5
Epistemic Pollution

In the last chapter, I argued that individual cognition, even the careful ("virtuous") cognition of those who possess genuine expertise, is a lot less powerful than we tend to think. In fact, it may actually be worse off than the forgoing discussion suggested. I've been arguing that knowledge is dependent on the social and institutional context in which beliefs are acquired and transmitted. But I've paid little attention to the properties of the actual contexts we find ourselves in. We live in *epistemically polluted* environments: deliberately and inadvertently, other agents shape our environments in ways that leave individual cognition even worse off than it might have been. In this chapter, I'll sketch some of the pollutants and how they work to undermine virtuous cognition. The epistemic world has been allowed to degrade, I'll suggest, because we've been unaware of how crucial it is to rational thought. Just as we urgently need to repair and to manage our natural environment, I'll argue, we must repair our epistemic environment.

My focus will be on the so-called "novice-expert problem"; the problem of identifying a genuine or reliable expert among those taking conflicting stances on an issue within their sphere of (apparent) expertise. As Cassam recognizes, solving this problem is essential for laypeople if they are to be able to come to justified views on many important issues. I'll argue that this isn't a problem we should expect novices to solve when they live in epistemically polluted environments.

Novices and Experts

A number of philosophers have risen to the challenge of identifying criteria that ordinary people might use to distinguish reliable experts from unreliable (E. Anderson 2011; Blancke et al. 2017; Johnny Brennan

Bad Beliefs: Why They Happen to Good People. Neil Levy, Oxford University Press. © Neil Levy 2022.
DOI: 10.1093/oso/9780192895325.003.0005

forthcoming; Goldman 2001; Guerrero 2017). While there are important differences between them, they converge in identifying *credentials*, *track record*, *argumentative capacity*, *agreement with the consensus*, and *intellectual honesty* as criteria by reference to which we can choose between experts.

Genuine experts have good *credentials*. They have PhDs in the topic under discussion or in a closely related field. They have published peer-reviewed research in the field. Experts with an especially high degree of credibility set the agenda for their field, as reflected in their citation count, and are honored by their peers (E. Anderson 2011). They also have good *track records*, where "track record" consists in more than peer-reviewed publications. It also consists in a record of making predictions that have been borne out by events. Whereas scientific expertise is esoteric knowledge, whether predictions about future events come to pass is often publicly observable and therefore exoteric knowledge (Guerrero 2017).

Argumentative capacity consists in more than debating skill (which can dissociate from genuine expertise). Rather, genuine experts display what Goldman (2001) calls "dialectical superiority." One expert displays dialectical superiority over another if the first expert is able to rebut the claims and arguments of the second. *Intellectual honesty* is displayed by making data available to other researchers, retracting claims that have been refuted and declaring conflicts of interest; because people may be biased, we should heavily discount those experts who have an interest in the truth of their claims. Finally, an expert should be accorded greater credibility to the extent to which her claims are accepted by a *consensus* of her peers.[1]

[1] Goldman (2001) has influentially argued that consensus may not be a good guide to credibility, because the different sources for a claim may not be sufficiently independent of one another. A *non-discriminating reflector* holds whatever opinion their "guru" holds, regardless of its plausibility, and therefore their agreement adds no independent epistemic weight to the initial opinion. In the actual world, agents are never or almost never non-discriminating reflectors. Even young children filter claims for plausibility, and will reject testimony from a familiar person, even a parent, in favor of more plausible testimony from an unfamiliar informant (see Harris 2012). The degree of independence of individual informants from one another varies from case to case, but we can be confident that each filters testimony for plausibility to *some* degree. Of course, experts may nevertheless defer excessively, without being genuinely non-discriminating. Coady (2006) argues that this kind of excessive deference is rare.

While they all recognize that there are obstacles to utilizing these heuristics, the philosophers who have identified these markers of expertise largely accept that ordinary people are able to deploy them to identify genuine experts. I think they're far too optimistic. Ordinary people are well aware that these criteria pick out markers of expertise. But they're also well aware that we live in epistemically polluted environments, and that a major source of epistemic pollution consists in the mimicry of these markers, to inflate the *appearance* of expertise (Guerrero 2017). Our epistemic landscape is polluted, because the cues for expertise don't correlate well with its actual possession. This fact greatly reduces ordinary people's capacity to distinguish reliable from unreliable sources. At the same time, the fact that such deception is widely known to occur reduces trust in legitimate sources.[2]

Merchants of Doubt (Oreskes & Conway 2011) describes a cavalcade of examples of the mimicry of expertise in the service of science denial. Beginning in the 1950s, the tobacco industry responded to compelling evidence that smoking caused cancer by attempting to sow doubt on the science. It aimed not to present an alternative case or to refute the mounting evidence, but to leave ordinary people confused about who and what was reliable. As the infamous 1969 industry memo mentioned in Chapter 1 put it, "Doubt is our product, since it is the best means of competing with the 'body of fact' that exists in the mind of the general public." As Oreskes and Conway document, the tactics the tobacco industry pioneered subsequently spread to those who sought to cast

