

3

How Our Minds Are Made Up

In the previous chapter, I surveyed a range of evidence that—I claimed—shows that knowledge production is a highly distributed enterprise, and that each of us is epistemically dependent on others for much of what we know. In the past 30 years, philosophers have come to recognize one facet of this reliance—our dependence on testimony—but have overlooked others. We know a great deal of what we know not only due to testimony, but also due to the ways in which epistemic labor is distributed and due to the ways in which knowledge is implicitly embedded in our social, and even physical, environment. It will be this embedding, and how it drives belief update, that will be the focus of this chapter.

Enlightenment Myths

Here's a caricatured and brutally compressed summary of the received view regarding the past few centuries of Western (*sic!*) thought: In the Middle Ages, intellectuals labored under the shadow of the Church. Only a few brave souls dared defy its orthodoxies. Even Galileo had to recant his heliocentrism in the face of the threat of torture, or worse, while Copernicus' *Des revolutionibus* was banned. Slowly, however, we emerged from this intellectual repression and into the radiance of the Enlightenment. We learned to think for ourselves, rejecting all orthodoxies, all authorities, except the authority of our own reason. "*Sapere Aude!*" "Have courage to use your own understanding!"—that is the motto of enlightenment," as Kant (1784/1991) famously wrote.

This caricature isn't merely simplistic, it's dead wrong. Reliance on our own epistemic powers is epistemically paralyzing. Alone we understand nothing. Knowledge is a social product. In this chapter, I'll focus on our epistemic deference in belief acquisition and update. We defer to

others so ubiquitously and so routinely we fail to notice when it occurs. We fail to notice our beliefs are dependent on what others believe, and shift as theirs do. We fail to notice this even when the changes are abrupt. These shifts may appear arbitrary and irrational. I'll show that they're neither. In deferring, we follow the evidence, though that fact may not be apparent to ourselves or to others. I'll begin in an apparently unlikely place: support for Donald Trump.

"Never Trumpers" were prominent Republicans who professed obdurate opposition to the candidacy of Donald Trump. They declared him unfit for office, on the grounds of his character and his policies. But most (though by no means all) Never Trumpers withdrew their opposition, and many became fervent supporters. This looks like belief update in the absence of evidence: Trump himself never provided any reason for them to think that their opposition was ill-founded.

Lindsay Graham, a long-serving Republican senator and himself a former candidate for the nomination, is an irresistible example. In February 2016, Graham described Trump as a "kook" who is "unfit for office." By November 2017, he sang a rather different song, now expressing his deep concern at the way the media stooped to portraying the president "as some kind of kook not fit to be president" (Lopez 2017). By April of the following year, Graham had decided that Trump deserved the Nobel Peace Prize for his reconciliation with North Korea.

Of course, Graham's about-face may reflect nothing more than the banal hypocrisy of politics. No one will be shocked to learn that politicians abandon professed convictions for political advantage. In any case, we shouldn't build theoretical castles on individual cases. But this kind of case isn't uncommon. Consider the swift abandonment of strongly held principles by white evangelicals. In 2011, white evangelicals were the group most strongly committed to the view that personal morality is an essential prerequisite for elected office. By 2016, they were *least* committed to that proposition (Kurtzleben 2016). This is surely a consequence of their preference for Donald Trump, a man who was brazenly immoral in his private life. A number of Evangelical leaders who had signed a letter condemning Bill Clinton on the grounds of his personal immorality, arguing that such immorality was more important than policies, now supported Trump and offered his policies as their reason (D. Miller 2018).

Recent examples of dramatic and apparently unprincipled changes in belief tend to occur on the political right. That's probably due to recent trends in world events, not anything intrinsic to being on the right. The candidacy and subsequent election of Donald Trump represented a watershed moment in US (and world) politics, for multiple reasons, but one in particular made such shifts more likely. Like all major political parties, the Republicans represent a broad swathe of somewhat conflicting interests: social conservatism (for instance) sits uneasily with free market fundamentalism, because the market is corrosive of traditional institutions (as we noted in an earlier chapter). Because parties yoke together different interests, candidates typically attempt to represent themselves as sensitive to the aims and desires of all of them. Trump decisively broke with this kind of within-party compromise, rejecting both the commitment to global trade and to at least a veneer of civility and respect for tradition that had characterized the self-conception of much of the party. As a consequence, those who identified as Republicans were forced to choose between party and these values. No doubt if the left was faced with as stark a choice, some would repudiate proclaimed core principles just as quickly.

Sudden shifts like these seem to occur irrationally, and in the absence of any evidence that could justify them. Trump's policies and behavior didn't change over the period during which Graham shifted from determined opposition to fervent support; not in any way that could explain the reversal. I'll argue that shifts like this typically occur as a result of our adaptive disposition to outsource belief to other agents. They manifest our nature as social epistemic agents. This outsourcing, moreover, is *rational* and evidence-responsive: it manifests our nature as *rational* social animals.

Outsourced Belief

In the last chapter, we surveyed the extensive evidence that human flourishing is very significantly due to our capacity to engage in distributed cognition. Human beings are distinctively dependent on cultural evolution for the development of the tools, techniques and practices that

enable us to colonize a dizzying diversity of environments. Cultural evolution generates cultural knowledge. It generates knowledge-how and it generates a great deal of propositional knowledge (“caribou skin must be harvested in the fall, because it’s too thin earlier in the year”). Cultural evolution relies on (or equips us with) a suite of adaptations that enable us to acquire beliefs from others, and thereby acquire the knowledge needed for flourishing. We also saw that science builds on these dispositions. It institutionalizes collective deliberation (for example, in the form of peer review), with distributed groups of scientists both deferring to and testing the hypotheses and findings of other groups. It depends on pervasive, but selective, epistemic deference. The Martian model of science introduced in the last chapter is false, at least here on Earth and perhaps anywhere there are beings capable of engaging in science.

Scientists outsource belief production and they outsource beliefs themselves. It is obvious how they outsource belief production: they depend on other scientists, both those working in their discipline and those in other disciplines, for the data and theories that are inputs into and constraints on their work. The outsourcing of beliefs comes in two varieties. First, scientists may rely on others to believe things on their behalf. Second, they may outsource beliefs they themselves hold by relying on others, or even the social or physical environment, to ensure that cues that trigger those beliefs are made available at appropriate times. Scientists outsource beliefs in this kind of way when they work in interdisciplinary groups. Science institutionalizes the outsourcing of beliefs. Members of different research groups rely on one another epistemically, thereby ensuring that beliefs that are more salient to some members than others are given their appropriate weight and role.

