2
Culturing Belief

What kind of being are we? This of course is one of the oldest questions in philosophy. In earlier eras, answers were often non-naturalistic (we are animals with souls, for instance). Today, one of the oldest answers is also one of the most popular: with Aristotle, we often think we are distinguished from other animals by our rationality.

Looking for necessary and sufficient conditions is a fool's errand, and a search for essences even more so. There's no property that distinguishes all and only human beings from other animals, beyond facts about descent. But there's something right about the claim that we're rational animals. Our intellectual capacities and achievements are distinctive and impressive. At the same time, there's also something misleading about the common picture of ourselves as rational animals. We think of our rational capacities as realized by our big brains, and there are quite a few grains of truth to that thought. But our rationality also depends on our sociality, and thinking of ourselves as cultural animals is no less accurate than thinking of ourselves as rational animals.

Nineteenth-century history provides excellent, if sometimes grim, illustrations of how limited individual cognition is compared to cultural knowledge. Big brains, good education, and excellent preparation often weren't sufficient to allow nineteenth-century European explorers to survive in difficult environments. In 1846, two ships commanded by Sir John Franklin, on an expedition to chart the Northwest passage, became stuck in sea ice in the Canadian Arctic. The entire crew perished. But the area was regarded by the local Inuit people as rich in resources. Despite their training and experience (Franklin was on his fourth Arctic trip) and ample resources, they were unable to acquire the skills they needed to survive. A few decades later, Roald Amundsen spent two winters in the same region. He relied on the help of the Netsilik Inuit for his survival (Boyd et al. 2011).

Bad Beliefs: Why They Happen to Good People. Neil Levy, Oxford University Press. © Neil Levy 2022.
DOI: 10.1093/oso/9780192895325.003.0002

In a very different environment, the Burke and Wills expedition to cross the Australian continent also suffered catastrophe due to rejecting indigenous knowledge (Burcham 2008). Running low on food, members of the expedition accepted a gift of cakes made from the nardoo plant for sustenance. However, apparently as a consequence of unease with being reliant on people they saw as inferior, they spurned further assistance and attempted to make the cakes themselves. They ground the seeds into a powder, mixed it with water, and baked it. Unbeknownst to them, the local people roasted the seeds prior to grinding. This step is required to remove toxins from the plant. Because they missed this step, the explorers didn't get the nutrients they needed from the nardoo cakes. There was only one survivor: he accepted further aid from the Yandruwandha people.

Outback Australia and the Canadian Arctic are harsh environments. But indigenous people didn't just survive in these environments; they flourished. They flourished not only due to their big brains—which they shared with the explorers—but because of their cultural knowledge. Life in the Arctic depends on a rich range of accumulated innovations (Boyd et al. 2011; Richerson & Boyd 2008). Traditional Arctic life requires knowledge how to make special clothing, to manufacture and use special tools for hunting, to construct snow houses, to build fires (without access to wood!) for cooking and melting water, and much more. Each of these skills is complex and hard to learn. Take clothing. The Inuit stayed warm and comfortable by making clothing from caribou skin, which has better insulation properties than seal or polar bear fur. But not just any caribou skin will do: it has to be harvested at the right time of year, and then prepared by repeated stretching, scraping, and moistening. The hides then have to be shaped in ways that maximize heat retention while allowing moisture to escape. A ruff of wolverine fur, especially selected for length, is then added. Footwear is equally specialized, consisting of *five* separate layers: three different layers of stockings, each with a different design, then two different kinds of boots. The know-how needed to make just one of these items of clothing is difficult to acquire, let alone the whole package. Moreover, individual elements of the package (for example, knowing how to prepare one of the layers of stockings) are often little use by themselves: it is only when they play a role in the entire package that they make a significant difference to survival.

If each element of the package is hard to acquire, and many elements have little value on their own, how did the Inuit succeed in acquiring the whole package? Almost certainly, the answer involves multigenerational accumulation of knowledge, with many different group members each playing their small part in the acquisition of the propositional and procedural knowledge of the group as a whole. Cumulative culture, an irreducibly and deeply collective enterprise, is essential to the acquisition of the kinds of knowledge that allows human groups to flourish in highly diverse environments, from the deserts to the tundra and the tropics.

The fate of the Franklin expedition stands testament to how far this kind of knowledge exceeds the ability of a group of individuals to reconstruct from scratch. The Inuit *themselves* found it impossible to reconstruct their own knowledge when it was lost. An epidemic seems to have struck the Polar Intuit sometime in the 1820s, resulting in the death of many older members of the community; that is, they lost many repositories of cultural knowledge (Boyd et al. 2011). As a result, they lost important skills, and the group entered a decades-long decline in numbers. This decline was halted only when they reacquired the skills from another group around 1862. They were not able to reinvent the lost skills of kayak and snow house design in the intervening decades, despite strong motivation and possession of a suite of related skills.

Cultural knowledge solves problems that are intrinsically difficult. When feedback is quick, individual cognition is often up to the task of solving problems. We rapidly learn to avoid suspension-destroying potholes in the road or nettles in the bushes. But when the relationship between an action and its effects is slow to manifest and probabilistic, individuals do very badly on their own. Think of how long it took to demonstrate the effects of tobacco on health; for decades, people denied the link between smoking and cancer, because they were more impressed by salient cases of individuals who had lived to ripe old ages despite smoking heavily (and, of course, because merchants of doubt deliberately muddied the waters; Oreskes & Conway 2011). Science has developed mathematical tools for detecting signal in the noisy relationship between variables, like the relationship between smoking and cancer. Without such tools, individual cognition is highly unreliable. But cultural cognition often succeeds in identifying the signal amid the noise without the need for statistical tools.

