Neil Levy

Radically Socialized Knowledge and Conspiracy Theories

ABSTRACT

The typical explanation of an event or process which attracts the label 'conspiracy theory' is an explanation that conflicts with the account advanced by the relevant *epistemic authorities*. I argue that both for the layperson and for the intellectual, it is almost never rational to accept such a conspiracy theory. Knowledge is not merely shallowly social, in the manner recognized by social epistemology, it is also *constitutively* social: many kinds of knowledge only become accessible thanks to the agent's embedding in an environment that includes other epistemic agents. Moreover, advances in knowledge typically require ongoing immersion in this social environment. But the intellectual who embraces a conspiracy theory risks cutting herself off from this environment, and therefore epistemically disabling herself. Embracing a conspiracy theory therefore places at risk the ability to engage in genuine enquiry, including the enquiry needed properly to evaluate the conspiracy theory.

One of the markers of intellectual seriousness is widely taken to be the adoption of a very sceptical attitude toward conspiracies theories, the attitude summed up in the well-known saying "If you have to choose between explaining something as a cock-up or a conspiracy, choose cock-up every time". As a number of philosophers have now demonstrated, this general attitude is neither epistemically warranted, nor is it actually adhered to in the epistemic practices of responsible intellectuals. Conspiracies are a common feature of social and political life, common enough that refusing to believe in their existence would leave us unable to understand the contours of our world; moreover, warranted belief in conspiracies is widespread, even among the intellectuals who confidently reject this or that putative explanation of an event as 'just a conspiracy theory'.

I shall argue in this paper that the actual epistemic practices of intellectuals are more rational than their expressed adherence to 'Hanlon's razor' would suggest.[1] In practice, it is not conspiracy theories toward which intellectually responsible people evince a reflexive suspicion. Instead (and as Coady (2006) has suggested) it is conspiracy theories which conflict with (the right kind of) official stories that come under suspicion. Few responsible intellectuals reject the explanation of 9/11 that cites the conspiratorial actions of a group of terrorists under the direction of Osama Bin Laden. Despite the fact that this explanation is clearly a conspiracy theory (an explanation of an event in terms of the plans and coordinated actions of a secretive

https://doi.org/10.3366/epi.2007.4.2.181 Published online by Cambridge University Press

group of conspirators), most of us have little doubt that it is true. But we do reject alternative explanations that cite the secret plans and coordinated actions of agents of the United States or of Israel as 'mere conspiracy theories'. The manner in which we express our refusal to entertain these claims, our contemptuous dismissal of them as mere conspiracy theories, is unfortunate and misleading, but the rejection itself is epistemically rational. A conspiracy theory that conflicts with the official story, where the official story is the explanation offered by the (relevant) *epistemic authorities,* is *prima facie* unwarranted.

I have two main goals here. First, I shall explain why it is generally irrational to reject the account of an event or process offered by the relevant epistemic authorities. Because knowledge is *deeply* – constitutively – social, I shall argue, the epistemic authorities, when properly constituted, are far better positioned to explain events than are isolated agents. Individuals who are not appropriately hooked up to the relevant social network of knowledge production are ill-placed to take issue with official stories. The social distribution of knowledge production is not guaranteed to produce truths, but it is a powerfully truth-*conducive* mechanism of explanation production, such that when an explanation of an event or process conflicts with that produced by this socially distributed mechanism, we ought to reject it as unwarranted. Second, I shall explain not only why we ought to accept the official story over alternatives, but also why many people overestimate their ability to detect flaws in the official story, and therefore seek alternative explanations.

1. RADICALLY SOCIALIZED KNOWLEDGE

A number of philosophers have recently emphasized the degree to which knowledge production is a social enterprise (Schmitt 1994; Goldman 1999). This work serves as a very important corrective to the excessive individualism that is typically characteristic of epistemology. However, though the claims made by social epistemologists, as they style themselves, are (largely) true and (frequently) important, they do not go far enough toward socializing knowledge and knowledge production. Most social epistemologists seem to think that knowledge production must be a distributed enterprise in compensation for the epistemic limitations of human beings. Because we do not have the time to check every claim, or to gain the expertise to allow us to even make a start, we must defer to experts. While this is undoubtedly true, knowledge is much more deeply social than it suggests.