Scientists also outsource beliefs to machines and other tools. As is often the case, this kind of outsourcing is most easily seen when it goes wrong. In 1985, a team from the British Antarctic Survey reported the discovery of a hole in the ozone layer over the Antarctic.¹ It had been known for more than a decade that chlorofluorocarbons (as well as other

¹ I owe my understanding of this episode in recent scientific history to O’Connor & Weatherall (2019) who themselves report a debt to Oreskes & Conway (2011).

industrial chemicals) could deplete ozone. However, the mechanisms underlying ozone depletion had been intensively studied, and the consensus view was that the observations the BAS team reported were impossible. While scientists working on ozone depletion agreed that there was a need to phase out the use of CFCs, they didn't think that there was any urgency. A phased reduction would be sufficient and was well in hand. Hence, the BAS report was met with skepticism.

This skepticism was justified, in large part, on the basis that the theoretical models appeared to show that the rapid depletion required for the BAS report to be accurate was impossible. But it was bolstered by the fact that the BAS data conflicted with another set of data. The BAS team used measurement devices at ground level. A NASA satellite monitoring stratospheric ozone levels had failed to detect any significant change. Not only were there good theoretical reasons to think that the BAS data was wrong, there was also conflicting evidence; the likelihood was that the BAS team's data was due to some failure of measurement or processing.²

Nevertheless, Richard Stolarski, a physicist based at the Goddard Space Flight Center, decided to double check the satellite data. He discovered that in fact the satellite *had* detected the anomalously low levels of ozone the BAS team reported. But *because* these measurements were anomalously low, the software had removed them as outliers. This kind of preprocessing of data is a necessity in science, since bad readings and instrument failures are not all that uncommon. Building theoretical presuppositions into our equipment is therefore routine and a requirement of good science. We make implicit assumptions about what entities we will measure, and what the properties of these entities might be, and embody them in our instruments and our algorithms. The scientists had

² Anomalies like this are not all that rare in science. It is an extremely difficult question when the best response to an anomaly is to devote time and resources to attempts to resolve it and when it is better to ignore it. In this case, a relatively cheap and rapid review was sufficient to resolve the anomaly. If that investigation had not sufficed, it might have been rational to live with the anomaly, assuming that, given the theoretical models and the satellite data, the BAS measurements must be wrong, even in the absence of a satisfying explanation of how they came to generate their findings. Anomalies and unexplained errors are a fact of life: instruments glitch, people make mistakes in recording values or in adding them, and so on. If we make a mistake in ignoring an anomaly, the future progress of science will usually correct for this, because genuine data that cannot be satisfactorily explained will cause further trouble for our models.

encoded their theories about ozone depletion into their instruments, and because they'd done so, they had failed to detect data that conflicted with these theories.

The story has a happy ending. The theories turned out to be essentially accurate: the problem arose not from the scientists' understanding of how ozone interacted with industrial chemicals, but from a failure to factor in how extreme cold and the Antarctic winds influence the process. It also had a happy ending inasmuch as the problem of ozone depletion was solved by international agreement to phase out CFCs (despite a familiar story of industry resistance, epistemic pollution and US recalcitrance). In any case, it is an excellent illustration of how scientists outsource beliefs to instruments. Many scientists who relied on data from the NASA satellite had no idea how, or even that, the data was corrected. They could outsource beliefs about the need for correction (to account for measurement error) to the software.

Ordinary people outsource belief in both these kinds of ways too. We rely on others to believe things on our behalf, and we outsource our beliefs to the external world. Some of the ways in which we outsource belief are already familiar to epistemologists. Goldberg (2010) for instance has persuasively argued that we rely on others to maintain the stability of our beliefs concerning important events and well-known people. I can be said to know that the Queen is still alive because had she died today, I would have been told by now. But our epistemic dependence on others goes very much further. I rely on others to maintain my attitudes about all kinds of things. I even rely on them to help me know what my principles are. I outsource some of my most important and heartfelt values to the community of knowers. The fact that I outsource in this manner helps to explain how I could rapidly replace one principle with another.

Belief Is Shallow

Heartfelt beliefs (for present purposes) are those I feel strongly about and would defend fervently. They *feel* deep. Some of them *are* deep, in the sense that they're ingrained. They might be ingrained because they cohere with and are supported by other states that are constitutive of me

as a person.³ Perhaps they're ingrained in a different sense: they're somehow encoded such that it would be difficult to overwrite or even to inhibit them. Many beliefs are surely like that. But many fervently and sincerely expressed beliefs are much shallower, both in the sense that they are in some manner optional for me (my identity—again in a characterological sense—is not threatened by my abandoning them) and in the sense that as a matter of fact I would abandon them easily, perhaps without even noticing I was doing so, and even in the absence of (first-order) evidence against them.

The shallowness of belief is an instance of a more general phenomenon: the adaptive outsourcing of cognition to the world. Evolution is sensitive to small differences in costs, and it is cheap to outsource representations. Why go to all the trouble of building a model of the world, for instance, when the world is easily available to represent itself? As an added bonus, the world is a more accurate representation of itself than any model could be. It represents itself without any loss of detail, on a handy 1:1 scale, and updates in real time. As Rodney Brooks famously put it (1990: 5), “the world is its own best model.” So long as the costs of accessing the world are not significantly higher than the costs of accessing an internal model (and they may often be lower, not higher) we ought to expect the job of representing the world to be taken on by the world itself. That is, rather than consult an inner model, we should expect organisms to use sense perception to track the world and how it changes over time.

And that seems to be exactly what we find. Just as we are stubborn epistemic individualists, despite the fact that we defer to others pervasively, so we take ourselves to have rich internal representations, but actually retrieve much of our visual imagery from the world as and when we need it. We are subject to a fridge light illusion: we think our model of the world is rich and detailed because whenever we attend to any aspect of our model, we retrieve a rich representation of that aspect (see Chater (2018) for a lively presentation of the evidence; some of it will already be familiar to philosophers from Dennett (1992)). But our

³ Here “person” is used in a characterological, not a re-identification, sense; see Schechtman (1996) for this distinction.

internal models are sparse. For example, if a computer monitor is set up to display constantly altering garbage text, but with real words appearing and disappearing so that their appearance is timed to coincide precisely with our saccades as we read, we have the experience of reading an unchanging page (Rayner 1998). We take ourselves to have a rich and stable representation of the page, but we fail to notice how it changes. Change blindness experiments provide further evidence for the surprising poverty of our internal representations. In these experiments, images that are identical but for one (large) feature are presented to participants successively, interspersed with a flicker (to prevent retinal persistence). Even on repeated viewings of the same pair of images, the apparently obvious difference between the images is surprisingly hard to detect (Simons & Levin 1997). Change blindness has been demonstrated in the world outside the laboratory, with passers-by failing to notice when the stranger they were talking to was replaced by another who shared only the broadest similarity to the first (Daniel J. Simons & Levin 1998).