Detoxification methods like those developed by the Yandruwandha people (and which the Burke and Wills expedition failed to imitate) are good illustrations both of how difficult these problems can be, and of the spectacular success of cultural cognition at solving them. Because edible plants often evolve toxins as protection against herbivores, many staple foodstuffs are (or were, before very intensive selective breeding) toxic. Take corn. Corn is cheap to produce and high in energy. For this reason, it was exported from the new world, to which it is native, to the old quite rapidly after the arrival of the Spanish. It came to be an important food crop in Italy, Spain, and later the southern United States. But with corn consumption came pellagra. Pellagra manifests first as a skin disease, but untreated it can lead to dementia and even death.

Medical professionals quickly realized that there was a link between pellagra and corn consumption, but suspected some kind of contamination as the cause. It wasn't until the second decade of the twentieth century that the real cause was identified: niacin deficiency. Yet pellagra was very rare among the indigenous people who had relied upon corn for centuries. They avoided pellagra by cooking corn together with an alkali, which releases otherwise chemically-bound niacin. They used wood ash, or ground sea-shells, or lime, depending on what was available locally. But of course they had no concept of niacin or alkali. Asked *why* it is necessary to mix wood ash with corn meal, indigenous people often had no more to say than "it is our custom" (Henrich 2015). They may not have known *that* it is adaptive, let alone *why*. They were and are smart, big-brained primate like the members of the Burke and Will expedition, but it wasn't their individual cognition (alone) that had allowed them to develop the detoxification processes that helped them avoid pellagra. It was cultural evolution, which enables the detection of a signal in a very noisy background.

Cultural Evolution

We started this chapter by asking what kind of being we are. I suggested that an answer as good as any would be that that we're cultural animals. "Culture," as I use the term here, refers to information that is acquired from others, by vertical or horizontal transmission (i.e., from elders or peers) and which affects behavior. In this sense, culture is not unique to human beings. Famously, when a young Japanese macaque dubbed Imo learned to wash sweet potatoes left for her troop on the beach by dunking the potatoes in the sea, the behavior spread throughout the troop. Since Imo's innovation was first documented, a variety of other behavioral traditions have been observed among macaques (Laland 2017). Chimps have traditions that differ from troop to troop, ranging from different techniques for termite fishing to the use of stone tools to break open nuts; orangutans also use tools, with the tools and the method of exploitation differing from area to area (Schaik et al. 2003).

Human cultures differ from the cultures of other primates in complexity, of course. But more importantly, only human culture appears to be *cumulative* culture. In our species (perhaps alone; certainly to an extent that is dramatically greater than in any other), cultural innovations are not merely transmitted: they become a platform on which others can build. Human culture is subject to the "ratchet effect" (Tennie et al. 2009). The behavioral traditions seen in other primates and other animals can preserve individual innovations and transmit them to future generations, but only cumulative culture builds on these innovations, enabling cognitive achievements that go beyond what any individual or any generation can achieve. It is this development over time that allows cultural evolution to detect signal in noise when the noise exceeds the capacity of unaided individual cognition to parse. It also enables the detection of temporal fluctuations that exceed living memory (Shea 2009). Cumulative culture opens up horizons for knowledge that are closed to individuals, no matter how individually gifted they are.

There's good reason to believe that the mechanisms underlying cumulative culture are evolutionary. Evolution is substrate neutral: it needn't be limited to biological reproduction. Evolution occurs whenever (roughly) there is selection between individuals which vary in their characteristics and this variation is heritable. So long as traits are differentially rewarded, these traits are heritable, and the environment is sufficiently stable over time, we should expect evolution. For instance, prior to (and as a condition of) the emergence of life, evolutionary processes account for the emergence of organic compounds (Bada & Lazcano

2009). Evolutionary mechanisms also account for changes in human behavior across time, independently of changes in gene frequencies.

Cultural evolutionary theory is sometimes mistakenly identified with *memetics*. Memes, first proposed by Richard Dawkins (1989), are units of culture analogous to genes. Memes are subject to selection pressures and they are heritable. Whereas genes (almost always) get selected *en masse*, however (when organisms are selected), memes are selected one by one. They can be fit even when we, their hosts, are not—even when they *lower* our fitness. A classic example of a meme that replicates independently of its fitness effects on its host is an earworm: a catchy tune that gets lodged in our heads long after we hear it. Notoriously, earworms may be disliked by those who experience them, but they are good at replicating themselves. A person who dislikes an earworm can nevertheless be a vector for its replication: she may find herself humming it, for example, thereby contributing to its spread. If she loses friends as a result, we may see a dramatic dissociation between the host's interests and the interests of the meme.

Perhaps memetics explains some features of culture. However, it can't explain the emergence and transmission of cultural practices of any great complexity. The kind of cultural evolution I'm invoking is of a different sort. While it's non-genetic, its effects are primarily on the fitness of the organism (and perhaps the group), not on the fitness of the units of culture—if there are any, in any meaningful sense—themselves. Beliefs, technologies and practices make an obvious difference to our fitness (given that they make a difference to how we behave), and therefore affect our biological fitness. Believing that *that* plant is edible and *this* is not may be the difference between life and death; inheriting the capacity to make a boat of a particular kind may enable a better catch, which in turn enables the person to support more children. Even songs or rituals may make a difference to biological fitness, say by increasing bonds of solidarity that enable a group to avoid open conflict.

Cultural evolution produces adaptive changes in practices or beliefs without (or independent of) changes in gene frequencies.[1] One group

[1] Note that gene-culture coevolution may occur. The classic example is the development of lactose tolerance. A cultural innovation—dairy farming—brought about selective pressure for genetic evolution in those groups to which the practice spread (Cavalli-Sforza et al. 1994).

may be fitter than another without the groups differing genetically; the fitness may instead be due to cultural practices. Cultural evolution has likely been more significant than genetic evolution as a factor in human evolution over the past 50 millennia or so, and perhaps much longer, if only because it can occur very much faster than biological evolution (Perreault 2012). It usually takes many generations for biological evolution to occur, but cultural evolution can occur within a generation.