[2] I suspect markers of expertise are actually more useful to experts themselves than to laypeople. As we've seen, scientists are routinely in the position of needing to rely on the work of other scientists, without being able to assess their work for themselves. They are, however, able to utilize markers of expertise better than laypeople, because they know which journals are predatory, how to assess citations and h-indexes, the quality of particular departments, and so on. In saying this, I take issue to some degree with recent claims by Konrad Talmont-Kominski (2020). Talmont-Kominski argues that in science, the role of source vigilance is very much attenuated compared to other domains. That's not because scientists don't take anything on trust—quite the opposite. Rather, it's because trust is high, Talmont-Kominski suggests. I think the picture is somewhat more complicated. It might be true that the closer the claim is to the scientists' very specific area of expertise, the smaller the role of source vigilance. Nevertheless, source vigilance remains important to scientists. Talmont-Kominski points to the use of double-anonymous review in science as evidence of the bracketing of source vigilance. But that's a revealing mistake: in fact, science uses single-anonymous review much more often than many other areas of inquiry, and its use is often justified on the basis of explicit appeal to the need to know the source of a claim (Palus 2015; Walker & Rocha da Silva 2015).

doubt on ozone depletion, the viability of Reagan's "Star Wars" project and—most perniciously and harmfully—climate change. Indeed, these tactics have spread much further than that. It's not only the well-funded and coordinated industry groups that are the focus of *Merchant of Doubt* that pump out epistemic pollutants. So too do cranks and frauds across the ideological spectrum, whether to convince us to buy their jade eggs or healing crystals, or to believe their theory debunking Einstein.

Charlatans and their fellow travelers employ a variety of tactics to mimic credibility.[3] For instance, those with an interest in deceiving the general public may set up parallel institutions that seem to guarantee expertise. Oreskes and Conway recount how denialists set up the "Nongovernmental International Panel on Climate Change" to peddle myths about global warming. The NIPCC produced reports in identical formats, with identical sections, to those of the IPCC, with the aim of spreading doubt. The American College of Pediatricians utilized similar tactics. The ACP was set up by a small number of right-wing pediatricians to promote their views. Doing so is surely permissible: what's less permissible (and probably intended) was the effect of muddying debate by misleading people into thinking that the college spoke for the profession. When the ACP issued a statement condemning gender reassignment surgery, many people mistook the statement for the consensus view of pediatricians. But the peak body for US pediatricians, the American Academy of Pediatrics, has a much more positive view of gender reassignment surgery (LaCapria 2016). Insofar as the larger organization, with a broader membership base, can be expected to reflect a broader range of expert views and a higher degree of expertise, it is reasonable to give its views greater weight than those of the smaller organization. When the ACP allows or encourages the impression that it speaks for the profession, it introduces an epistemic pollutant.

Fake and dubious journals are also epistemic polluters. The recent growth in predatory publishers, who publish low-quality scientific research for a fee, has attracted a great deal of attention. But the

[3] These fellow travelers, as I label them here, are not seeking to deceive us. They may not even seek to *mimic* markers of credibility. Rather, they may see themselves as creating parallel (but genuine) epistemic institutions and outlets. Nevertheless, their efforts result in the introduction of pollutants into the epistemic environment.

phenomenon isn't new. *Merchants of Doubt* provides several examples: for instance, the climate denialist *Journal of Physicians and Surgeons*. This journal, and its predecessor the *Medical Sentinel*, also published articles questioning the link between HIV and AIDS, the consensus on DDT and papers alleging that abortions are much riskier than the consensus view maintains. For an even more egregious example, consider the case of pharmaceutical companies cooperating—conspiring?—with the publishing giant Elsevier to produce promotional materials designed to mimic peer-reviewed journals (Grant 2009). These fake journals were able to leverage the prestige of Elsevier to give the "research" they published an air of reliability. When the deceit was uncovered, however, the effect was just the reverse: rather than make the findings published look more legitimate, the deception made Elsevier—and by extension, academic journals—look less legitimate.

Predatory journals have the same effect: reducing confidence in the entire system. Even those who work in academia are sometimes unsure whether a particular journal is legitimate or not, and there are genuine borderline cases. For example, the *Frontiers* stable of journals appears (to me) to be legitimate, despite the fact that authors are required to pay a publication fee.[4] But some *Frontiers* journals appear questionable. *Frontiers in Public Health* controversially published articles linking vaccines and autism (Chawla 2016) and questioning the link between HIV and AIDS (Ferguson 2015); the first was subsequently retracted while the second was reclassified as "opinion." Perhaps in response to these incidents, the librarian Jeffrey Beall decided to add the publisher to his influential (but now sadly defunct) list of questionable journals (Bloudoff-Indelicato 2015). The controversy surrounding Beall's decision indicates how difficult such judgments are even for professionals. If it's hard for academics with expertise in the relevant fields to assess whether a particular journal or a particular publisher is legitimate, we can't reasonably expect ordinary people to make such judgments. If their confidence in scientific findings is lowered across the board as the result of such epistemic pollution, we can hardly blame them.

[4] Since conflicts of interest are a reason to discount expertise, it is incumbent on me to note that I have published in *Frontiers* journals on several occasions.