Daniel Simons, one of the researchers responsible for the work on change blindness just mentioned, has also been involved in work on inattention blindness that further illustrates how sparse our perceptual representations may be in the absence of attention. In a now famous experiment, Simons & Chabris (1999) had participants watch a video of two teams of players passing a basketball. In the easy condition, participants had to count all the passes; in the hard condition they had to keep separate tallies of passes for each team. While the teams were passing the ball, another person walked through the players, (apparently) in full view of the watchers. In one condition, the person wore a gorilla costume; hence the paper is called *Gorillas in our midst*. In the hard condition, only about half the participants noticed the gorilla. Even in the easy condition, around a third failed to notice the gorilla (in both conditions, they were more likely to notice a woman carrying an umbrella than the gorilla, perhaps because such an intrusion is less unexpected; nevertheless, collapsing across all conditions around a third of participants failed to notice the woman). Again, this seems evidence that our visual representations are less detailed than we'd have expected.

These experiments provide persuasive evidence of what we should, on theoretical grounds, have expected in any case: we offload aspects of

cognition onto the world. As Andy Clark, who has done more than anyone else to highlight the extent of our cognitive dependence on the world and on artifacts has argued,⁴ we believe in accordance with the 007 principle: keep things on a “need to know” basis:

evolved creatures will neither store nor process information in costly ways when they can use the structure of the environment and their operations upon it as a convenient stand-in for the information-processing operations concerned. That is, know only as much as you need to know to get the job done. (Clark 1997: 64)

Clark’s focus is on how tools and artifacts extend our cognitive capacities. There are, however, good reasons to think that beliefs are extended in analogous ways.

A central piece of evidence that our visual representations are extended consists in their *shallowness*, where a representation is shallow if it is easily uprooted. If beliefs are extended, then we should expect to see evidence that they are shallow in similar ways. We take ourselves to have rich and detailed internal representations of the visual scene, but careful experimentation shows that we actually retrieve a great deal of visual information by rapid and unconscious saccades as and when we need it. As a consequence, our visual representations are shallow: under appropriate conditions, we may fail to notice the substitution of one person for another, because our visual representation is (largely, though of course not entirely) stored outside our heads. Similarly, we take ourselves to have stable beliefs, but if our beliefs are shallow then we should expect them to be easily uprooted through analogous processes.

Cognitive dissonance experiments (reviewed in Cooper 2007) provide experimental evidence that our beliefs can be easily uprooted. In one

⁴ Clark’s account of cognitive dependence—which pays relatively little attention to social dependence—comes in defense of the extended mind: the thesis that minds are not confined to the skull of agents, but instead leak out into the world (Clark 2003; Clark & Chalmers 1998). The extended mind hypothesis is highly controversial, with a number of philosophers arguing that our cognitive dependence is better explained by our minds being *embedded* rather than *extended* (Rupert 2009). I’m not taking sides on this dispute. Holding that our cognition is heavily dependent on the environment and on other agents doesn’t commit me to holding that our minds are literally extended (nor to rejecting that claim).

variant of the induced-compliance paradigm, participants are asked to write essays defending a conclusion that they can be expected to find unpalatable. Often, college students are asked to defend the conclusion that their tuition fees should rise. Participants in the control group are paid to write the essays, while those in the experimental group are prevailed upon to write them (the experimenters might tell the participants that they can argue either for or against tuition rises, but that they don't have enough essays arguing in favor of a rise, and it would be nice of the participant to write such an essay). The manipulation results in those in the experimental group being significantly more likely than those in the control group later to express support for the conclusion they defended in the essay.⁵

Why should the manipulation induce those in the experimental group to be significantly more likely to endorse the conclusion they defended in the essay? The induced-compliance manipulation is designed to leave participants feeling free to refuse to write in favor of tuition rises (course credit or payment, for instance, doesn't depend on what they write). The perception they chose freely to write an essay defending tuition fees puts the participants in a very different situation to those in the paid group: they find it harder to explain their own behavior to *themselves*. Because they can't attribute their choice of topic to the effect of financial inducement or to constraint, they take their actions as evidence that they believe, or at least are not violently opposed to, the conclusion (Carruthers 2013). A simple manipulation of their own behavior is sufficient to induce a shift in their beliefs.

⁵ The classic work on cognitive dissonance stems from the bad old days in psychology, when sample sizes were often very small and questionable research practices abounded. For that reason, we shouldn't place too much weight on these findings. A multi-lab preregistered replication (using the essay writing task and an induced compliance manipulation) is nearing publication. Previous preregistered replications of some of the basic findings (e.g., Forstmann & Sagioglou 2020) provides grounds for guarded optimism with regard to the replication. Moreover, the researchers behind the multi-lab replication attempt report that they believe that the effect is real, albeit inflated in the published literature by the file-drawer effect (Vaidis & Slegers 2018). As we noted in the first chapter, researchers tend to have an accurate sense of which effects are real (Camerer et al. 2018), so their expectation is also (some) grounds for optimism. More recent work on choice blindness—soon to be discussed in the main text—suggests a mechanism for belief revision akin to that at work in these experiments, which is a third reason for optimism, or at any rate for believing that a mechanism like this one helps to cause belief revision.

The ease with which we are led to self-attribute beliefs by these kinds of manipulations suggests that we lack detailed internal representations of our beliefs. Just as our internal representations of the visual scene lack detail, and are easily swamped by changes in the external world so long as gross features are retained (*white man dressed as a construction worker*, as in Simons and Levin 1998), so our beliefs can be swamped by quite weak evidence that we believed something else all along. Notice, though, that this isn't just good evidence that we lack detailed internal representations. It is also good evidence of how we make up for this lack: by reference to cues for belief. In this case, the cues are our own behavior. In other cases (we'll soon see), these cues are the behavior of other people. Under a variety of circumstances, we will tend to construct or reconstruct our beliefs on the spot, rather than recalling them, with attention to cues to belief central to this reconstruction.⁶