Social epistemology, as it is generally understood, remains individualist in the following sense: it assumes, or at least is fully compatible with, a picture of knowledge production in which cognitive processing takes place within the heads of individual agents. It is then transferred, by means of testimony, to other agents, in the form of symbolic communications deciphered in the heads of these agents. The social, on this view, is the communicative link between individuals. Granted, the social makes possible a division of cognitive labor that makes knowledge production proceed far more quickly and efficiently than it otherwise could. Each individual is able to be a cognitive

https://doi.org/10.3366/epi.2007.4.2.181 Published online by Cambridge University Press

specialist, working on his or her own small section of the cognitive fabric, confident that their efforts will, if successful, enter into relations of mutual support with the fruits of the cognitive labor of other agents. It is this picture of the social distribution of knowledge production that motivates Keeley's (2006) worry that too ready a belief in conspiracy theories leads all too easily to scepticism: if we begin to distrust the network that transfers knowledge from agent to agent, we shall be left unable to tap into the knowledge of others, and thrown back upon our meagre epistemic resources. Keeley's worry has been given short shrift in the literature on conspiracies (Clarke 2006; Coady 2006). It is, I think, more serious than those who dismiss it recognize: acceptance of the kind of conspiracy theory that Keeley has in mind tends to promote a generalized paranoia and distrust (hence the common observation that people who accept one unwarranted conspiracy theory frequently accept further such theories, and increasingly wild ones). In any case, the kind of social epistemology from which the worry stems is nevertheless only social in a shallow way. In fact, knowledge production is deeply social, in a way that few epistemologists have recognized.

Individualist epistemology presupposes individualist philosophy of mind. On the individualist picture, cognitive agents each possess rich and detailed internal representations and models of the world; cognition refers to the set of ordered state-to-state transitions between these representations and models. On this picture, cognition is usually thought to be computational, and internal representations are thought of as having their own syntax in virtue of which these computational processes can get to work on them. On the socialized epistemology sketched above, the results of these rule-governed processes might then be transferred from individual to individual. Now, it may be that human cognition really does proceed, at least in part, by way of representations in a language of thought (though I think the evidence for this is weaker than is generally recognized). But it is certainly false that each of us possesses internal models of the physical or social world that are anything like as rich or as detailed as individualism suggests. Or rather, if we have any such models, they are not located where the individualist thinks they are. Rather than being found within the head of the individual, they are located *outside*, in the world. Each of us typically thinks he or she (in isolation from others) possesses a far richer model and a far more detailed explanatory scheme than we do. Individualism is intuitive, but it is wrong.

Consider, in this context, the phenomenon of change blindness. We each think that we possess a detailed internal representation of the visual scene. For instance, you probably think that right now you have a pretty detailed internal representation of the page in front of you. In fact, you do not. Rayner (1998) showed that subjects reading text on a computer screen were entirely unaware that the only genuine words on the screen at any one time were those they were reading. If alterations in the text were timed carefully, so that the real text always fell within the subjects' reading window, the rest of the page could be replaced by gibberish without the subjects noticing anything. Simons and Levin (1998) provide a far more graphic illustration of the same phenomenon. They stopped passers-by and asked for directions. While the subject was answering the query, two confederates dressed as workmen carried a door

https://doi.org/10.3366/epi.2007.4.2.181 Published online by Cambridge University Press

between the experimenter and the subject, giving the experimenter time to slip away and be replaced by an entirely different person. Most subjects failed to register the substitution of one person by another; even when they were prompted they denied noticing anything unusual. This, despite the fact that standing side-by-side no one would have trouble distinguishing the original and the substituted person.

What these experiments, and literally hundreds more on change blindness, apparently demonstrate is that though we seem to have detailed inner representations of the outer world, we do not. We do have detailed representations, but they are not *inner* representations. We represent the world to ourselves not by way of an internal image, but by using an external model: the world itself. Rather than take a snapshot of the scene and store it internally, we rely upon the actual stability of the world. We store our representation *outside* us. We are not aware of the fact that we do this because we are not aware of the way in which our internal representations are constantly updated by our eye movements. The human eye has a very small area of high-resolution vision: less than 0.01 percent of the entire visual field. But our eyes constantly dart about, moving this window of high resolution across the visual scene. These movements, called *saccades*, are intelligent; they are not random, but instead gather information relevant to the tasks currently confronting the person. They are also very fast, averaging about three per second. Our frequent and repeated saccades allow us to inspect the world and to update our picture of it, so that it seems to us that we have a rich representation of it. And so we do, but it is not an internal representation. In a sense, this is unsurprising. Why build a model of something when the original is there to be used as its own best model (Clark 1997)?