Epistemic pollution may be emitted by legitimate institutions of knowledge production, as well as from those who mimic such institutions. For example, it may arise from attempts to game systems that are supposed to track expertise. Institutions like universities, the bar association and peer review have as one of their functions the certification of expertise. But they have other functions, too, and these functions sometimes conflict. From this conflict, pressure to inflate credentials can arise. For example, universities have a financial incentive to overstate the expertise of their academic staff (thereby increasing their rankings, and attracting grant money and students). Systems that assess expertise may therefore be manipulated. There are many such cases: for instance, a Malaysian University was recently revealed to be urging faculty members to cite one another to boost citations (McCook 2017). Institutions, including the most prestigious, may be slow to investigate accusations of fraud or try to keep its discovery in-house, to protect the university's reputation.

To these sources of epistemic pollution we can add problems internal to the conduct of science, some of which have recently been widely publicized. Consider the so-called replication crisis, which we briefly discussed in the introduction. While much of the publicity to date has focused on social psychology, many of the problems seen in social psychology are just as common in other disciplines. For example, *publication bias* and the *file drawer effect* are certainly and notoriously problems in medicine. Publication bias is a kind of distortion in what gets published. It occurs when journals are more likely to publish certain kinds of finding than others, even though the intrinsic scientific merit of the favored kinds don't warrant the bias. Findings might be published because they are surprising, or because they're on certain topics (as Kitcher (1987) notes, it is far easier to publish work on human sexual behavior than on less exciting topics, and standards are accordingly lower). Perhaps the single biggest source of publication bias in science is a bias in favor of positive findings. Journals are full of papers that report that there is a significant correlation between two variables (framed as suggestive of a causal relation between them), or that a particular intervention significantly reduced the incidence of some pathology, and so on. Some of these findings are due to chance or would evaporate were

some factor controlled for; others are the result of what have come to be called questionable research practices (John et al. 2012) that torture the data for significance. These findings may go uncorrected due to publication bias: because papers that fail to replicate the finding are less likely to be published, or are published in less prominent places and fail to attract attention (see Fine 2013 for a discussion of how this dynamic plays out to amplify claims of sex differences and downplay evidence against such differences).

Publication bias may affect not only what gets published but also what research is conducted in the first place. Knowing that a failed replication will struggle to be published at all, and that if it is published the venue will not be high profile (and therefore will do relatively little to advance authors' careers) may discourage researchers from undertaking such research at all. For similar reasons, researchers may decide not to further pursue research when initial results are negative. This results in the file drawer effect: when negative results are filed away rather than submitted for publication. More pernicious still, researchers may repeat experimental protocols until they get the results they wanted. Selective publication of positive trials and suppression of negative findings may lead to an overestimation of the efficacy of new treatments (which may in fact be no better or even worse than currently accepted treatments). Unsurprisingly, this is a more common problem in industry-funded trials than in those conducted independently of industry (Every-Palmer & Howick 2014).

Industry funding is a general and central source of epistemic pollution. The tobacco companies, for instance, spent millions on funding research by university scientists. Surprisingly perhaps, the research they funded was often (though far from always) legitimate. The aim was often not to produce spurious findings but instead to draw attention away from tobacco as a cause of cancer by highlighting genetics, indoor pollution, and a host of other (in fact genuine but rare) causes. Industry funding also had an added benefit for the tobacco companies: the production and promotion of a cadre of (genuine) experts who were friendly to the industry and who could be called upon to testify in public forums in its support (Oreskes & Conway 2011).

Recent modeling work has demonstrated how powerful the promotion of genuine science can be in spreading epistemic pollution. Well-conducted science produces findings with a predictable statistical distribution: though good experiments are a reliable means to discover the truth, sometimes even the best experiments will throw up spurious results (that's why replication and supportive evidence from other kinds of scientific work are essential to the interpretation of science). Propagandists can take advantage of this fact to promote spurious findings.[5] Modeling work by O'Connor and Weatherall (2019) shows that even on the assumption that propagandists fund or selectively promote only genuinely well-conducted (but misleading) science, policy makers who attend to them can come to be more and more strongly persuaded of a false view, even as the scientific community converges on the truth; this remains true even if the policy makers also receive information directly from representatives of the scientific consensus.

Identifying Experts in a Polluted Environment

Goldman, Anderson, and other writers are optimistic that ordinary people can identify experts, using the criteria they set out. I think their optimism is misplaced. The epistemic pollution identified in the previous section makes the task of distinguishing reliable from unreliable sources too difficult for ordinary people to reasonably be expected to accomplish it.

The markers of expertise can play their certifying role only if they are not themselves excessively polluted. But these markers are polluted and they're known to be polluted. Ordinary people know that universities don't merely certify expertise. They know that universities also aim to attract funding and to manage public perceptions, and that these aims may conflict. Ordinary people know that peer review is conducted by people with their own interests and biases. They may reasonably (if

[5] Since small datasets are more likely to generate false positives and false negatives than large, the propagandists can ensure best bang for their buck by funding a larger number of smaller trials rather than the reverse (O'Connor & Weatherall 2019).

usually wrongly) conclude on that basis that certain views are not getting a fair hearing. Recall the example of the bacterial origin of stomach ulcers discussed earlier in this chapter. The medical community was slow to give the evidence due weight. Ordinary people may therefore wonder what other dissenting views don't pass peer review, due to the bias of reviewers, or what dissenting hypotheses are not investigated because granting agencies won't fund them. What should we make of the fact, for instance, that Michael Behe's own university department posted a disclaimer on their website, disavowing his "intelligent design" (a theory widely regarded as creationism in scientific dress)? Such a disavowal is predicted both by the view that Behe's claims are not well-supported by evidence, *and* by the view that scientists close ranks against dissenters.