We saw above how change blindness experiments provide striking evidence for the sparseness of our perceptual representations. *Choice* blindness experiments provide parallel evidence for the sparseness of our beliefs, and, like the evidence from cognitive dissonance, suggest that beliefs about our own values and commitments are much shallower than we'd have expected. In choice blindness experiments, participants make a series of choices between pairs of options, and their responses are recorded. For instance, the options might be represented by cards and the chosen card placed in a pile. After the participants have completed the series, the experimenters show them the options they picked one by one, and ask them to justify their choices. On a minority of questions, however, the choices have been switched and participants are asked to justify a choice they didn't in fact make. Participants rarely notice the switch (despite being given the opportunity to correct their responses). Instead, on most switched trials the participant goes on to defend choosing an option they had actually rejected. They attribute

⁶ Constructing beliefs rather than recalling them is in fact routine. As Gareth Evans (1982) influentially noted, we (typically) don't answer questions like "do you believe there will be another World War?" by searching through a repository of our beliefs, but instead by considering the world. In the kind of case I have in mind just as much as in Evans' case, we construct our beliefs by assessing the evidence for them. Evans focuses on first-order evidence (are geopolitical tensions rising?), but higher-order evidence (what do people like me think? What do *I* think?) is evidence too.

to themselves a belief that is at odds with one they'd expressed just a few minutes earlier.⁷

Choice blindness has been demonstrated in multiple spheres, from judgments of facial attractiveness (Johansson et al. 2005) to moral judgments (Hall et al. 2012). Choice blindness has been demonstrated with regard to political opinions too. In the lead up to a Swedish general election, researchers approached members of the public and asked them about their attitudes to the actual policies of the parties contending in the election (Hall et al. 2013). They were asked to express their attitudes to each policy by marking a point on a 100mm line, with one end indicating complete agreement and the other complete disagreement. They were then asked to justify their choices, including some choices which had been switched (the median alteration changed the point marked by 35.7mm). Only 22 percent of manipulated responses were detected and more than 90 percent of respondents offered at least one justification of an altered choice. Self-reported political engagement and the use of extreme ends of the scale didn't correlate with likelihood of detecting a switch. Similarly, most participants in a study of hot-button moral issues offered justifications of choices they hadn't in fact made (Hall, Johansson, & Strandberg 2012).

In these studies, just as in cognitive dissonance experiments, belief revision was induced by manipulating cues to what the person themselves believes. In others experiments, belief revision is induced by manipulating cues to what *others* believe (taking advantage of our disposition to engage in social referencing). In one experiment, Maoz, Ward, Katz, & Ross (2002) presented Israeli Jews with a real draft peace proposal. Jews who were told that it had been drawn up by Palestinian representatives were significantly less favorable to it than those who were told that it had come from the Israeli side (which was in fact the truth). In a second experiment, Maoz et al. showed that Palestinians too were susceptible to devaluation of a proposal on the basis of information about who authored it. Cohen (2003) reports even stronger evidence: for his sample, information about whether welfare policies were supported by House Democrats or House Republicans was a more powerful

⁷ Note that choice blindness has been replicated with a massive sample (Rieznik et al. 2017).

predictor of attitudes to it than policy content. Much more recently, Barber and Pope (2019) showed that endorsement by Trump powerfully influenced Republican attitudes to a policy, independently of its content. Liberal policies supposedly endorsed by Trump were 15 percent more likely to be supported by those on the political right than those which were not (supposedly) endorsed by him (conversely, being told that Trump *opposed* a policy didn't move liberals).

These studies appear to demonstrate that many beliefs are shallow, in much the same kind of way as perceptual representations are shallow. I take myself to have rich internal representations of the world around me, but in fact these representations are quite sparse: the fact that I will easily accept substitution of one obtrusive element in my visual scene for another without noticing is evidence that my state lacks the detail I take it to have. I find it hard to hang on to the visual scene, because in fact my representation of the world is *out there*, in the world: I rely on the stability of the world to ensure my representation doesn't wobble. Similarly, I take myself to have fixed and definitive beliefs. But I rely on the world (where the "world" includes my own body) to tell me what I believe and to ensure both stability and detail in my beliefs. If the world changes, my beliefs change with it, sometimes without my noticing.

Outsourcing and Belief Shift in the Real World

The evidence just cited is evidence that our belief states are often much less rich and much less stable states than we would have guessed. While of course we have internal belief states, they are remarkably shallow. When we ask ourselves what we believe, we look as much to the world (especially the social world) to answer the question as to our own internal states. In the face of evidence that *that I believed that p*, or that *people like me believe p*, I conclude *I believe p*, and I may do so even if previously I had quite a different belief, and without noticing the shift.⁸

⁸ No doubt, alternative explanation, consistent with our having rich internal models, might explain the data. But our interpretation of experimental results should be constrained by other data and the theories that explain them. We should expect to see outsourcing to the world, when conditions are appropriate (when, for instance, our perceptual access to the world is

Let's turn from the laboratory to the real world, to see these mechanisms of belief attribution at work. Take the case of Never Trumpers once more. This is a disparate group of individuals, and no doubt different mechanisms are at work in different members of the group. But social referencing might explain belief revision in some of these cases (or it might partially explain it, by explaining how the revised belief came to be more palatable or more accessible). When Never Trumpers made their initial declarations, Trump was one candidate among others for the Republican nomination (Graham's February 2016 proclamation was typical). At that point in the nomination process, social referencing didn't result in a strong signal of support for any one candidate. Instead, insofar as someone who strongly identified with the Republican party asked themselves "what do people like me believe?" the response was divided among the possible candidates. Because the signal was divided, other sources of evidence and motivation had to take up the slack: Republicans had to make up their minds who to support by reference to a familiar array of considerations (electability, personality, policies, personal identification, and so on) with social referencing playing a more minor role. As the primary season progressed, however, candidates fell away one by one, and the signal constituted by the support of other Republicans for the remaining candidates strengthened. The signal supporting Trump in particular may have grown in strength as a consequence of the numbers attending his rallies (and as a consequence of the media's constant highlighting of these numbers).

Once Trump was declared the presumptive nominee in early May, the signal returned by social referencing was unified: *people like me (largely) support Donald Trump*. Given that we discover what we believe, in important part, through such referencing, the signal constituted pressure to fall into line. And as Never Trumpers fell away, the signal increased in strength. Once Trump was elected in November 2016, the signal gained additional strength: *people like me support the President*, especially when he comes from my party.

sufficiently rapid and reliable). We saw in the previous chapter that there is a strong case for thinking our species owes much of its success in colonizing a great variety of environments to cultural evolution, and the way in which it distributes epistemic labor. It is in light, too, of this picture of ourselves that we ought to interpret the evidence.

