Concluding Thoughts

Rational Animals After All

The Enlightenment celebrated the power of human reason. In more recent times, a number of psychologists and naturalistic philosophers have attempted to replace the Enlightenment conception of our cognitive powers with a deflationary and allegedly more realistic picture. On this picture, often presented within the framework of dual process theory, the kind of rationality the Enlightenment highlighted, impressive as it genuinely is, is only a small part of the fuller picture. Type 2 cognition is a scarce resource, and we must be selective in deploying it. The rest of the time—most of the time—we must rely on fast and frugal Type 1 cognition (Bargh & Chartrand 1999). Even when we do deploy Type 2 cognition, we remain reliant on Type 1 cognition for inputs, which limits the power of Type 2 cognition to correct errors. This is bad news for our rationality, because Type 1 cognition is inflexible and sometimes atavistic.

Pushback against the claim that our cognition is lazy, inflexible and pervasively irrational has come mainly from advocates of ecological conceptions of rationality. They emphasize how well adapted such cognition is to the threats we actually faced in the environment of evolutionary adaptiveness. They concede, of course, that we pay costs for relying on it: for example, we respond to stimuli as if they were threats while knowing that they're harmless. Even when we deliberate carefully, we may be subject to prejudices we disavow, because the reasons we consciously process have already been assigned weights prior to deliberation (Uhlmann & Cohen 2005). But these are costs a well-designed thinker would pay, proponents of ecological rationality maintain. It's better to respond quickly to a possible threat then to deliberate longer:

Bad Beliefs: Why They Happen to Good People. Neil Levy, Oxford University Press. © Neil Levy 2022.
DOI: 10.1093/oso/9780192895325.003.0007

we'd do better to pay the cost of a false positive then to run the risk of being bitten or mauled. If we're irrational due to our reliance on fast and frugal cognition, we're rational to be irrational in this way.

In this book, I've presented a very different picture. I've suggested that we're more rational than naturalistic philosophers have tended to think. We've failed to see how rational we are because we've been looking for rationality in all the wrong places. We've been looking at individual cognition and at first-order evidence to vindicate our conception of ourselves as rational agents. Both of these things matter, of course: they matter a great deal. But apparent failures to rely on them often don't indicate departures from rationality. They indicate a rational outsourcing of our cognition, a reliance on the division of epistemic labor, and the appropriate use of higher-order evidence.

With this picture in place, we can now see that we're *individually* more rational than we sometimes seem, though our individual rationality doesn't take the form we expect. The behavior of other agents is higher-order evidence for us, and individually we respond to it appropriately. Deference to experts is an appropriate use of higher-order evidence; so is the use of the conformity bias and the prestige bias. The use of environmental cues is the use of higher-order evidence: it renders options salient to us. Most of the time, we respond appropriately to the communicative cues in our environments. Our individual rationality doesn't consist in our processing of first-order evidence alone. It consists, also and importantly, in the use of higher-order evidence, made available to us through distributed and outsourced cognition.

Our flexible and intelligent response to higher-order evidence is rational, whether it's the product of conscious deliberation or of automatic processing. Of course, deliberation (conscious or not) about first-order evidence is also essential to our epistemic success. My deference to experts is appropriate because *they* have produced a body of knowledge, and their knowledge production is heavily reliant on first-order evidence. But no significant epistemic achievement is the product of deliberation about first-order evidence alone. Even scientists must defer, even on their own terrain, because no individual is able to grasp all the methods and tools and data she must nevertheless rely on for her work. Within the narrow sphere of our expertise, our reliance on first-order

evidence is relatively heavy; elsewhere, higher-order evidence plays a much greater role.

At this point, a worry about self-defeat obviously arises. This book is itself the product of individual deliberation, and argues against an orthodoxy that emphasizes first-order evidence, intellectual autonomy and individual deliberation. If I'm confident that I'm able to see the problems in the orthodox view, then I seem committed to thinking that individual cognition can successfully strike out on its own, contrary to my own claims. Aren't I committed to making an exception of myself, in a way that is at best unprincipled? Deference for thou; not for me.