Now, there is a sense in which our use of external representations might be seen as a concession to our epistemic limitations. We do not build detailed internal models of the world because doing so would be costly, where 'costs' are measured in the kinds of considerations to which evolution is sensitive. Had we unlimited attentional and cognitive resources, we might build detailed internal models. But using external representations is not *just* a concession to our limitations: it is an effective and efficient strategy for representing the world, which has a number of benefits apart from the fact that it requires fewer resources to implement and maintain than the individualist alternative. For one thing, an external model of the world is far more responsive to changes in the world than an internal one would be: all alterations in the modelled world are made in real time, for the simple reason that the modelled world is its own model. Moreover, the efficiency gains are so great that we ought to expect this kind of approach to be selected even by beings with far greater cognitive resources than we possess: given that certain conditions are satisfied (in particular, the world is stable enough to serve as its own model) these resources are far better expended elsewhere.

When we store our representations externally, we use a publicly accessible medium. Our mental representations are *shared*. In this sense, they are far more deeply social than anything envisaged by traditional social epistemology. Moreover, the storage of our representations of the world outside us, in the space we share with others, is but one example of a ubiquitous phenomenon. Cognition is very often dependent upon

https://doi.org/10.3366/epi.2007.4.2.181 Published online by Cambridge University Press

environmental resources, and the more technical and complex the kind of cognition, the more it is reliant upon such external scaffolding. Human beings are able to perform very simple arithmetic in their heads; indeed, we may have an innate number sense which allows us to automatically and effortlessly process small numbers (Wynn 1998). But even slightly complex arithmetic – not to mention more arcane mathematics – is heavily dependent upon external resources. Most of us cannot multiply even quite small numbers – say, three digit numbers – without relying upon calculators or pieces of paper. Some people can perform such multiplications, but it's simply a party trick: there is nothing to be gained from performing these calculations internally. It makes sense to offload performance of the algorithm, and save our cognitive resources for where they're needed (Rowlands 1999). Notice, moreover, that this extended process is not merely dependent upon physical scaffolding, pieces of paper, manipulable tokens, and the like. It is also essentially dependent upon other *people*. The algorithms we use in performing these calculations (perform a series of single digit multiplications, carry the one, and so on) are not things we worked out for ourselves. Typically, we simply accepted, on faith, that they work. Thus, our cognitive process is reliant not only on pencil and paper, but also on the intellectual products of previous agents, the agents, long since dead, who developed the algorithms (as well as the chain of agents who handed them down across the generations).

Of course, this feature of our algorithms is typical: most of our significant environmental resources for thinking are the product of many hands. They generally embody the expertise of an uncountable number of agents – uncountable, because each of these agents was reliant upon yet other agents, both directly and indirectly. Think of the expertise embodied in the computer: the expertise not only of computer scientists, but also of electrical engineers, physicists, mathematicians and logicians, industrial scientists who manufacture the components (not to mention those who manufacture the machines used in the manufacturing), and many others, in a network which extends across the world and also reaches back into time, to the beginnings of science and systematic knowledge.

There are two important points here. First, that our cognitive abilities are not just shallowly social, not merely dependent for their success upon cumulative repositories of knowledge and on the division of cognitive labor, but deeply social, in that they are actually *constituted* by their embedding in the social world. Second, that our reliance upon these external resources is so complete, so much a part of the background of our thought, that it is often entirely missed. Individualism is intuitive, though it is deeply mistaken. I shall comment briefly on both these points in turn.