Ordinary people assess expertise in an epistemically polluted environment, in which fakes, flakes, and frauds are promoted by merchants of doubt, by commercial interests and by a media in thrall to "balance" and to the need for sensation. They look to the markers of expertise to certify it, but these markers are themselves regularly manipulated. Moreover, they're aware that *all* sides—legitimate and illegitimate—are subject to extrinsic pressures. These worries affect every marker of expertise. Take track record. Cassam argues that David Irving's track record of misrepresentation ought to alert us to his deceptions. But track record is very often intrinsically difficult to assess. As we saw, Guerrero (2017) advises us to look to the *exoteric* record: predictions (or retrodictions) made by an expert that can be verified by the non-expert. Of course, sometimes experts (or putative experts) make predictions that can be easily verified or falsified. One well-known example concerns the neoconservative political pundit Charles Krauthammer. In response to the failure of the US military to find evidence of an active weapons of mass destruction program in Iraq—the ostensible existence of which had formed the central plank of his case for the war—he noted that the team had had only five weeks to find the WMDs. "Come back to me in five months. If we haven't found any, we will have a credibility problem," he wrote. He thereby provided an exoteric test for his credibility; a test he failed badly (Farrell 2013). But exoterically assessable predictions and retrodictions are the exception, not the rule. In many areas, what exactly

is being predicted and how it would be falsified is often hard for the non-expert to assess.

In specialist science, it often takes specialist knowledge to understand just what is being predicted, let alone to verify the prediction. Climate denialists, for example, seem to be committed to the prediction that global temperatures will not correlate with concentrations of CO_2. Climate scientists will tell you that this prediction has been falsified: as a matter of fact, temperature rises are well correlated with CO_2. But the denialists have a response, or a number of responses. They may maintain that the apparent correlation reflects manipulation of the data, rather than genuine change. It has been alleged, for example, that the National Oceanic and Atmospheric Association tampered with temperature data inconvenient to the warming narrative (Richardson 2017); that the Climate Research Unit at the University of East Anglia deleted its data to hide anomalies, and so on. All of these claims have been multiply debunked, of course. The interested reader would do well to consult the blog *Skeptical Science* for both entry-level and advanced discussions of all these myths.[6] But these debunking efforts have themselves been met with (attempted) debunking.

Liberals tend to think of denialists as ignorant, unintelligent or hucksters. In fact, sophisticated denialism is easy to find. *Watts Up With That*? describes itself as the "The world's most viewed site on global warming and climate change" and promotes denialism with (apparent) facts. Judith Curry, a climate scientist with a solid track record of well-regarded publications, uses her scientific skills to promote doubt on her own blog (as well as the blogs of others, in the media, and in front of US House committees). Assessing her claims, as well as those of other sophisticated denialists, is far beyond my capacities. Are the empirical claims she makes true (e.g., that a particular technique has been misapplied, or that there are statistical errors in a paper)? Assuming they are true, do they support to some degree her skepticism regarding the scientific consensus? While I find the responses on *Skeptical Science* more plausible than her posts, it's likely that my disposition to defer to those on my side helps explain that fact. I'm confident this point generalizes:

[6] https://skepticalscience.com/

our being swayed by the arguments is due in part to our responding to other cues. I'm not of course claiming that there's no fact of the matter here, or that one side doesn't have very much better evidence than the other. I'm suggesting, rather, that *we* (non-specialists) lack the capacity to identify which side is right by reference to argument quality alone (or, typically, even primarily).

Argumentative capacity—the possession of what Goldman (2001) calls "dialectical superiority"—fares no better as a marker of expertise (D. Coady & Corry 2013). The ability to rebut arguments and the *appearance* of having this ability may dissociate. As many scientists who have debated creationists have learned to their cost, well-rehearsed debaters can seem to neutral audiences to be dialectically superior by having an apparent response to every objection, even if the response is only smoke and mirrors. They can also appear to evince dialectical superiority by raising so many objections and making so many points so quickly that their opponent is unable to rebut more than a small fraction of them (this is known as the Gish gallop, after a creationist who specialized in the technique). Reference to dialectical superiority enables us to distinguish those who have spent a lot of time on a topic from those who haven't, but it's insufficient to allow us to distinguish genuine experts from pseudo-experts who have also spent a great deal of time on the topic.