Some otherwise puzzling facts about human beings might be explained by the centrality of cultural evolution to our flourishing. Human beings have unusually long periods of dependency on caregivers. This long period of dependency is also an apprenticeship in the local culture. We live in a dizzying diversity of environments, and what is adaptive in one may be highly maladaptive in another. The beliefs and behaviors that are adaptive in the Amazonian rainforest are very different from those that are adaptive in the Arctic, which are different again from those required in Karachi or in Copenhagen, or in sub-Saharan Africa, or the Australian outback. Genetic evolution can transmit behaviors in animals (nest building, hunting techniques, song patterns—though even in these cases some kind of learning usually plays a role too). It can even transmit conditional behaviors: that is, it can encode instructions for one set of behaviors in one kind of environment and a quite different set in another (so-called facultative adaptations). But it can't encode for the enormous diversity of complex behaviors needed for flourishing across the range of (ever-shifting) environments in which our species lives. Hence the long period of apprenticeship: we need the time to acquire the set of behaviors we'll need in the specific environment we're born into.

Our early malleability and long dependence isn't the only adaptation we have for the acquisition of culture. We have a whole suite, whether as a product of genetic evolution or (as Cecilia Heyes (2018) has argued) due to cultural evolution itself.[2] A well-known (albeit somewhat controversial) example is our disposition to imitate. We use the verb "to ape" to

[2] Even on Heyes' account, we have some dispositions that are not themselves explained by cultural evolution but which facilitate it. For instance, we are peculiarly tolerant of others and especially the young, allowing them to observe our behaviors.

describe slavish imitation. In fact, no other animal—not even the other apes—apes as much as we do. There's experimental evidence that we're *overimitators*: whereas other animals copy behaviors when they recognize they're instrumentally rational, human beings are disposed to copy even those components of behavior that don't appear to be required for goal pursuit. Nagell, Olguin, & Tomasello (1993) demonstrated a novel technique to human children and chimps. They used a rake, *tine side down*, to draw sweets that were otherwise out of reach toward themselves. Using a rake that way is very inefficient: many sweets slip through the gaps in the tines. Given the opportunity to perform the task themselves, chimps flipped the rake so that the flat side acted as a more efficient tool, with fewer sweets escaping. But human children tended to imitate the action just as demonstrated.

Later experimental work demonstrated that the disposition to overimitate is selective: children overimitate when the behavior appear to be intended by the model, regardless of whether they see the point of the actions. For instance, while children who observed a model turning on a light by butting it with her head tended to do the same, rather than use their hands (Meltzoff 1988), imitation dropped significantly if the model's hands were occupied, suggesting that the decision to use her head was not a component of how things are supposed to be done (Gergely et al. 2002). This sensitivity to whether the behavior is intentional allows children to distinguish between those behaviors that are constitutive of the culture they are acquiring and those that are incidental.

It's worth pausing to appreciate the (prima facie) oddity of this contrast between chimp and human behavior. We humans pride ourselves on our intelligence; we're supposed to be rational animals. Yet when humans and chimps are confronted with the task of procuring a valuable good using a tool, it's chimps and not us who analyze the causal structure of the task to accomplish the goal more efficiently. Children copy the behavior demonstrated if it appears intentional; chimps drop irrelevant features or inefficient techniques. We are *over*imitators: we imitate more than seems justifiable given our goals and the nature of the task. Even when some steps are very obviously irrelevant, and even in the face of prompting to drop superfluous elements, children are reluctant to deviate from the demonstrated sequence of actions (Lyons et al. 2007).

Nor is overimitation confined to children: Flynn & Smith (2012) found that only when adult participants were told that a demonstrator was themselves a novice did they not overimitate.

Why are we prone to ape more than apes? If we're so smart, why do they seem to outperform us in identifying efficiencies and more successful routes to a goal? It's because *we* are cultural animals, and *they* are not (not at least, to anything like the extent we are). Imitation is an adaptation for culture. It allows us to acquire knowledge and practices developed by multiple individuals, individuals dispersed across space and time. It allows us to acquire, and then to build on, deeply social knowledge: adaptive behavior that could not have been developed by any individual de novo, no matter how gifted and insightful that person might be. Sir John Franklin and his party could not develop the techniques that would have allowed them to survive in the Arctic: this suite of techniques and this knowledge must be developed and refined across multiple generations and multiple individuals. What they lacked wasn't intelligence or physical capacity. They lacked the requisite culture: it is only by appropriate imitation that we can acquire the requisite knowledge, and that usually takes long enculturation.

Deeply social knowledge and practices may be partially opaque to those who inherit them. They may be deployed by people who know *that* and *how* they are to be used, but who have mistaken ideas about how they work, or no ideas at all. As we saw, indigenous Americans often gave ethnographers no further justification of the practice of preparing corn with an alkali beyond "it is our custom" (Henrich 2015). Sometimes practices are justified in supernatural terms. For example, Naskapi hunters decide where to hunt by using the shoulder blade of a caribou, heated over hot coals so that it cracks and burns, as a kind of map. There's evidence that this kind of divination is adaptive because it effectively randomizes behavior, overcoming our disposition to detect illusory patterns (Henrich 2015). This kind of adaptive use of divination is widespread across multiple cultures. In yet other cases, behavior is given a naturalistic but false rationale. For example, the more toxic marine foods are tabooed for pregnant women in Fiji, significantly lowering their risk of miscarriage (Henrich & Henrich 2010). But justifications for the practice offered range from custom to the idea that the child would take on the properties of the animal eaten (e.g., that it might have rough skin if its mother ate shark).