This kind of mechanism supports rational belief update, moreover. We'll have a lot more to say about the rationality of mechanisms like this one (which I'll argue are common) in later chapters. For the moment, let's just note that this kind of mechanism manifests *ecological* rationality (Todd & Gigerenzer 2012). A cognitive process is ecologically rational when it is designed to get the right answer, no matter how it does it. Mental shortcuts are the paradigm cases of ecologically rational mechanisms: it might not look rational to guess that one foreign city has a larger population than another because the first is more familiar to you, but the recognition heuristic works, at least in some domains (Goldstein & Gigerenzer 2002). Social referencing is rational in this kind of way. Recall from the last chapter that groups of deliberators are (under appropriate conditions) more accurate than individuals. Recall that knowledge is distributed across individuals, and that deference to group norms is adaptive given this fact. Recall, too, the prestige bias: other things equal, it is adaptive to imitate the beliefs (and behavior) of prestigious individuals. Given all these facts, there was *rational* pressure for someone like Graham to update his beliefs. The majority of people like me support Trump; prestigious individuals with whom I identify (or up to whom I look) support Trump; therefore *I* support Trump. In later chapters, I'll argue that social referencing is rational directly: it's not just a way of reaching the right result; it's a way of reaching the right result by responding to evidence.⁹

In many cases, this shift will occur without any phenomenology or other indication of a change of mind. Instead, the person may take themselves to have always believed what they now believe. Our memories are highly fallible. Perhaps this fallibility arises from the same causes that explain the fragility of our beliefs and our perceptual representations: because they're very sparse. There is extensive evidence that memory is not a snapshot of the past, with details preserved, but is instead

⁹ The recognition heuristic itself might work in this way and therefore qualify as directly rational. A mechanism is ecologically rational if it is well-designed to generate the right response (in an appropriate environment). It is directly rational if the cues to which it responds are not merely correlated with getting the result; they're *evidence* for it, sufficiently strong to justify the response given. Recognition might be evidence for city size; hence relying on it may be relying on justifying evidence.

sensitive to the circumstances of retrieval. As a consequence, memory exhibits shallowness parallel to that of other representations. Loftus (2003) has shown how vulnerable our memories are to false information, with participants in her experiments incorporating into their memories information that was actually presented after the events. Schacter (1996) cites evidence that this kind of effect is not confined to the lab, but plays a role in the confabulation of sometimes spectacularly false and highly consequential memories: memories of being the victim of abuse, and even of being an abuser, for instance. It would not be entirely surprising, therefore, were some people to shift from fervent opposition to equally fervent support of a politician without taking themselves to have changed their minds.

Of course, a prominent Never Trumper like Graham won't be allowed to forget he was once a fervent opponent of the president. The media and political foes will delight in reminding him. But the sparseness of memory and of representations more generally may nevertheless aid someone like him in reconciling himself to his shift. He may imagine all sorts of mental reservations, for instance. Sparseness leaves room for such confabulation: probing his mental states, he'll fail to find a unified, obdurate, *Never Trump!* representation. Rather, he'll find little content beyond the words themselves, and he'll have room to confabulate all kinds of caveats (*Never Trump, unless* he were to become more presidential; *unless* he's endorsed by Ted Cruz, or what have you). For others, who are less likely to be reminded of previous views, the shift comes more easily. Those White Evangelicals who suddenly dropped the view that excellence in character was essential in a politician may well see themselves as holding a single position over time.

It bears repeating that this mechanism of belief update is not partisan. It's one we should expect to see deployed across the political spectrum. People will differ in who they defer to, and what counts as a good cue to belief, but they will engage in analogous kinds of social referencing. Social referencing may help explain relatively rapid shifts in opinion on left and right. Take the relative swiftness with which gay equality has achieved widespread acceptance in most developed nations. It wasn't long ago that gay people were routinely mocked and caricatured on mainstream television. Within my memory, such mocking became unacceptable to most

people, and being gay became relatively unremarkable, as the candidacy of Pete Buttigieg for the Democrat nomination, and his subsequent cabinet selection by Joe Biden, illustrates.¹⁰ Polling data on marriage equality provides quantitative evidence for this decline in prejudice (Pew Research Center 2019). Opposition to marriage equality halved, from 60 percent of US adults in 2004 to 31 percent in 2019, while support doubled over the same time period (31 percent/61 percent). The shift is too big and rapid to be very significantly explained by a cohort effect, and in fact we find shifts in every cohort, including the oldest. Of course, there are multiple explanations for this shift, and (first-order) evidence and argument surely played some role. But social referencing probably had some effect too. Insofar as people settle what they believe by asking themselves *what do people like me believe?*, changes of mind can snowball. The more people shift to believing p , for whatever reason, the more *people like me* believe that p , for any value of “people like me.”¹¹ Shifts in belief can help to cause further shifts in belief, and we may see very rapid changes in attitudes.

All that said, it’s possible that this particular mechanism tends to operate more strongly on the right than the left, today at any rate. It is sometimes held that the contemporary left is characterized by identity politics, rather than by identification with a broader movement. If this is right, or if (for some other reason) identification with the Democratic party or any other broad institution is weaker,¹² we might find that the effects of social referencing are weaker, because for those on the left there are fewer people (sufficiently) *like me* to provide a strong signal for

¹⁰ This is of course not to suggest that anti-gay prejudice is no longer a huge problem. Rather, the point is that it has decreased to a very significant extent in the population, in a relatively brief period of time.

¹¹ Of course, this is the case only when the issue is not strongly polarized. If only one side of a debate shifts, then asking oneself *what do people like me believe* will yield divergent answers for (say) Republicans and Democrats, and we may see snowballing on one side and intransigence on the other (think, for example, of how attitudes to Russia moved in lockstep on the left and right until recently, when they began to diverge sharply; Suls 2017). While attitudes to same sex marriage remain somewhat polarized, the degree of polarization is relatively small and Pew reports large shifts on both sides of politics.

¹² One piece of evidence in favor of thinking that identification with a single bloc is weaker on the left than the right: media consumption patterns are much broader on the left than the right, with the former accessing both highly partisan media and more mainstream sources, while the latter confines itself much more to partisan sources (Benkler et al. 2018).

belief update, or because signals are spread across multiple groups. If that's the case, of course, it's not because people on the left are more rational or that their beliefs are explained by different kinds of mechanisms. The same mechanisms will operate in different ways and with different effect in different contexts.