Intellectual honesty fares no better, because it's not appropriately independent of the issues disputed by the experts themselves. All sides accept, with Anderson (2011), that a putative expert acts dishonestly if she doesn't withdraw claims that have been refuted. But the fakes and fellow-travelers hold (sincerely or duplicitously) that it's the genuine experts who are intellectually dishonest because it is *their* claims that have been refuted. Similarly, accusations of conflict of interest are often unhelpful, because (as Guerrero notes) such conflicts typically appear on all sides. It is of course common for anti-vaxxers to accuse their opponents of being in the pockets of "big pharma," and climate denialists cite the attractions of grant money to explain the appearance of consensus among scientists. Both can point to genuine scandals in the relations between scientists and pharmaceutical companies, such as the phenomenon of medical ghost-writing (Langdon-Neuner 2008), where

a prominent physician or researcher puts their name to a paper that has been largely or even entirely written by company representatives.[7]

Claims concerning the existence of a consensus on a topic are also of limited help, insofar as such claims also fail to be appropriately independent of other issues. If credentialing bodies will not grant PhDs to dissenting researchers, for instance, we should expect to see a consensus of appropriately credentialed scientists on a topic, regardless of whether the consensus is well supported. If data that conflicts with the consensus view is suppressed (deliberately or just because it's difficult to publish), the consensus will not have much evidential value. This point is a generalization of Goldman's claim that the concurrence of additional experts with a claim adds no additional evidential weight to it unless they are sufficiently discriminating in what they believe. Goldman worries about excessive deference to opinion makers, but there are other ways in which an unreliable consensus could be generated. If institutions that grant credentials use inappropriate criteria in assessing expertise, the resulting consensus will not be truth conducive.

Claims of intellectual dishonesty are also symmetrical. Climate scientists routinely (and to my mind rightly) accuse some of the denialists of deliberate deception.[8] Sometimes, persuasive evidence of such deception emerges, when memos and emails never intended for public consumption comes to light (such as the tobacco industry memo that provided *Merchants of Doubt* with its name). But such revelations are rare, and *their* side thinks that they have such smoking guns too. The most famous here is "climategate." In November 2009, a server at the Climactic Research Unit at the University of East Anglia, was hacked

[7] One of the clearest cases of intellectual dishonesty in recent medical history is surely the Andrew Wakefield story. In 1998, Wakefield and his co-authors published a paper alleging a link between the MMR vaccine and autism. After other researchers failed to replicate his findings, Wakefield was found to have undisclosed conflicts of interest. The British General Medical Council then investigated further and found a litany of other problems, from performing unnecessary and invasive procedures on children with autism to suppressing data. The paper was retracted and Wakefield was struck off the medical register. Those who trust the relevant institutions will take Wakefield to be discredited and his research invalidated. But if you are disposed to distrust these institutions, you might see them as closing ranks against a brave truth-teller.

[8] Only some. Some of the dissenters are sincere, and some of the sincere dissenters are competent. Indeed, their dissent may be explained, in part, *by* their competence: their mistake (or one of them) is to overestimate the powers of individual human reason.

and a trove of documents and emails stolen. Excerpts from the emails were subsequently published on denialist blogs, which alleged that they show that researchers were fabricating and manipulating data to support their political line. Denialists seized on one email in particular, in which Phil Jones (a leading climate scientist) said he used "Mike's *Nature* trick…to hide the decline." Here was the smoking gun! The mainstream media reported the revelations with varying degrees of credulity: for the *Telegraph*, for instance, it was "the worst scientific scandal of our generation" (Booker 2009). Three separate inquiries found no evidence of scientific misconduct. Unsurprisingly, the denialists allege the inquiries were themselves fraudulent. What would you expect when the scientists close ranks?

For each claim by a scientist or a group of scientists, it seems that there is a rebuttal by an opponent, a response to that rebuttal and a further response to it in turn. Good luck keeping up! While one side may feature better credentialed experts than the other, such a pattern of distribution is exactly what one would expect if the better credentialed side suppressed dissenting research (such suppression might be conspiratorial, but it need not: it could even be produced by well-meaning but biased scientists trying and failing to give their opponents a fair hearing). Claims of intellectual dishonesty abound on both sides, but for the most part they don't help, because the accusations are symmetrical and we can often adjudicate the claims only by adjudicating the first-order issues on which they turn. We can't, therefore, utilize concerns about intellectual dishonesty to identify reliable experts: the criteria aren't sufficiently independent of one another.

The Efficacy of Epistemic Pollution

So far I've been long on argument and assertion about the effects of epistemic pollution, and short on empirical evidence. We are, as I've already noted, epistemic individualists, and we tend to be confident of our intellectual powers. Readers of a book like this one are particularly likely to have a high (and probably well-founded) opinion of their capacities.

Surely I exaggerate the degree to which epistemic pollution is an obstacle to belief? Surely you (dear reader) can, with sufficient effort and application, sort through the lies and the fog, and come to an accurate assessment of the evidence?

You are (very probably) in a much better epistemic position than most people. It's not just that you are well-educated and (again, very probably) more intelligent than average. It's not just that you probably have research skills that most people lack. You are also (very probably) epistemically luckier than most. As a consequence of your socialization (from family through to prestigious academic institution), you have acquired dispositions to trust reliable sources. You know enough to distinguish legitimate institutions from diploma mills; you have some idea of the degree of legitimacy conferred by a publication in *Nature* or *Science*. You are alert to signs of predatory publishers and on the lookout for industry funding. You are therefore protected, to some degree, from epistemic pollution.