There is, I acknowledge, a tension between the message of this book and the very act of advocating that message. To some degree, this tension is endemic to philosophy itself, insofar as it tries to show "how things in the broadest possible sense of the term hang together in the broadest possible sense of the term" (Sellars 1962). This ambition requires stepping back and trying to find what is common to disparate areas of inquiry. In an age of hyperspecialists, philosophy still pretends to the role of the generalist. Inevitably, the generalist runs the risk of a failure of understanding when they take as their subject matter the findings or the nature of specialized disciplines. Many philosophers, including me, have responded to this kind of worry by backing away from generalist pretensions and instead limiting ourselves to more specific domains, where we can hope to make headway. But this book attempts much more general claims about human knowledge and its acquisition and thereby runs full tilt into these concerns, and in a particularly pointed way due to its advocacy of deference.

I take comfort from the fact that I'm not entirely out on my own here. In this book, I'm leveraging and building on the work of many others: work in cultural evolution, in psychology, in social epistemology, and in other fields. To a large degree, I *am* deferring to these thinkers (moreover one of my major aims, albeit one that's in the background, has been to make us more accepting of a science—climate change—that I lack any real capacity to understand). I do take my own advice in this regard, if not as extensively as perhaps I should. Nevertheless, to some degree I am caught in worries about self-contradiction. I present what I take to be first-order evidence in favor of relying on higher-order evidence,

arguments in favor of deference at odds with an orthodoxy that denigrates it.

Of course, I haven't advocated relying on higher-order evidence *rather than* first-order evidence. I've stressed that higher-order is reliable, in very important part, because it's generated by people who are grappling with the first-order evidence. All cognition is dependent on both first-order and higher-order evidence, to varying degrees in different contexts, depending on our capacity to assess each kind of evidence. To some extent, that fact mitigates the tension. Still, I can't complain to have the kind of expertise that would warrant me in setting aside the near-universal consensus in favor of individual reasoning and striking out on my own. The tension remains unresolved. Here I can only appeal to you for help. Come over to my side. Once there are enough of us, I can comfortably advocate deferring to the new orthodoxy. If you're on the fence about my arguments, I hope concern for my comfort will tip you over into accepting them.

Let me finish with a few words about the Enlightenment. Must we abandon its legacy, if we accept the picture of ourselves I've urged here? Perhaps not. *Sapere Aude!*, Kant's injunction, is most naturally interpreted individualistically. Kant calls on us to emerge from "immaturity," which he characterizes as "the inability to use one's own understanding without the guidance of others." He thus calls on us to use our "*own* understanding" (Kant, 1991: 54; emphasis in original). I, too, advocate we use our own understanding. On my picture, though, there is no conflict between such use and apt deference. We should *not* use our understanding without the guidance of others; instead, a primary function of our understanding is in orienting us well toward such guidance.

Kant called on us to change our epistemic strategies, to rely more on our individual judgment and less on the judgments of others. Insofar as I have advice for each of us, as individuals, it's to rely on others more and better (and of course, insofar as we're able, to engineer the epistemic environment to support such reliance). We err in overemphasizing individualism, not in deferring too much. Does that entail abandoning the legacy of the Enlightenment? Not necessarily.

First, perhaps Kant was right to call on his contemporaries to think more for themselves and to defer less. It is no part of the picture

I've presented here to claim that individual deliberation over first-order evidence isn't a central component of our cognitive success. I've argued we ought to defer to the scientists, and that scientists must defer to one another. But deference has its limits. We should defer to scientists in very important part because they've deployed their individual cognition in the domain of their expertise. Admittedly, they've deployed it in a way that is socially and institutionally supported, and in ways that are heavily imbricated with deference; they've employed it nevertheless, and their reliability is partly (but only partly) due to that fact. In Kant's historical context, perhaps people had too little opportunity, or too little motivation, to deploy their individual cognition in these ways (I leave that as a question for historians).

Second, we needn't see the legacy of the Enlightenment as exhausted by this heavy emphasis on *individual* rationality, in what Kant regards as its mature form ("without the guidance of others"). Rationality, in its fullest sense, is, roughly, the deployment of cognition in the effective service of truth by appropriate response to the evidential content of information. Many psychologists and philosophers see us as rationally irrational: we deploy our cognition in the effective service of truth but we do so through the use of heuristics and other fast and frugal processes that do *not* respond appropriately to the content of our information. We respond irrationally—in ways that are not warranted by our evidence—but we're rational to do so. On my account, we are rationally rational. We respond to the *higher-order* evidence encoded in our environment and in the assertions of others, by deferring to them or even self-attributing beliefs. We do so in the service of truth. We're rational animals after all, even if our rationality is somewhat different to how we imagined it. We need to have the courage to use one another's understanding as well as our own.