The Pervasiveness and Rationality of Outsourced Belief

Ever since Putnam (1975) drew our attention to the phenomenon, it has been a familiar fact that we rely on experts to fix the reference of some of our concepts. I don't need to know how to distinguish elms from beeches because, when relevant, I can exploit the expertise of relevant experts to make the distinction for me. In fact, this kind of deference is very common. My "elm" concept is adaptively indistinct. Rather than having fully specified contents, our representations often include *placeholders* or *pointers*, where a "placeholder" is a gap or indistinctness that requires specification (if it becomes necessary to deploy a specified concept) and a "pointer" is a placeholder that indicates where specification may be sought (Rabb et al. 2019). As Rabb et al. recognize, scientific concepts are typically adaptively indistinct in just this kind of way.

It's not just scientists who defer like this, of course. On the basis of testimony from their parents or teachers, children come to accept many claims that they can't distinctly represent. They may accept that *germs cause diseases* while having little idea what "germs" are, for instance. These indistinct representations may scaffold the later acquisition of more precise first-order representations (Sperber 1996), but many of us never come to replace our indistinct first-order representations with anything more precise. Most of us have little idea what "viruses" are, even while the world is gripped by a pandemic we know to be caused by one. Many people believe *that the theory of evolution is true*, but experimental evidence indicates that many of those who take themselves to "believe in" evolution have little genuine understanding of the theory or the mechanisms involved (Shtulman 2006). This indistinctness is not a bug, but a feature: given that genuinely understanding the science is extremely difficult (such that even experts can master only some of the

relevant knowledge), our having indistinct concepts allows us to defer to the experts to provide specification. But experts engage in this deference too: theories outside their area of expertise constrain their work, sometimes in ways that require them to allow others to provide details in their own representations.

Indistinctness of content is also common in the domain of political and moral beliefs. One reason to think that our political commitments are indistinct is that most of us sign up to political programs that are amalgams of several conflicting currents, apparently without really noticing the conflicts. For instance, mainstream conservatives support political parties that are strongly in favor of free markets and of traditional institutions like the family. But as we previously noted, free markets are solvents that undermine traditional institutions. Mainstream center-left parties often combine irreconcilable currents within them too (a commitment to the environment and to traditional, highly polluting, blue collar jobs, for instance). Even at an individual level, there is great indistinctness of content. As Schwitzgebel (2009b) notes, concepts like “freedom” and “brotherhood” and propositions like “all men are created equal” are widely affirmed, but are “only half-filled or quarter-filled with a real thought” (57). They are grasped indistinctly, and await outsourcing to others to get a distinct content.

This outsourcing leaves us vulnerable to sudden shifts, of the sort seen in the case of Evangelical Christians or Never Trumpers. Our disposition to outsource in this kind of way constitutes a vulnerability, because once others know what cues we respond to, they can seek to exploit us by exploiting these cues (more on this in later chapters). But *vulnerability* doesn't entail *defect*. Being vulnerable to others in this kind of way may be adaptive. There are no free lunches, and every adaptation has costs. Our vulnerability may be an inevitable cost we pay as epistemically social animals, and being epistemically social is responsible for much of our capacity to generate significant knowledge. Outsourcing is ecologically rational. It is also *directly* rational: being open to cues for what others believe is being open to *reasons*. A proper defense of this claim must await later chapters. For now, let me emphasize how recognizing the rationality of these kinds of dispositions should change our attitudes to our fellow citizens.

The media delights in telling us how ignorant we are. Often this takes a partisan form (we've all seen the videos of Trump supporters making outrageous claims; we rarely pause to wonder how many people they had to ask to get the response they wanted), but it also occurs in a non-partisan, if equally self-congratulatory, way. The surveys reported are often badly designed—or well designed to get the kind of answers that will generate headlines—but the underlying finding is robust: most of us know surprisingly little about history, geography, or contemporary politics. Some people know a lot, but “most people know nothing, and many people know *less* than nothing” (Brennan 2016: 24). Brennan (2016) surveys some of the lowlights of this ignorance: even during election years, most people cannot identify the candidates standing in their district; most cannot name the party that controls congress; they routinely report the opinion that the US spends too much on foreign aid, instead suggesting that a percentage of the budget that vastly *exceeds* actual spending would be appropriate; 40 percent can't name the United States' enemies during World War II; at the height of the cold war, only a minority correctly stated that the Soviet Union was not a member of NATO. More than a third of actual voters are “know-nothings” Ilya Somin (2013) claims. Such ignorance of basic facts concerning the issues central to elections is by no means confined to the United States. A survey of people who voted in the United Kingdom's Brexit referendum revealed that neither Leave nor Remain voters knew much about the EU; on some questions they performed *worse* than chance (Carl et al. 2019).

Voters' ignorance concerning the issues at stake in elections is central to Brennan's (2016) case for epistocracy: government by the informed. It is easy to sympathize, in the light of election results in Brazil, the United States, India, Turkey, and elsewhere over the past five years: isn't this what government chosen by the ignorant looks like? But it would be a mistake to blame the apparent rise in populism on voter ignorance alone: voters have *never* been well-informed about the issues or the parties; instead they make their decisions by looking to signals from elites (J. R. Zaller 1992). That's not irrational: that's adaptive outsourcing of belief. Voters may not know how their representatives promote their values, but they know who to look to for such promotion. They may not know what policies the parties offer, but they often have “metaknowledge”

(Rabb et al. 2019); pointers to repositories of information that fill in the gap. Even if they lack such metaknowledge (having mere “placeholders”), they can outsource cognition by deferring to experts who have more detailed knowledge (sometimes, deference may appropriately be to those with metaknowledge, rather than those who know).¹³

Possession of metaknowledge isn’t necessarily inferior to possession of knowledge. We live in a complex world, and we all daily engage with systems and use concepts we cannot fully specify. None of us is expert in anything more than a small range of domains. Metaknowledge unaccompanied by knowledge may not merely compensate for our deficits; it may often be *better* than knowledge. While the genuine expert may be cognitively better equipped than the person who has metaknowledge without knowledge, the person with metaknowledge alone may be better equipped than the person who has *some* knowledge, but falls short of expertise. The person with metaknowledge may track changes in expert opinion much more reliably than the person with knowledge short of expertise, since the former defers to the experts for the contents of their beliefs.