Causal analysis, chimpanzee-style, is not required for the acquisition and deployment of cultural knowledge. In fact, since the products of cumulative culture are complex and the contribution of some steps obscure (why roast the nardoo seeds? why add wood ash to the corn meal?), such analysis risks degrading the value of cultural knowledge. The cultural apprentice does better taking the technique on trust. *That's* how we do things, so that's how I'll do them. Second guessing the technique is appropriate when it's a component of shallow cultural knowledge, like the behavioral traditions seen in other primates: when what is transmitted is an innovation that is within the cognitive grasp of a single individual, tinkering with it may easily reap rewards. When what is transmitted is deep cultural knowledge—when the culture is cumulative culture, and successive innovations have become the platform for further elaboration—then it is maladaptive to attempt to innovate by analysis. *Deference* to custom is the appropriate attitude.[3] We might say that we owe our success to the fact that we are in some ways less—or at any rate less directly—rational animals than chimps. We defer to tradition (relatively) unthinkingly, in conditions in which they would analyze causal structure and innovate (in later chapters, I'll suggest that this deference is more rational than it might seem).

Our disposition to overimitate is just one of our adaptations for the acquisition of local norms, practices, and conventions. Conventional ways of behaving are essential to the coordination of action in complex

[3] Of course, too much and too effective deference would put an end to cumulative culture by preventing further innovations. To be cumulative, innovation must occur. It remains somewhat mysterious how we pull off the difficult balancing act required—faithful transmission combined with some degree of innovation—insofar as the two elements are in conflict with one another (Sterelny 2012). Fridland (2018) argues that imitation supports faithful transmission and explicit pedagogy supports innovation. As she notes, when skills are causally opaque, imitation is required for transmission while innovation threatens to be fatal. But when skills have a causal structure that can be discerned, we may do better to transmit them by breaking them down into steps. This proposal leaves the central problem unresolved, however: how do causally opaque complex practices develop in the first place? Fridland's proposal seems to require that such practices must first develop through explicit pedagogy, and subsequently become the target of imitation. But that proposal cannot account for deeply social practices, which have components that those who transmit them may not be able to grasp. *Pace* Fridland, some kind of more genuinely evolutionary process, involving mutation and selection, must play a significant role in the production of deeply social knowledge.

societies, but have an arbitrary content (D. Lewis 1969). Think of driving conventions: it doesn't matter whether we drive on the left or the right, but it matters very much that we all agree on which side we drive on. Because conventions are arbitrary, we can't guess or infer what the local convention may be. It's not only conventions that differ from place to place; so do norms, some with very complex contents (think of how whole books get written about norms of etiquette, not to mention moral norms). The *conformist bias* is a disposition to acquire the local ways of behaving, enabling us to acquire the right set of conventions and norms (Henrich & Boyd 1998). The conformist bias enables us to pick up adaptive ways of behaving without paying the costs of exploration.

While the conformist bias helps us to acquire the local conventions and norms, the *prestige bias* leads us to imitate particularly successful individuals (Chudek et al. 2012; Henrich & Gil-White 2001). The link between their success and their behaviors may be opaque to us. Greater success could be due to their hunting techniques, their social network, their religion, their diet, their dress, and so on. Since the relationship between success and behavior is often causally opaque, we do better to copy successful individuals relatively undiscriminatingly. The prestige bias (like all heuristics and biases) leaves us vulnerable to certain sorts of exploitation: advertisements that link a celebrity or a well-known athlete to a watch or a fragrance take advantage of our disposition to copy successful individuals, even when there's no apparent link between their success and the behavior imitated. But it also helps to drive the evolution of behavior by allowing more successful behaviors to spread within and across groups. I'll call the prestige bias and the conformist bias instances of *social referencing*—looking to others within our social group for cues for what to believe and how to behave. As we'll see in later chapters, social referencing is a very widespread means for people to acquire and update their beliefs.

Imitation, the prestige bias and the conformist bias are adaptations for culture that help to explain why we're less discriminating in what we copy and less disposed to analyze imitated behaviors than chimps. Just how discriminating the mechanisms underlying cultural evolution are is controversial. I've presented a picture that leans very heavily on one side of these debates. The scientists I've cited (Peter Richerson, Joe Henrich, Robert Boyd) are sometimes said to constitute the Californian school (Sterelny 2017). The so-called Paris school (led by Dan Sperber) takes a different view on many issues, and in particular on how discriminating the mechanisms of transmission are. This isn't the place to address this dispute in any detail, but I will say a few words in justification of my reliance on the Californians (but—importantly—not in opposition to the Parisians; the two are far more compatible than the polemics suggest).

Sperber and his colleagues argue that relatively undiscriminating imitative mechanisms play a smaller role in the transmission of culture than the Californians suggest. Instead, the Parisians argue that culture is reproduced largely through mechanisms that are *reconstructive*, with the learner contributing a great deal to the final product. The Parisians maintain that the reconstructive mechanisms are intelligent and strategic, rather than unthinking and deferential. For example, they cite evidence that children and adults are far from undiscriminating in whom they imitate (Mercier 2017). We are also highly selective in whose testimony we accept. We filter out testimony from out-group members, when there is evidence that it conflicts with the consensus, and when it comes from those that have track records of unreliability or who have shown lack of benevolence to us in the past (P. Harris 2012; Mascaro & Sperber 2009; Sperber et al. 2010). In doing so, we filter testimony by reference to cues that correlate with reliability: consensual testimony is (other things being equal) more likely to be accurate than dissent; malevolent testifiers are more likely to mislead us than benevolent, and so on. Filtering testimony in these ways is the behavior of a rational animal, not an unthinking conformist.