The constitution claim can be expressed as follows: entirely new cognitive horizons open up when we learn to rely upon extended resources. We couldn't do mathematics if we did not rely upon the embodiments of the knowledge of others. For a dramatic illustration, consider how metarepresentational abilities become available to apes when they lean upon environmental supports. The unadorned chimp brain is able to learn to categorize pairs of objects on the basis of their similarity or difference from one another. So, for instance, they can be trained to put any pair of identical objects

https://doi.org/10.3366/epi.2007.4.2.181 Published online by Cambridge University Press

– two cups, say, or two bananas – into one box, while placing any pair of dissimilar objects into another (say one cup and one banana). But without external aids, they cannot sort *pairs of pairs* by similarity or difference. Two pairs of pairs are identical just in case they share their first-order properties: they are either both identical, or they are both dissimilar. Otherwise, they are a dissimilar pair. So the pair (of pairs) apple-banana and cup-shoe is identical, while the pair cup-cup and apple-banana is different. This higher-order task is difficult enough for the human brain. Chimps can learn to accomplish it by, in effect, turning it into a first-order task. They do this by learning to associate tokens with the first-order pairs. For instance, they might learn to associate a plastic triangle with a pair of objects that are identical to one another, and a plastic square with a pair of objects that are dissimilar. Once they have accomplished that task, the higher-order task is simple. If you want to know whether a pair of pairs is similar or dissimilar, simply compare the tokens associated with them: if they are identical, so is the higher-order pair, if not, they are dissimilar (Thompson et al. 1997). Now, just as this metarepresentational task first becomes possible for chimps when they learn to associate concepts with tokens, so for us too higher cognition becomes possible when we learn to turn metarepresentational tasks into first-order tasks, or complex tasks into simple ones. We perform complex mathematics by turning the task into a series of simple arithmetical calculations (single digit multiplications, for instance). Our extended resources simplify our cognitive landscape, allowing us to reach peaks and valleys that would otherwise remain hidden to us.

But, and this is the second point, we systematically underestimate the extent of our reliance upon extended resources; correlatively, we overestimate our own internal resources. Folk responses to demonstrations of change blindness is one illustration of this: people are surprised to learn that they suffer from it. Here's another illustration, one that is more directly relevant to the way in which knowledge is reliant not only upon the external environment, but also upon other agents and their expertise. People suffer from an almost ubiquitous *illusion of explanatory depth*: they systematically overestimate the extent to which they understand and can explain the workings of everyday objects and natural phenomena (Rozenblit and Keil 2002). Subjects express confidence in their ability to explain how flush toilets and piano keys work, or the causes of tides and rainbows, which greatly outstrips their actual ability to explain these things. Wilson (2004) argues that part of the explanation for this illusion stems from our unthinking reliance upon the division of cognitive labor. Our reliance upon the expertise of others is so extensive and so automatic that we take ourselves to possess the expertise itself. In a sense, I think we *do* possess the relevant expertise: just as I really *can* perform five digit multiplications, though I would be at a loss if I had to rely upon my unadorned brain, so I really do know what causes tides if I am able to rely upon Wikipedia to summon the explanation. But of course *that's* not the sense of 'knowledge' that's usually in question when we ask whether someone knows what causes tides.

We now have in hand all the ingredients to explain (a) why it is usually irrational to disbelieve the official story, and (b) why there is a strong temptation to disbelieve it.

https://doi.org/10.3366/epi.2007.4.2.181 Published online by Cambridge University Press

2. OFFICIAL STORIES

The official story is the story promulgated by the authorities. Coady, who introduced the phrase into the debate, did not distinguish between the different kinds of authorities that might promulgate an explanation of an event. Clearly, it is often rational to heavily discount the official stories offered by some authorities. In totalitarian countries, people learn to read the official news media with a jaundiced eye, and this attitude is often warranted. Recent events in Anglophone Western democracies demonstrate that this kind of attitude toward the official stories promulgated by governments and by their sycophants in the media is all too often warranted in non-totalitarian countries. But there are other kinds of authorities beside governmental authorities, and other kinds of official stories. Though it is not epistemically irrational to reject official stories *per se*, there is a class of official stories that, other things being equal, we ought to accept. Responsible believers ought to accept explanations offered by *properly constituted epistemic authorities*.

Moreover, I claim, when intellectuals evince, in their actual practice, suspicion of conspiracy theories, it is generally theories which conflict with the explanations offered by these epistemic authorities they reject. The same intellectuals who dismiss such explanations as mere conspiracy theories may frequently reject official stories offered by government; they may even go on to postulate conspiratorial explanations of why these official stories are given wider coverage than they deserve. The official story offered by the United States government is that global warming is 'just a theory', in the pejorative sense of 'theory'. This official story has influenced policy debates, in the United States and elsewhere, though there is evidence that constructing and disseminating it required the cooperation of a group of well-placed conspirators in government and in private industry.[2] Responsible intellectuals, even when they are themselves citizens of the United States, do not accept *these* official stories, nor do they cast opposition to them – even opposition which, as I have just done, postulates the existence of a group of conspirators as a causally necessary element in their promulgation and dissemination – as 'just a conspiracy theory'.