For all these reasons, you're indeed more likely than most to get things right when you (attempt to) judge for yourself. But that's not because you're a counterexample to my claims: it's because you fit my model so well. It's because you defer well that you do well. When you attempt to judge for yourself, you actually engage in social cognition; and that's why you tend to get things right. You can reliably adjudicate between David Irving and his many critics, between climate scientists and denialists, between anti-vaxxers and genuine experts. But while it may seem to you that you do so well (epistemic individualist that you are) through the power of your unaided reason, a very important part of the explanation for your success is that you defer so fluently and appropriately. You owe your success to the way in which you are embedded in epistemic networks.

Even so, I bet even you sometimes go wrong. Your capacities, and your disposition to defer, only get you so far. You live in an environment that is unreliable, in which frauds and fakes mimic the cues to reliability you rely on. Sometimes—I bet—you fall for their tricks. I certainly have.

One of the examples featured in *Merchants of Doubt* is "acid rain." The phenomenon was first recognized in the mid-nineteenth century,

and regulations were introduced by the British parliament to address it. The problem returned to scientific and public consciousness in 1974, when Likens and Bormann (1974) published a paper in *Science*, showing that acid rain was a serious problem in large areas of the United States. The National Academy of Sciences and the EPA both launched investigations, and both concurred: acid rain was a "serious hazard to human health" (O'Connor and Weatherall 2019, 37). The Carter administration moved to regulate the power plant emissions that were largely responsible. But implementation was left to the incoming Reagan administration. And then the merchants of doubt moved in.

There's no need to rehearse, here, how some of the very same people who had been involved in defending tobacco and who would later obfuscate the science of climate change hijacked the process and ensured that it stalled. *Merchants of Doubt* tells the story much better than I can. Here I want to mention just one thread in the broader narrative. Edward Krug, a soil scientist at the Connecticut Agricultural Research Station, was promoted by the denialists and by institutions sympathetic to them as offering a view contrary to those who called for increased regulation. Krug argued that changes in soil acidity were largely the product of natural processes, not acid rain (Krug & Frink 1983). His claims were assessed and quickly refuted: acid rain was swamping such processes (Galloway et al. 1984). But despite the refutation, Krug's work was picked up beyond the science journals and weaponized in the fight against regulation. It was presented in *Policy Review*, *Reason Magazine*, and even on *60 Minutes*. The pollution seeped in everywhere: in 1990, NPR reported that the scientific consensus on acid rain was that the issue was complicated.

Surely, though, sophisticated people, people with a background in science and who are responsive to the right cues and read the right sources, surely *they* are able to see through the fog? In *Merchants of Doubt*, Naomi Oreskes has a confession to make: in the early 1990s she "used Krug's arguments in an introductory earth science class at Dartmouth College to teach 'both sides' of the acid rain 'debate'" (Oreskes and Conway 2011, 103). In an epistemically polluted environment, even the most sophisticated people risk being taken in.

Restoring Trust in Science

If we're to bring people to believe better, it won't be by asking them to behave more responsibly or by inculcating the epistemic virtues in them; not primarily and—I bet—not very importantly either. Epistemic humility, open-mindedness, care in evidence-gathering—these all good things (in their place). But they're no solution to the problem of believing better, largely because it's extremely difficult, and perhaps impossible, reliably to judge when they're called for and when they're not. They're dispositions that can as easily lead away from the truth as toward it (Levy & Alfano 2019).[9] More pointedly, it's simply false that the epistemic virtues and their responsible application enable the person reliably to track truths. To the extent she succeeds, it is her embedding in appropriate epistemic and social networks that enables her success.[10]

[9] Those people who generate conspiracy theories—as opposed to those who consume them—exhibit a great many of the epistemic virtues (K. Harris 2018). Conspiracy theories often begin from the identification of an anomaly: a piece of data that appears to conflict with the official explanation. The theorist exhibits open-mindedness in looking for an alternative explanation. She looks to alternative epistemic communities, thereby displaying epistemic humility. She certainly can't be faulted for a failure to look for evidence: conspiracy theorists may be voracious in their consumption of reports. In many ways, conspiratorial ideation looks like science: it is science gone wild, science no longer constrained by the epistemic networks within which mainstream scientists work.

[10] Eric Winsberg (2018), who accepts (and strengthens) the case for the claim that nonexperts can't hope responsibly to assess climate science for themselves, nevertheless argues that individuals can and should engage in careful assessment prior to accepting consensual scientific claims. We should assess not the basis of the claims the science makes, but the structures that underlie the generation of the consensus. Indeed, Winsberg thinks it's uncommonly *easy* to engage in this kind of assessment in the case of climate science. This is for several reasons. First, climate science is the product of a multiplicity of different disciplines and this fact ensures robustness against corruption in one area of the science and cross-checking of findings. Were one discipline's contribution suspect, the others would detect the problem when its results impinged on theirs. Second, climate science has an institution—the IPCC—that summarizes and assesses the science. Third, climate science is subject to well-funded hostile scrutiny. These facts ensure that the consensus is robust, and these facts are easily discerned by laypeople.