In emphasizing the intelligence of the mechanisms underlying the transmission of culture, the Parisians offer an important corrective to the perception that it is undiscriminating and unthinking. But it's the *perception* they correct; intelligence doesn't enter the picture only when Paris-style mechanisms are in play. Imitation, California-style, is not reflexive and automatic. Instead, it manifests a great deal of intelligence (see Boudry 2018; Buskell 2016). In fact, even our apparently automatic imitation *itself* manifests intelligence (it's a major aim of this book to show that's true). We don't face a choice between Parisian intelligence and Californian automaticity: the mechanisms emphasized by both sides should be seen as intelligent.

A second reason to think that the conflict between Paris and California is smaller than is sometimes thought is that the two schools focus on the explanation of different aspects of culture. The Parisians' focus is on cultural lineages (see, especially, Morin 2016). High-fidelity imitation is not necessary for the preservation of these lineages: folk tales, painting styles, and myths. In these domains, reconstruction (coupled with what the Parisians call "cultural attractors") may suffice for the transmission of culture.[4] But unobvious technological innovations and sophisticated environmental knowledge—the kind of knowledge that depends on the detection of a signal in a noisy environment—do depend importantly on such imitation.

Agents, adults as well as children, default to imitation when success depends (or seems to depend) on following a precise sequence of steps, when mechanisms are causally opaque, and when the demonstrator is presented as an expert (Acerbi & Mesoudi 2015). It is in these sorts of contexts that reconstruction risks the loss of cultural knowledge. If agents intelligently reconstruct the causal process, they may easily leave

[4] Cultural attractors are species-typical cognitive dispositions that play an important role in stabilizing cultural traditions. Because certain features of a narrative (for example) come naturally to creatures like us, transmission from teacher to learner can be noisy: the novice will reconstruct ambiguous or partial information in ways that match the transmitted narrative except in the (unlikely) event that the narrative is excessively unintuitive. But in the domain of deeply social knowledge, reliance on such attractors would be risky and therefore is likely to be minimal. Such attractors risk distorting the knowledge transmitted, because deeply social knowledge is often highly *un*intuitive. It is also worth noting that the cultural attractor story itself plausibly requires culturally specific, rather than species-typical, machinery to explain content. Take minimal counterintuitiveness theory (e.g., Atran 2002; Boyer 2000), often hailed as one of the great success stories of the approach. According to this account, minimally counterintuitive entities have an advantage when it comes to cultural epidemiology: they are more likely to be recalled and more likely to be transmitted. However, Sterelny (2018) has persuasively argued that the institutional and ritual scaffolding of religious belief is central to its survival. These factors are not species-typical, of course—not in their contents, at any rate—and they are neither invented by individuals nor understood by them as belief transmission mechanisms. If we can generalize the lesson, the attractors themselves may turn out to be far more deeply cultural than the Paris school would be comfortable with.

out crucial steps, since their contribution is often opaque to us (recall the fate of the Burke and Wills expedition). While the telling of folk tales may be adaptive, their precise content typically doesn't matter much, and reconstruction can be given free reign. But when the precise content is crucial, deference is required. We're therefore sensitive to cues for switching from reconstruction to imitation. These cues include the presence of experts, causal opaqueness, ostensive communication ("look at what I'm doing here") and other cues that indicate conventionality (Acerbi & Mesoudi 2015). In a slogan: *deeply* social knowledge depends on deference. In some domains, reconstruction—Paris-style—is probably the primary means of cultural transmission. But in the domains of causally opaque technology, such as food preparation techniques, and the detection of signals in noisy environments, imitation reigns supreme.

Up to this point, we've focused on cultural knowledge: knowledge of the behaviors that we need to function as members of a particular society and to flourish in sometimes harsh environments. The mechanisms we've examined have the function of enabling us to distinguish signal from noise in causally opaque systems, or to identify regularities at temporal and geographical scales that exceed the grasp of an individual. I noted that prior to the development of statistical tools, these mechanisms were the only means we possessed for the detection of such signals. Now, of course, we possess the tools of science, which allow us to achieve the same sorts of ends much more quickly, efficiently, and accurately. These tools can be deployed by individuals. Does that entail the end of the millennia-long age of deeply social knowledge?[5] The next section will assess the extent to which knowledge is social today, in environments far removed from the environment of evolutionary adaptiveness.

[5] Roughly speaking, the claim that knowledge of the causal regularities essential to our flourishing has transitioned from being deeply social to being individual inverts the major claim of Elijah Millgram's *The Great Endarkenment* (Millgram 2015). Millgram argues that as more and more domains of knowledge becomes hyperspecialized, our capacity to understand specialist knowledge shrinks and knowledge becomes fragmented. Millgram provides good reasons to reject the idea that knowledge has recently become less social. But he dramatically overestimates how individual it was prior to the age of hyperspecialization. Human knowledge has *always* been deeply social, and individuals have never had the capacity to genuinely grasp the kinds of knowledge that have always been distinctive of our species.

Science on Mars

Until the last decade of the twentieth-century, epistemology was a largely individualistic enterprise. It was primarily recognition of our pervasive reliance on testimony that altered the landscape (C. A. J. Coady 1992; Lackey & Sosa 2006). We are dependent on testimony, implicit and explicit, for our knowledge about the temporally and geographically distant, and for much of our knowledge about the unobservable posits of science. We learn the history of our country and about major events in the past few centuries via explicit testimony in schools and colleges, and about political events, current affairs, celebrity gossip, and much more through media (increasingly, social media). Sources of implicit testimony include fictional narratives (which encode a lot of information they don't explicitly assert: for instance, about social norms, or about kinds of things and states of affairs) as well as all the subtle cues which convey to us how we should speak, behave, and even think.[6] All this knowledge is social: it's conveyed to us by others and largely taken on trust. For much of what we know about the world, we are deeply dependent on others.