It is sometimes, perhaps often, rational to doubt the official story when the issuing authority is the government. But responsible intellectuals do not contemptuously dismiss a proffered explanation as 'just a conspiracy theory' when it conflicts, merely, with the government line. Instead, we reject such explanations when they conflict with the official story of the properly constituted *epistemic* authorities. It is not because the government tells us that the attacks of 9/11 were carried out by Al Qaeda operatives that we dismiss rival explanations, according to which the government itself was behind the attacks, as just a conspiracy theory. It is because the relevant epistemic authorities – the distributed network of knowledge claim gatherers and testers that includes engineers and politics professors, security experts and journalists – have no doubts over the validity of the explanation that we accept it. When there is a conflict between official stories, between the explanation offered by the political authorities and that offered by the epistemic authorities, responsible intellectuals are

https://doi.org/10.3366/epi.2007.4.2.181 Published online by Cambridge University Press

ready to believe the latter (regardless of whether either explanation cites the actions of conspirators). We reject the former explanation even if (as is typically the case) we cannot ourselves even begin to evaluate the rival claims for ourselves. Simply knowing that a proffered explanation conflicts with the official story (where, once again, the relevant authorities are epistemic) is enough for us rationally to reject the alternative.

Who are the properly constituted epistemic authorities, and why is it epistemically irrational for us to reject explanations which conflict with those they offer? Talk of an authority being 'properly constituted' suggests a political or institutional process. I mean the talk of constitution to be taken more literally: an epistemic authority is properly constituted when it has the right kind of structure. I cannot hope to fully elucidate this structure here (indeed, I doubt anyone has proposed a entirely satisfactory account of such a structure anywhere). I shall have to content myself with a few remarks. The right kind of structure is that exemplified by science: knowledge claims are the product of a socially distributed network of inquirers, methods and results are publicly available (especially, but not only to other members of the network), inquirers are trained in assessing knowledge claims according to standards relevant to the discipline, and rewards are distributed according to success at validating new knowledge *and* at criticizing the claims of other members of the network. Epistemic authorities are properly constituted to the extent to which they consist in a distributed network of agents, trained in assessing knowledge claims, who make their evidence and processes available to scrutiny, within and beyond the network.[3]

Knowledge produced by such a network of inquirers is *deeply* social. The distributed cognitive enterprise produce knowledge that is *in principle* inaccessible to isolated individuals. It is not merely a concession to our cognitive limitations that we pursue our most highly prized knowledge through such networks. We utilize these networks because they open up parts of the landscape that would otherwise remain forever inaccessible. Since these distributed networks are a necessary means to certain kinds of knowledge, cutting ourselves off from the networks and means of knowledge production is not merely cutting ourselves off from testimony, and it does not merely breed scepticism and distrust. It is, far more radically, cutting ourselves off from our own best epistemic techniques and resources.

Of course, cutting ourselves off from the properly constituted epistemic authorities is not cutting ourselves off from our intellectual inheritance: many of our epistemic resources are merely *genetically* social, in the double sense that their development required the work of many agents, widely distributed in time and space, and that their acquisition requires that the child be appropriately socialized. Once they are acquired, we are no longer reliant on social mechanisms for the maintenance of cognitive resources that are only genetically social: we can perform arithmetic on a desert island.

But newer techniques and resources remain *constitutively* social: we do not have access to the cognitive landscapes that they open up if we cut ourselves off from them. The division of cognitive labor, of which traditional social epistemology has made so much, is in fact just as illustrative of the way in which knowledge is constitutively social

https://doi.org/10.3366/epi.2007.4.2.181 Published online by Cambridge University Press

as it is of the more traditional point that knowledge is more efficiently pursued when results are shared and research is specialized. The scientist, or indeed the philosopher, who can simply assume that the results of others have been sufficiently tested to be accepted without further examination is able to push on into new areas. Indeed, this paper is an illustration of how this works: by assuming (without checking) the claims of cognitive scientists, I am able to make claims about the temptations of conspiracy theories, and why it is rational to resist these temptations.