While Winsberg is surely right that these facts about the structure and the institutional setting of climate science entail a high degree of credibility, I am both much less skeptical that consensus (however generated) is good evidence and much more skeptical of ordinary people's ability to discern and understand the facts about structure that help to ensure reliability. Climate science *denial* is also the product of multiple disciplines. It, too, has institutions (including a shadow IPCC) that claim to summarize its results and identify the work that is reliable. It, too, is subject to well-funded hostile scrutiny. Just as markers of expertise can be mimicked, so can institutions and structures. Of course, the merchants of doubt haven't been able to mimic the extent and depth of the structural and institutional network underlying climate science. But I only know *that* by testimony! I don't know how to verify these claims in any

Philosophers often advocate the teaching of critical thinking skills, such as the capacity to identify argumentative fallacies, as a partial solution to our epistemic crisis. But such teaching has small and short-lived benefits (Mercier et al. 2017). A broader general education, including scientific education, also doesn't seem to reap any benefits. In fact, it may hurt: as we've already seen, better educated Republicans are *less* likely to accept the consensus view on climate change than less well educated Republicans (Kahan 2015). Better educated Republicans are also more likely to think that Obama is a secret Muslim (Lewandowsky et al. 2012). Better education and more tools for argumentation may enable those who distrust the institutions of science and the universities to counter their claims more effectively. This may arise from what Taber and Lodge (2006) call the sophistication effect, whereby being more knowledgeable provides more ammunition (and more skills) with which to counter unpalatable claims.

The restoration of trust in science and scientific institutions is likely to make a bigger and longer lasting difference to the goal of better belief formation. On the right, trust in these institutions has ebbed significantly in recent decades (Gauchat 2012). This distrust has generalized to the universities as a whole: a majority (58 per cent) of Republicans now say that colleges and universities have an overall negative effect on the United States (compared to 19 per cent of Democrats). The same survey shows that 85 per cent of Republicans have a negative view of the news media (Doherty et al. 2017). If we are to promote better belief, we need to promote better deference, and that requires the restoration of trust in these institutions. Central to doing so is reduction of epistemic pollution.

Epistemic pollution *rationally* reduces trust in institutions. If you know that the same institutions that credential science—universities, for

other way and I'm skeptical there *is* any other way. If testimony is good enough for knowledge about the structure and institutional setting of climate science, then it's also good enough for acquiring knowledge about the claims of climate science.

I'm also skeptical that knowing the facts about the institutions of climate science positions me to assess the degree to which these institutions are knowledge-conducive in any case. What counts as a knowledge-conducive structure and what counts as group think or corruption is a difficult issue, and not one that most people can assess for themselves. I am, in fact, skeptical that *anyone* can accomplish this on their own.

example—are also involved in gaming the credentialing system, your trust in them is rationally lower than otherwise. If you know that the funders of science often have conflicts of interest that might (and do) lead them to suppress unfavorable data, then you should reduce your trust in them. If you know that scientists themselves sometimes engage in questionable research practices, you should be somewhat slower to endorse their findings. While we can't realistically hope to eliminate these practices, we can and should take steps toward their reduction.

Pollution of the traditional kind is often hard to tackle due to a collective action problem: while everyone might be better off if no one pollutes, no individual can make a significant difference on their own, and any individual who pays the cost of clean-up locally is worse off than others who don't cooperate. Collective action problems are solved by mechanisms that ensure that (almost) everyone contributes to the goal. There are multiple ways this can be done, but often (and especially in cases when some of the actors don't share the goal), some degree of coercion is required. Epistemic pollution is also a collective action problem: while most of us would be better off if it were significantly reduced, individuals can't make a significant difference to it by themselves, and anyone who acts alone is worse off than others who don't cooperate.[11] It is, moreover, a collective action problem made worse by the fact that some actors don't share the goal most of us would like to achieve: merchants of doubt, purveyors of predatory journals and peddlers of expensive and ineffective drugs may prefer to go on polluting to having a clean epistemic environment. Reducing epistemic pollution will almost certainly require some degree of coercion, from government or other institutions with the clout to impose costs on those who don't cooperate.

While I'm not the right person to develop policy proposals, some preliminary steps toward restoring trust are obvious. We need to vastly reduce the number of predatory or fake journals, or ensure that such journals are effectively confined so that they aren't (and aren't seen to be) contaminants in the scientific ecosphere. Doing this requires that legitimate open access journals are clearly distinguishable from

[11] Everett and Earp (2015) suggest that the replication crisis is a tragedy of the commons; I think it's plausible to generalize this point across a range of epistemic pollutants.

illegitimate. This is a task for the scientific community as a whole. Universities should refuse to pay publication fees for journals identified as illegitimate and researchers who publish in them should not receive credit (in the form of citations, promotions or grant funds) for such publications. Such a move would starve the illegitimate journals of funds and should lead to the closure of most. Beall's list was a great start, but it would be better done collectively.[12] As we saw above, the list was controversial, in part because it erred on the side of considering journals illegitimate. It would be better to use less-demanding criteria. A consensus of the scientific community on the vast majority of such journals could easily be reached and a very significant epistemic contaminant drastically reduced.