Science, however, may seem crucially different. Doesn't science demand we take nothing on trust? The motto of the Royal Society, the oldest national scientific society in the world, is often taken to express

[6] Epistemologists may quibble with my very expansive use of "testimony" to describe how information is conveyed from one or many agents to others. Certainly implicit testimony often fails to satisfy *any* of the criteria put forward by Coady (1992: 42): for instance, it may not be directed toward those who are in need of evidence. A case can be made that implicit testimony doesn't even satisfy the criteria set down by Lackey's (2008: 35–6) more permissive disjunctive account of testimony, though this is a harder issue to judge (much depends on how we interpret "is reasonably taken as conveying that p"). I don't intend to offer a definition of testimony. It's enough for me that the instances of what I'm calling implicit testimony here clearly have features that warrant treating them as or alongside testimony: they involve the transmission of information from one set of agents to another, whether the first intends to convey the information or not, and independent of whether the second set of agents is conscious that information has been conveyed to them. The information thus conveyed functions as evidence for the second set of agents and they respond to it accordingly. To me, these features seem sufficient to justify calling these instances of information transfer "implicit testimony." If the reader prefers some other term, so be it.

the scientific attitude: *Nullius in verba*—take no one's word for it. Science seems to require that every claim be regarded with a skeptical eye, and nothing be accepted unless and until it has been adequately tested. In science, our naked intelligence and our capacity to test the facts is all that really counts—right? Call this the *Martian model* of science, after the film (and novel) in which a lone individual's ability to "science the shit" out of things enables him to survive in an unforgiving environment, and call intuitions that accord with this model *Martian intuitions*.[7]

Martian intuitions are powerful, and not only as an explanation of the success of science. But they are misleading. Under a variety of (reasonably undemanding) conditions, group deliberation outperforms individual deliberation (Mercier & Sperber 2017). Consider our performance at reasoning tasks like the Wason selection task. In this task, participants are presented with four cards, and asked which cards must be turned over to test if some rule is being violated. The rule is a conditional: if *p* then *q*, and the cards represent *p* or not-*p* on one side and *q* and not-*q* on the obverse. For example, the rule might be "if there is a D on one side of a card, then there is a 5 on the other side," and the cards presented might be the following:

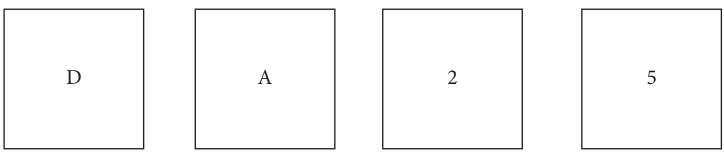

Figure 2.1 The Wason Selection Task.

[7] The line "I'm going to have to science the shit out of this" occurs in the film version of *The Martian*, but not the original book. It is, however, an accurate encapsulation of the flavor of the book as well as the film. Interestingly, the film and the line resonated strongly with many scientists (Neil deGrasse Tyson tweeted that that was his favorite line in the movie), despite the fact there's good reason to think that the success of science is explained by its structure and the division of cognitive labor it exhibits—by the way in which it instantiates features of the cultural model, the rival of the Martian model—rather than by the way it instantiates the Martian model.

Though the task is logically simple, most people do badly at it: only about 10 percent of people select the right cards on the task (in this instance, the first and third).[8] But *groups* of individuals do very much better. In fact, groups of deliberators may manifest the *assembly bonus effect*, where the group performs better than the best individual within it.

Despite the better performance of groups on the Wason selection task, and on many other reasoning problems (see Sunstein 2006; Surowiecki 2004 for accessible reviews), the Martian intuition remains powerful. Groups have a bad name (think of books with titles like *Extraordinary Popular Delusions and the Madness of Crowds* (Mackay 2011)). We badly underestimate group performance on reasoning tasks; we often think groups will do worse on such tasks, rather than better. Academic psychologists with a specialism in human reasoning, who are well aware of the dismal individual performance on the Wason selection task, underestimate the benefits of group deliberation on the task to the same extent as do laypeople. Managers of teams, individuals from East Asia and WEIRD people, all alike underestimate the benefits of group deliberation (Mercier, Trouche, Yama, Heintz, & Girotto 2015).

We are, it seems epistemic individualists who are reliant on social networks and culture for our epistemic success. What explains our individualism, in the face of the pervasiveness of epistemic dependence? Our obstinate individualism may be partially (and only partially) due to the fact that it is sometimes epistemically fruitful to insist on one's private evidence. In fact, some degree of epistemic individualism might be conductive to *group* deliberation. Group deliberation is subject to characteristic limitations and pathologies. Information cascades can mislead the group; powerful individuals can carry disproportionate weight and people may self-silence in the face of prejudice or anxiety. All of these

[8] One common mistake consists in picking cards representing the consequent (in this example, the 5). That's a mistake, because the rule is a conditional and not a biconditional. It might therefore be thought that people are guilty not of a reasoning error but of misunderstanding the nature of the task. One problem with this suggestion is that individual performance on the Wason selection task improves dramatically if the rule is a social contract rule (Cosmides & Tooby 1992). Since changing content while keeping the instructions fixed improves performance so significantly, it seems that the original results are not driven by a confusion brought about by the wording.

problems can be mitigated if people are disposed to give private information and their individual deliberation disproportionate weight.