I am now in a position to sketch why we ought generally to accept official stories, where the relevant authorities are the epistemic authorities. By "we", I mean those of us who aim to be epistemically responsible and who do not possess expertise *directly* relevant to assessing the official story. We non-experts – and we are all non-experts with regard to most official stories – ought simply to defer to properly constituted epistemic authority. We cannot responsibly assess them: we often lack the tools even to understand them, much less to poke holes in their fabric. We may often find them counterintuitive or implausible, but that's only to be expected: since socialized epistemic resources open up features of the cognitive landscape that remain hidden to those who rely upon their unadorned brain – however smart that brain – we should expect that there will frequently be a conflict between true explanations and our folk intuitions. It is rational for us to accept the testimony of the relevant experts (Levy 2006). The (directly relevant) expert, on the other hand, generally ought not to defer to the experts. She ought to maintain a high degree of epistemic humility, recognizing that when her conclusions in her precise area of expertise conflict with those of her peers, the chances are that she's wrong. Nevertheless, she may have a *ceteris paribus* duty, to herself and to the truth, to pursue her own conclusions, even when she possesses the requisite degree of epistemic humility.

The intellectuals who embrace explanations of the kind that we typically and pejoratively label conspiracy theories are almost never in possession of the directly relevant expertise. They may be experts in *something*, but rarely do they belong to the class of enquirers with the authority to issue official stories regarding the event to be explained.[4] It is this class of intellectuals for whom the consequences of embracing conspiracy theories are most serious.

When the layperson embraces a conspiracy theory, she does not risk her cognitive resources. Most of the resources she employs are of the kind that are only genetically social, or relatively shallowly social. Believing that, say, the Holocaust is a myth, or that the Oklahoma bombings were carried out by US government agents, does not threaten her command of these cognitive resources. But for the expert, much more is at stake: she threatens to cut herself off from cognitive resources that are deeply social, and which (alone) make possible cognition at the cutting edge. Since these deeply social resources open up new cognitive landscapes, cutting oneself off from these resources disables one for conducting genuine research. It cuts oneself off from the very means one requires to be a responsible intellectual, and therefore, ultimately, to be able to evaluate the conspiracy theory one accepts.

Since there is a holism of knowledge claims, and the official story enters into

https://doi.org/10.3366/epi.2007.4.2.181 Published online by Cambridge University Press

relations of mutual support with other knowledge claims, doubting the official story tears a hole in the web of distributed knowledge. It places at risk not merely the social relations of testimony and trust, as Keeley has argued, but also the very techniques and resources of knowledge acquisition. It leaves the conspiracy theorist unable to rely upon a growing body of research and tools: not only the studies that support the view she rejects, but also the studies upon which it relies, and the studies it supports in turn. It leaves her doubtful of the techniques those studies employed. It throws her back upon her own cognitive resources – and, no matter how clever she is, no matter how educated, these resources are meagre.

In fact, the conspiracy theorist herself may implicitly recognize the feebleness of the unadorned human brain. She typically seeks alternative networks of knowledge production. But these networks are shortlived and vulnerable to all kinds of distortions. They cannot compensate for the riches that our shared epistemic resources, embodying millennia of history and the labor of many thousands of agents, make available to the less sceptical agent.

The radically socialized view of knowledge sketched here provides an explanation for why it is irrational to reject the official story. Moreover, it also helps to explain why conspiracy theories are tempting. Recall the illusion of explanatory depth: we consistently underestimate the extent to which our knowledge depends upon our location in the socially distributed network of epistemic authorities. We take ourselves to be able to understand more, far more, by ourselves than we are really capable of. Hence we take ourselves to be able to detect flaws in the official stories, flaws that the epistemic authorities have either, inexplicably, overlooked, or from which they have deliberately turned. We take the conflict between our intuitions and the explanations offered by the epistemic authorities as evidence that the latter are stupid or base rather than recognizing that the conflict is the predictable consequence of our lack of access to the relevant cognitive tools.

Prime the mechanism with one or more of the heuristics and biases – the confirmation bias, the clustering effect, and the availability error all seem likely candidates – and the feeling that the official story is a fabrication that we see through can be almost overwhelming. The result can be the promulgation of alternative explanations that are wildly at variance with the evidence.