Problems in the conduct of legitimate research must also be addressed. Incentives should be put in place for the replication of research; such incentives should be combined with a greater willingness to publish failed replications (alternatively, institutions can mandate such replications as part of the training of graduate students; see Everett and Earp (2015) for a proposal along these lines). Hypotheses and methods should be preregistered to ensure that researchers don't engage in questionable practices *post hoc* to ensure significance. Preregistration also eliminates the temptation for selective reporting of results: if only and all preregistered studies are published, we can be confident that we have the full array of data. Statistical techniques can be utilized to compensate for the file drawer effect and thereby generate more realistic effect sizes. Such techniques can also identify evidence of data manipulation, such as p-hacking. These proposals are by no means novel: in fact, many are already being implemented. Prestigious journals in psychology, for example, have implemented changes to their practices, requiring bigger sample sizes (lowering the risks of chance findings) and

[12] Beall's list of Predatory Journals and Publishers was maintained and updated by Jeffrey Beall, an academic librarian, from 2008 to 2016. Though controversial, it was widely respected and consulted. While it remains unclear why Beall chose to shutter the service, there is evidence that pressure from predatory publishers played a part in his decision (Straumsheim 2017). Individuals are less able to resist such pressures than collectives, in which pressures can be shared and mutual support offered. Moreover, collective decisions may be less controversial, especially if the decision-making body includes individuals with different perspectives and interests.

encouraging preregistration of hypotheses and methods (see Lindsay 2015). If the more prestigious journals all follow suit, ambitious scientists will be forced to adopt these standards and the degree to which science is unreliable should fall.

We should also reduce the incentives for science by press release and the extent to which new research is presented in the mass media (and to a lesser extent in the journals themselves) as revolutionary and earthshattering. This, too, is a collective action problem. Researchers likely prefer a world in which everyone refrains from hyping their research to the current situation.[13] Most might also prefer that media attention was not a significant determinant of prestige, promotions, and grant success. But given that media attention is valued by institutions and granting agencies, and no individual researcher can change the culture by themselves, each feels that they have to play the media game, which means representing their research as more important and revolutionary than it really is. The result is that consumers of the media are left with the impression that yesterday's research findings have been overturned by today's, and that today's will be overturned by tomorrow's—all of which has the result of reducing trust in any particular finding or claim.

Of course, actually implementing the agreements needed to solve collective action problems is difficult, especially given that science is an international enterprise. There are at least two possible routes to effective regulation. One is through governmental action: if the United States and the EU ensured that funding of science was tied to responsible media engagement, norms might change across science (given the proportion of science funded by them). Bringing China on board would be even better, and might be possible. Admittedly, in many domains the record of government is not encouraging: when questions come to be politicized, policy often ignores expert opinion. We can reasonably be more optimistic about self-regulation from within epistemic

[13] What if the research genuinely is earthshattering? I strongly suspect that rules regulating science and its reporting should not be written in ways that make explicit allowance for such eventualities. The genuinely earthshattering is sufficiently rare that we do better to design regulations that assume that the research governed by them is normal, not revolutionary, science. If research is sufficiently significant, this can be expected to be evident without the need to hype it.

communities, insofar as policy makers within these bodies must remain responsive to the expert opinion of their members. National peak organizations, for instance, could regulate the conduct of science in ways that could produce the same effects (self-regulation by epistemic communities may also reduce worries about overreach, given that such communities have a much smaller domain over which they exercise power than do governments). Perhaps no official mandates are necessary: norms within science are changing rapidly, and many researchers now look askance on unregistered hypotheses and small numbers of participants. Social pressures may go a long way toward fixing many problems.

Changing media norms is a tougher nut to crack. In the contemporary environment, the media is fractured and cooperation unlikely. The collapse of traditional funding models has left media organizations—where they survive at all—chasing clicks, which encourages sensationalism. It might be possible to produce many of the desired effects without the media on board: if reputable scientists withdraw cooperation with sensationalistic media, they may come to be known to feature only charlatans and the likelihood that the public will ignore them will rise.

None of the measures mentioned above, some of which have already been introduced (albeit patchily), would solve the problem of distrust in science in the short term. When trust is lost, it's difficult to restore, and measures taken by the very institutions distrusted are likely to be regarded with a skeptical eye. Over the longer term, however, removing epistemic pollutants from the environment should increase trust in reliable sources of information, and thereby improve belief formation.

In this and the previous chapter, I've argued that individual cognition—unaided—is much less powerful than we tend to think. Without heavy duty scaffolding and heavy reliance on others, we're very much less reliable than we hoped, even in the best cases. Worse, we're rarely in the best cases: we live in an epistemically polluted environment, in which others seek to misdirect us. When we do well (as we frequently do), we tend to attribute our success to our own individual cognition, but that's only because our deference is so smooth and automatic, we fail to notice it.

Showing that thinking for ourselves is less powerful than we thought or hoped is one thing, however; it's quite another to show there's an

alternative that is both more successful and also ethically permissible. It's to that task that I turn next. The epistemic engineering I advocate is, or is closely akin to, nudging, and nudging is hugely controversial. It's widely seen as impermissible, or at least undesirable, on the grounds that it subverts individual autonomy. I'll argue that while nudging *can* be used to subvert autonomy, that's not because there's any distinctive problem with influencing people's behavior in this kind of way. Nudging is the presentation of information, and done appropriately it's no more subversive of autonomy than is giving (explicit) arguments for a conclusion or a course of action. In making the case for the permissibility of nudging, I'll also complete my case for seeing ourselves as rational animals. I'll provide a fuller defense of a claim I've made several times already: that deference is fully rational: not merely *ecologically* but also *directly* rational.