It’s in this light that we should interpret our use of concepts like “freedom” or “equality,” when they’re quarter-filled with a real thought. It’s intrinsically difficult to fill them with a thought—that is, to specify what they mean. Philosophers who specialize in these topics disagree with one another, and would probably agree that their preferred view has costs, inasmuch as it fails to capture dimensions that others emphasize. Only specialists can hope to make these concepts anywhere near fully distinct. Yet we non-specialists may *appropriately* be committed to them. Outsourcing specification to the experts may be the rational way for us to proceed. It’s not a limitation on our rationality; it’s a way in which we manifest it. In a slogan, meta is (often) better: outside the very limited sphere in which we can acquire detailed knowledge, we often do better and cognize more rationally by knowing how to find out whether *p*, or knowing who knows whether *p*, rather than by knowing

¹³ Again, this outsourcing leaves us open to exploitation by those who can make use of it, of course. The degree to which we’re vulnerable to exploitation is limited by the fact that we remain sensitive to first-order evidence regarding how well we’re doing: voters turn on the governing party when economies falter, for instance. We integrate our higher-order evidence with the first-order evidence that is sufficiently near and clear for us to make use of it.

much about *p*. Sometimes, we even do fine by deferring in the absence of metaknowledge.¹⁴ We are all very ignorant, and we should be fine with that.

Of course, there are kinds of ignorance that really are costly. Our ignorance of our pervasive reliance on the outsourcing of cognition and on metaknowledge might be a costly kind of ignorance. In fact, we seem to mistake our outsourced knowledge for individually possessed knowledge and think of distributed cognition as inferior. Paradoxically, our sense of how much we know *individually* is sensitive to our metaknowledge: in response to cues that others know, and cues that we can access testimony, we inflate our estimate of how much we know *without* the need to rely on others (Rabb et al. 2019). Access to the internet increases people's confidence that they can answer questions *without* using the internet; the knowledge that one's partner in an experiment knows something inflates a person's confidence that *they* know it; people rate *their* sense of understanding higher when they are led to believe that experts understand something compared to when they are not (importantly, however, this effect was seen only when they were told that this knowledge could be accessed), and so on. We are subject to a pervasive *illusion of explanatory depth* (Mills & Keil 2004; Rozenblit & Keil 2002), arising from a similar mechanism. The illusion of explanatory depth refers to our tendency to overestimate the extent to which we understand and can explain the workings of everyday objects and natural phenomena. People express confidence in their ability to explain how flush toilets and piano keys work, or what causes tides and rainbows, which greatly outstrips their actual ability to explain these things. Keil & Wilson (2000) suggest that this illusion arises from the pervasiveness of the division of epistemic labor and our effortless facility with it. So ingrained is deference and so fluent that we confuse knowledge possessed by our community with knowledge possessed by us.¹⁵ At the

¹⁴ Consider how we defer to the geographical knowledge of taxi drivers or pilots. We may lack metaknowledge in some such cases: we may not even know *who* is in charge.

¹⁵ I once suggested that conspiratorial ideation may arise, in part, from the same mechanism: our vulnerability to mistaking community knowledge for individual knowledge leaves us with an inflated sense of our capacity to explain complex events; in some individuals, this leads to suspicion of a coverup when the official narrative does not strike us as fully satisfying (Levy 2007).

conclusion of this chapter, we'll return to this kind of ignorance. While first-order ignorance may not be a problem, higher-order ignorance—ignorance, perhaps, of the very fact that first-order ignorance is not a problem—may play a role in our current epistemic malaise.

How to Make Up Your Mind

I've suggested that the pervasive outsourcing of belief should lead us to rethink our attitude to the apparent ignorance of our fellow citizens. The most consequential example of such apparent ignorance is, of course, ignorance of climate change. Understanding the mechanisms underlying belief acquisition and update should lead us to rethink ignorance of climate change too. Contrary to a widely accepted story, it doesn't arise from stupidity or irrationality. It's a product of mechanisms responding to evidence. Understanding these mechanisms will provide us with levers for changing people's beliefs, including their beliefs about climate science.

Why do so many people reject the science of climate change? Because the social mechanisms of belief update provide them with good reasons to do so. Centrally, they deploy social referencing, asking *what do people like me believe?* Multiple cues tell them that people like them reject the science (think of how advertisers and merchants of doubt play on cues to identity and identification: true Americans drink Coke; true Americans are individualists who won't be told what to think by outsiders). Kahan thinks of mechanisms like this as identity-protective and philosophers have followed his lead. For instance, Carter and McKenna (2020) argue that we have strong grounds to be skeptical about beliefs acquired via these mechanisms, because they are responsive to "the epistemically irrelevant property of being socially approved." But that's a mistake: the fact that a proposition is socially approved is higher-order evidence that bears on its truth, and there's nothing irrational in being guided by it. The primary purpose for which we deploy these mechanisms is to get things *right*, not (just) to fit in.¹⁶ So too with other mechanisms

¹⁶ In defence of a view somewhat similar to Kahan's, Williams (2021) suggests that an epistemic account like mine cannot explain why we see worse performance from those who score

that seem partisan—for instance, giving more credence to testimony from those who are like me than those who are not. There is experimental evidence that we give more weight to the explicit testimony of those we see as benevolent toward us (P. Harris 2012; Mascaro & Sperber 2009; Sperber et al. 2010). This bias is rational because those who don't share my values may seek to exploit me, and those on my side are likely to be more trustworthy (toward me). This kind of mechanism may generalize beyond explicit testimony to the implicit testimony constituted, for example, by *failing* to attend to certain messages (thereby signaling they're not worth attending to), or simply by getting on with one's life without worrying about the climate. If those I identify with don't worry about climate change, that's a reason—a genuine *reason*—for me not to worry. They provide me with evidence.

Asking oneself *what do people like me believe* is one mechanism whereby people come to reject (or to accept) climate science. There are, of course, others, and they work in concert for many skeptics. They also ask *what do prestigious people believe*? Different people accord different degrees of prestige to different individuals. For liberals, the relevant individuals may be scientists (think of the role Dr. Fauci played in the pandemic), as well as politicians on their side. Conservatives get a very different message about climate change from opinion leaders on their side: most notoriously from Donald Trump, but also from most other major figures in the Republican party, as well as from the apparent experts they see on Fox.