Take information cascades. Such cascades may occur when agents are led to disregard private evidence because the public evidence outweighs their own. Consider a case like this (based on Anderson and Holt 1997). You're given the opportunity to draw a ball from one of two urns. One urn contains white balls in a ratio of 3:1 to red; the other the reverse. Your task is to identify which urn is which, by drawing a ball from one. Obviously, if you draw first you should bet that the urn you draw from contains predominantly balls of the color you've drawn. Suppose, however, you draw after other agents. You don't know what color ball they've drawn, but you do know how they've betted. With all agents acting rationally, unlucky initial draws can give rise to an information cascade and rising confidence in a falsehood. Suppose the first and the second agent both draw a red ball, and bet accordingly. Their behavior is rational, but their evidence may be misleading: by chance, they've drawn red balls from the predominantly white urn. The agent choosing third in sequence may now rationally bet that the urn is predominantly red, *no matter what color she draws*, because the evidence stemming from the betting behavior of the earlier agents suggests the urn is predominantly red. From this point, the total evidence available to each agent—their private evidence plus the evidence of the betting behavior of the previous drawers—favors red. As the sequence continues, successive agents become increasingly confident that the urn, which is actually predominantly white, is predominantly red; this despite the fact that the majority of their collective evidence supports the correct conclusion. The problem of information cascades can be reduced if individuals are overconfident: if they take their private information to have a disproportionate weight, relative to public information. In some conditions, groups can deliberate better due to the individualism of some of their members.

Even preferring one's own opinions to those of epistemic *superiors* can be epistemically productive, because self-silencing by those who recognize their inferiority may lead to "hidden profiles": information relevant to deliberation going unshared. In fact, groups may benefit from the arrogance of intellectual inferiors even when the information they insist on turns out to be misleading (Surowiecki 2004). Mercier and Sperber (2017) suggest that some of our individual-level reasoning pathologies may be adaptations for collective deliberation: the confirmation bias, for example, which leads us to overvalue evidence in favor of hypotheses we are well disposed to and undervalue contrary evidence, may conduce to an effective division of epistemic labor, given a diversity of views within the group. Similarly, I suggest, some degree of epistemic individualism—a disposition to underweight the views of others and overweight our own—may be an adaptation for collective deliberation (Levy 2019a).

Whatever the explanation for our epistemic individualism may be, that it's strong—perhaps especially among us WEIRD people (Henrich et al. 2010)—is beyond doubt. We're especially prone to embrace it with regard to science: our cultural paradigm of the scientist is the lone genius, working in isolation in his [*sic*!] lab, often misunderstood or ignored by his contemporaries. But this paradigm is deeply misleading. Scientists, too, are dependent on testimony, explicit and implicit, even *within the domain of their own expertise* (Rabb et al. 2019). Working scientists use tools they didn't develop (and that they may not be able fully to understand), often applied to data they didn't gather and which they can't verify, to test hypotheses that are constrained by theories they may not grasp. These constraints *enable* them to do science. Since climate change denial is a principal exemplar of bad belief in this book, let's take climate science as an example.

Climate science is heavily interdisciplinary: the models and the data that climate scientists use are the product of multiple researchers and research groups across multiple fields. The tools, the techniques, and the fields of expertise of climate scientists vary widely, from taking ice core samples to constructing mathematical models, from studying the physics of air circulation to estimating the extent of previous sea rises from geological evidence. Climate scientists calibrate their findings and refine their models in the light of evidence from other fields. Like every other scientist, moreover, they rely on tools they didn't design and which they lack the skills to fully understand. Central examples include the computers they use, and the mathematical tools they apply to build their models. They rely on these tools, despite having a limited understanding of how they work (Winsberg 2018). Moreover, each climate scientist works on a narrow area. A paleoclimatologist, for instance, has little expertise in Arctic sea ice or glaciology. But she may depend on experts in those areas for interpretation of some of her data. When it comes to assessing evidence that lies squarely within her competence, she will assess rival hypotheses in the light of prior probabilities that are sensitive to findings of other scientists in other fields (think of how a psychologist might assess evidence in the light of prior probabilities sensitive to information about evolution: if it is difficult to see how a particular psychological mechanism could have evolved, she'll think it unlikely that a mechanism of that kind can explain her observations).

Of course, the spectacular degree of interdisciplinarity seen in climate science is not evident across all sciences. But no science uses only tools it developed, or that its practitioners fully understand. Every science uses mathematical tools. While some capacity to understand these tools is a prerequisite for their competent use, it is routine for many practitioners not to fully grasp their details and near universal for their development from scratch to be beyond their capacity. Working scientists inherit these tools, which often owe their origins to work outside their field. Scientists also use physical tools they could not develop and often do not fully understand. Few neuroscientists can build or even troubleshoot fMRI machines. Less obviously, they rely on physicists and mathematicians for the development of the algorithms used to transform blood oxygenation signals into images of the brain. Every scientist working in the field of biology presupposes the truth of evolution ("nothing in biology makes sense except in the light of evolution," as Dobzhansky (1973) famously wrote). Epistemic dependence is routine in science: dependence on others for data, for tools and techniques, and for theories.

Science produces knowledge not *despite* this dependence, but *because* of it. The spectacular success of science is largely due to its structure as a set of epistemic institutions. This structure plays a central role in neutralizing, sometimes even harnessing, the influence of individual scientists' biases, whether their ideological preferences or the biases identified by psychologists (the confirmation bias, recency effects, the salience bias, and so on). Peer review, for instance, helps to ensure that scientific findings are subject to scrutiny from a variety of perspectives, each with its own biases. Since these biases often conflict (especially insofar as reviewers may be hostile toward the findings of the paper), their influence is mitigated. The confirmation bias—roughly, our disposition to seek confirming evidence for a hypothesis, and overlook disconfirming—isn't just mitigated: it's harnessed. Because each of us is motivated to search for evidence that favors our preferred theory, together we ensure that all our evidence is brought to bear (Mercier & Sperber 2011, 2017).

