

attempt to save the penalty. They don't stay upright, because they believe (possibly rightly) that they will be blamed less if they made a spectacular and demanding, if futile, attempt to save the penalty than if they engage in the less spectacular strategy of guarding the center. They have an incentive to dive, even if diving is less successful, on average, than not doing so. Similarly, governments may have an incentive to engage in spectacular interventions in the face of a public health crisis. The penalty, in terms of public opprobrium, for underreacting might be much greater than the penalty for overrreacting.

Again, I'm not taking sides here on what the right response should be. I don't know. But even expressing doubts in this way seems quite different to the kind of response I've urged we take to science and scientists generally. What happened to the deference I've argued we should display?

I don't think the pandemic is a counterexample to my claims. The first thing to note is that in expressing doubts about the appropriateness of governments' responses, I'm not expressing any doubts about *epidemiology*. Identifying the appropriate response to the pandemic is not a matter for epidemiologists alone: rather, it's a policy question, on which multiple different kinds of expertise bears. Epidemiologists are not experts in economics or in mental health or social policy or politics or behavioral science, and all these disciplines—and more—are relevant to the right response. To the extent that governmental policy is guided predominantly by epidemiology, it's permissible to worry that it reflects only some of the relevant expertise. Second, much of the debate concerns modeling, and modeling is not the province of any one discipline. People with a range of backgrounds may be qualified to weigh in on models, and those with yet other backgrounds may be qualified to weigh in on their assumptions.

Of course parallel points hold true for climate science as well. The expertise of many different disciplines (economics, sociology, anthropology, political science, even psychology) is relevant to how we should respond to the climate crisis. Why should we be more deferential with regard to climate science than with regard to the pandemic? The central reason is that climate science is mature in a way that our thinking about *this* pandemic is not. This virus differs—in infection fatality rate, in the kind of burden it places on healthcare systems, in the profile of those

who are especially vulnerable to it—from other viruses, and the context of the pandemic is dramatically different from the context of previous outbreaks. Previous pandemics occurred in a world that was vastly different from ours. Think, for instance, of the ways in which the internet has made lockdowns much easier to implement than previously: many people can work or study from home, we can order food to our door, we can entertain ourselves and reach out to others, all very much more easily than ever before. So there's no body of evidence and expertise ready-made to bear directly on *this* pandemic. In contrast, climate science presents us with a consensus which has *already* been tested and retested multiple times, for several decades. A great variety of experts from a great variety of disciplines have *already* contributed to the climate consensus. It's not because the coronavirus is different from climate science that it's appropriate for people to second guess the science. It's because it's the same: there was a time when such second-guessing was appropriate for climate science too. That time has long passed.

There's an important lesson here. A scientific consensus is reliable when it has been stress-tested, by all the disciplines relevant to the topic, for an extended period of time. Only under these conditions is the consensus reliable. Any consensus on the pandemic doesn't meet these conditions. As Schliesser and Winsberg (2020) put it, "there is currently no well-ordered scientific community studying COVID-19 and its impact, so the emerging consensus could be the result of any number of all-too-human biases." These differences between climate science and the state of knowledge over COVID-19 make an epistemic difference: there's no properly generated consensus to defer to in the latter case.

All that said, I'm skeptical that the pandemic is a case in which any of us does better epistemically by making up our own minds. Individual cognition is limited and biased, for reasons that are by now familiar. At this point in the development of knowledge, we may appropriately contribute to the establishment of a consensus through stress-testing, but for each of us it's very unlikely that our considered view is better than that of the epidemiologists advising governments (say). Even in this case, and in the absence of a justified consensus, almost all of us probably do better by deferring than by dissenting—though here the state of knowledge as a whole may benefit if we dissent.