Evidence of consensus is significant for both sides. Of course, there's a consensus *among scientists* about climate change, but the consensus that is most weighty for an individual might be the apparent consensus that is reinforced every day: the apparent agreement of those around them. Who we interact with day to day plays an important role in what we believe. Bob Inglis, a former Republican congressman who lost the backing of his party and his seat over his endorsement of climate science, credits his conversion from skepticism to belief in climate change

higher on tests of knowledge and capacity. One possible explanation is that greater sophistication allows for better deference: those who know more know better what they ought to believe in virtue of their ideology. See Martin and Desmond (2010) for a suggestion along these lines.

partly to visits to US bases in the Antarctic, where scientists showed him the evidence (Cohn 2013). There's every reason to think that the (first-order) evidence played an important role, but living in close proximity to the scientists and thereby coming to establish a basic level of trust with them likely also played a role. We pay attention to cues for consensus (P. Harris 2012), but since our best evidence for consensuality will often be what those around us say, those nearest to us will once again have an outsize role in what we take to be consensual (mere repetition can induce the impression that something is widely believed; Weaver et al., 2007).¹⁷

The environment also contains multiple non-verbal cues to belief; cues that are also genuine evidence. It's hard to sustain belief that we face a crisis when those around us are relaxed. It's hard to sustain the belief that we must dramatically reduce our consumption of fossil fuels when the gas station is central to your daily life, or when your job or those of your friends depends directly or indirectly on oil. It's also hard to sustain belief in a climate crisis when it is cold outside, and a number of studies have found that belief in climate change increases during hot spells and falls during cold, even among those who distinguish between weather and climate (Borick & Rabe 2017; Hornsey et al. 2016). Perhaps the sheer solidity and robustness of the built and natural environment makes it hard for us to believe we face a crisis (perhaps the coronavirus pandemic might unsettle our taken-for-granted assurance that the world will simply continue as before sufficiently to make us more receptive to worries about the bigger challenge represented by climate change).

Bad belief in climate science has—formally—the same source as better belief. Those who accept the science see scientists as prestigious and the institutions that employ them as reliable. Those who do not have much less positive views, both of scientists (trust in science has been

¹⁷ The most effective content-based intervention for producing better beliefs on climate change seems to turn on making perceptions of the scientific consensus more accurate (since a consensus of genuine experts is strong higher-order evidence, that's not surprising). Perception of the consensus is a "gateway" to more accurate belief (S. L. van der Linden et al. 2015, 2019). But it's not sufficient simply to present that information: in the face of multiple dissenting voices, informing people about the true degree of consensus doesn't correct people's misperceptions. It's necessary to "inoculate" people by first presenting information rebutting—or prebutting—denialist talking points (S. van der Linden et al. 2017).

declining on the right since the 1970s; Gauchat 2012) and of universities (58 percent of people who are Republican or Republican-leaning now say that colleges and universities have an overall negative effect on the United States, compared to just 19 percent of those who are Democratic or Democratic leaning; Pew Research Center 2017). We pay attention to those we trust, those whose values we share (or who we take to share our values), those who are prestigious or simply near to us, and we believe what they believe. Doing so is rational: it's a mechanism for calibrating belief to evidence, albeit indirectly. Like all mechanisms, however, it can misfire, and in the case of climate change it has gone spectacularly wrong for very many people.

If you want to change minds—your own or those of others—you need to look to the social and institutional cues on which beliefs depend. Change those, in the right ways, and minds follow. If social referencing yields a particular answer to the question what do *people like me believe about X*, then I'll tend to accept that answer as my own. If cues to belief are distributed across the social and institutional environment in certain ways, then those who are sensitive to cues like these will tend to accept the associated beliefs. Change minds by changing the world: physical and social. That, I suggest, is not only how minds are most effectively changed. It is how minds have *always* been changed, at least with regard to matters outside the immediate purview of the agent.

Outsourcing Religion

I'll finish this chapter with a brief digression from the main argument, to address a question that was posed, but not answered, in Chapter 2. In that chapter, we briefly surveyed evidence of behavioral inconsistency among religious believers. Christians appear to be more likely to bid in a charitable auction on a Sunday than on other days of the week (Malhotra 2010); consume less pornography on a Sunday than other days (Edelman 2009) and think of the dead as living on only when given a religious prompt (P. Harris & Giménez 2005). Christians are by no means unique in this kind of apparent inconsistency, which is also seen in Hindus (Xygalatas 2012), Muslims (Duhaime 2015) and in the Vezo people of

Madagascar, who practice an ancestral religion (Astuti & Harris 2008). Van Leeuwen (forthcoming, 2014) thinks these inconsistencies are best explained by differences in the kind of attitude people take to religious claims. In response, I pointed to evidence that religious credences tend to govern behavior more broadly than Van Leeuwen thinks, and that therefore they qualify as beliefs. But if behavioral inconsistency is *not* explained by the nature of the underlying attitudes, what does explain it? I suggest that it might arise from the outsourcing of cognition.

We outsource not only to people and cognitive artifacts, like books and the internet, but also to the environment: for example, we outsource perceptual representations to the world around us. We surely outsource religious representations to other people: *I* don't need to understand the nature of the Trinity, because I can rely on theologians to do so on my behalf. But perhaps we outsource religious representations to features of the environment too. Whereas when I outsource my visual representations, the world takes on the role of (partially) constituting the representation for me, outsourcing of religious representations is more likely to play the role of providing prompts for more purely internal representations (call such triggers *cues*). It's adaptive to outsource in this way, because it increases the efficiency with which cognitive resources can be managed. If I can be confident that my representations will reliably be cued when they're needed, I don't need to be vigilant for situations to which they're relevant. This kind of outsourcing can explain inconsistencies in behavior across time: in the absence of cues, my behavior won't be guided by a representation that would otherwise be relevant to me.

Religious representations can be expected to differ in the degree to which they're cue-dependent and the range of cues to which they're sensitive. They might also differ in how long they persist once triggered, and these differences may themselves be a function of how reliably they're cued. Consider the apparent difference in persistence of the Sunday effect (the dispositions of Christians to behave consistently with their religious creed to a greater extent on Sundays than other days of the weeks) compared to the effects of the call to prayer on charitable giving. The call to prayer appears to have short-lived effects on behavior, with a rapid drop off in its effects on charitable giving (Duhaime 2015).

A mere 20 minutes after it ended, the percentage of those choosing the most charitable option had dropped from 100 percent (while the call was audible) to less than 50 percent. In striking contrast, the Sunday effect appears to persist long after Church services have ended.

It's unlikely, I think, that this difference represents a deep difference in Islam versus Christianity. More plausibly, it is the result of a difference between the availability of cues to religion in a very secular society, like the contemporary United States, versus a more religious society like Morocco. In the contemporary US, reminders of religion are much less frequent than in Marrakesh. Because reminders of religion are less frequent, believers need to internalize them to a far greater degree: once they're cued (for instance, by the realization that today's the Lord's Day), they persist longer and decay slower than elsewhere. For those who live within the hearing of a mosque, by contrast, the call to prayer is available five times a day, every day, to take up the representational slack. There's much less need to internalize the representations. In medieval Europe, I speculate, church steeples and bells might have played an analogous role, and the Christians of Paris or Prague may have had fast-decaying religious representations, like contemporary stallholders in the Souks. Conversely, Muslims in Melbourne or Manitoba would have persistent religious representations, more like their Christian counterparts than their co-religionists in other countries.