Of course, peer review subjects scientific findings to scrutiny by a small number of people, with a restricted range of disciplinary expertise. But the de facto institution of *post* publication peer review is very much more powerful. Every high profile paper is read by experts across a range of fields and is assessed for plausibility in the light of their disciplinary expertise; this collective scrutiny increases the likelihood that flaws in the gathering and interpretation of data are identified alternative hypotheses considered, and overlooked evidence considered. If the findings survive this collective scrutiny, they're transmitted to others and given some degree of credence.[9] Other researchers and research groups may attempt to build on them, or constrain their formulation of new hypotheses and their interpretation of data in the light of the findings. Others, of course, may not be convinced, and instead seek to develop and test alternative hypotheses. Even for these researchers, the findings constitute testimony: they take on trust that the reported experiments were conducted, with the populations described in the methods section, and the results are accurately reported.[10]

Peer review and post publication review is (somewhat) conflictual. But the success of science is also heavily dependent on cooperation. In climate science, we see different disciplinary specialisms tackling a

[9] This might best be seen as a special application of Goldberg's (2010) principle "if that were true, I would have heard it by now"—if there were serious problems with this paper, they would have been identified. Of course, many papers don't receive adequate post-publication review, so the principle applies only to higher-profile papers.

[10] This trust is not undiscriminating, of course. Statisticians have developed tools for the detection of p-hacking; that is, the manipulation of data to produce statistical significance (see, for instance, Head, Holman, Lanfear, Kahn, & Jennions 2015). Since p-hacking involves departures from the advertised methods (for example, by adding additional participants to tip a p-value below significance, or the failure to report unsuccessful trials), these tools ensure that data does not have to be taken entirely on trust.

common problem. Collaboration is also typical at the level of the individual lab or paper. In the sciences, single-authored papers are unusual, and the number of authors per paper is rising (Mallapaty 2018). Co-authors bring different skills and ranges of expertise to papers; often the skills and expertise of different fields. Some authors are engaged in data collection, others in experimental design, yet others in interpretation of the evidence. Some may be experts in techniques that the paper utilizes, and have little or no expertise in the topic itself. They must *trust* one another in order to collaborate, because they always lack the time, and usually lack the expertise, to assess the contribution others make. This is true even if they all belong to the same field. Hardwig (1985) gives the example of a paper in particle physics with 99 authors; each with slightly different expertise. Since then, the trend in physics has been toward much larger groups of authors; the (relatively) recent paper reporting the detection of the Higgs boson had more than 5000 authors (Castelvecchi 2015).[11]

The extent of relatively undiscriminating deference exhibited by scientists is best illustrated by their attitudes to a well-established protocol. They often have little idea of the causal contribution (if any) of every step. Why *that* amount of a solvent, rather than 5ml more, or 5ml less? Why *that* amount of time in the centrifuge? Especially if the experiment is expensive to run (in terms of resources or time), scientists have little incentive to attempt the kind of trial and error experimentation required

[11] Winsberg et al. (2014) worry that deeply distributed science threatens *accountability* for its epistemic achievements. When science is distributed across many individuals and across time, and no one fully understands the methodological choices made, no one "is in an epistemic and social position to explain and defend the methodological and epistemic standards employed in arriving at such a claim, as well as how they were met" (16). They do not, however, give a reason why any individual *should* be accountable in this kind of way. What is lost when accountability goes? This is a pressing question, because (contrary to what they seem to think) the deep distribution of knowledge production is not a recent phenomenon: if accountability requires being able to give an account of how knowledge was generated, or even what is consists in, then it was lost some time in the Paleolithic. Of course there are significant challenges in deeply distributed science, arising from the fact that (as they put it) there is no guarantee that the methodological choices made by the different individuals and groups who contribute to a larger research project, each of which involves epistemic trade-offs, "will form one coherent justification" (19). But that's a question about what they call the *reliability* of the research, not its accountability. It is a pressing question at the level of the individual research project, but in a well-organized and properly stress-tested science (one in which problems are tackled from multiple angles by groups with different interests, aims and agendas), we can expect these issues to be identified and resolved.

to discover whether certain steps might be dropped or compressed. Unless they have a particular interest in some step, they are likely simply to stick to the protocol. As Shea (2009) puts it, "all sorts of techniques and steps are copied without any appreciation of whether or why they are necessary to achieve the goal – following an experimental protocol can feel rather like following a magic spell" (2434). Faced with an established protocol, scientists behave like children shown a technique, rather than like chimps. Rather than innovate, they imitate.

Contemporary science doesn't individualize knowledge; if anything, it distributes it ever more deeply (Millgram 2015). It is because multiple individuals work on a common problem that we make progress on its solution, and these individuals are always themselves embedded in broader research cultures and research groups. The myth of the solitary genius is just that: a myth. In *The Enigma of Reason* (2017) Hugo Mercier and Dan Sperber probe a paradigm of that myth: Werner Heisenberg's withdrawal to a small island to think through deep problems in quantum theory. Heisenberg was, of course, a brilliant thinker. But he did not develop the uncertainty principle on his own. The advances he made were the product of collaboration and dialogue. They emerged through confrontation with the work of multiple other thinkers—Schrödinger, Bohr, and Dirac most prominently, but also other less famous physicists—and through a voluminous correspondence with Wolfgang Pauli. Heisenberg on his island was far more solitary than most scientists ever are, but even he was not alone: not intellectually. It was through dialogue and collaboration that he made his breakthrough.

We are the type of rational animal we are because we are cultural animals. We possess a suite of dispositions that orient us toward others epistemically. For us, knowledge production is in essential respects a product of the distribution of cognitive labor. Distributed cognition is more than merely a very efficient way of exploring the logical and empirical space, and mitigating, or even harnessing, our biases. It opens up new cognitive horizons that would otherwise be entirely inaccessible to us. Science does not free us, the animals we are, from epistemic dependence; if anything it increases it.