There are no simple formulas that can be applied to yield an overwhelming preponderance of (significant) truths over falsehoods. Every rule we might suggest to guide responsible belief formation is liable to have exceptions. Nevertheless, the following maxim will, I suggest, guide us far better than most: adjust one's degree of belief in an explanation of an event or process to the degree to which the epistemic authorities accept that explanation. Sometimes, of course, the epistemic authorities will be wrong. Nevertheless, since they will usually be best positioned to discover the truth, and because the alternative is cutting oneself off, to a greater or less extent, from the very resources one needs to correct one's beliefs, accepting the official story is almost always rational. Implementing this strategy is not a trivial task, for two reasons. It requires, first, that one be able to discover the degree to which the relevant explanation

https://doi.org/10.3366/epi.2007.4.2.181 Published online by Cambridge University Press

is accepted by the epistemic authorities (which requires, in turn, that one can identify these authorities, or identify those who can), and that is not always easy. Second, it requires a measure of epistemic humility that is far greater than we are accustomed to, or accustomed to counting as a virtue. When it comes to knowledge, we each do best by cultivating only our own garden: that relatively small sphere in which we can claim some expertise. The big picture we ought to leave to take care of itself.[5]

REFERENCES

Bloch, A. 1980. *Murphy's Law Book Two: More Reasons Why Things Go Wrong!* Los Angeles: Price Stern Sloan.

Borger, J. 2005. "Ex-oil lobbyist watered down US climate research." *The Guardian* June 9. Retrieved September 19, 2007 from http://www.guardian.co.uk/oil/story/0,11319,1502486,00.html

Clark, A. 1997. *Being There: Putting Brain, Body, and World Together Again.* Cambridge, MA: The MIT Press.

Clarke, S. 2006. "Conspiracy Theories and Conspiracy Theorizing." In D. Coady (ed.), *Conspiracy Theories: The Philosophical Debate*, pp. 77–92. Aldershot: Ashgate.

Coady, D. 2006. "Conspiracy Theories and Official Stories." In D. Coady (ed.), *Conspiracy Theories: The Philosophical Debate*, pp. 115–28. Aldershot: Ashgate.

Goldman, A. I. 1999. *Knowledge in a Social World*. Oxford: Clarendon Press.

Keeley, B. L. 2006. "Of Conspiracy Theories". In D. Coady (ed.), *Conspiracy Theories: The Philosophical Debate*, pp. 45–60. Aldershot: Ashgate.

Levy, N. 2006. "Open-Mindedness and the Duty to Gather Evidence." *Public Affairs Quarterly* 20(1): 55–66.

Rayner, K. 1998. "Eye Movements in Reading and Information Processing: 20 Years of Research." *Psychological Bulletin* 124(3): 372–422.

Rowlands, M. 1999. *The Body in Mind: Understanding Cognitive Processes*. Cambridge: Cambridge University Press.

Rozenblit, L. and F. Keil. 2002. "The Misunderstood Limits of Folk Science: An Illusion of Explanatory Depth." *Cognitive Science* 26(5): 521–62.

Schmitt, F. E. (ed.) 1994. *Socializing Epistemology: The Social Dimensions of Knowledge*. Lanham: Rowman & Littlefield Publishers.

Simons, D. J., and D. T. Levin. 1998. "Failure to Detect Changes to People During a Real-World Interaction." *Psychonomic Bulletin & Review* 5(4): 644–9.

Thompson, R. K., D. L. Oden and S. T. Boysen. 1997. "Language-Naive Chimpanzees (Pan troglodytes) Judge Relations Between Relations in a Conceptual Matching-to-Sample Task." *Journal of Experimental Psychology: Animal Behavior Processes* 23(1): 31–43.

Wilson, R. A. 2004. *Boundaries of the Mind: The Individual in the Fragile Sciences.* Cambridge: Cambridge University Press.

Wynn, K. 1998. "Psychological Foundations Of Number: Numerical Competence In Human Infants." *Trends in Cognitive Sciences* 2(8): 296–303.

https://doi.org/10.3366/epi.2007.4.2.181 Published online by Cambridge University Press

NOTES

1 Hanlon's razor is a more general claim, of which "cock up before conspiracy" is a specification: "Never assume malice when stupidity will suffice." This saying is attributed to Robert J. Hanlon ('cock up before conspiracy' is itself due to Bernard Ingham) (Bloch 1980).

2 See Borger (2005) for one piece of evidence among many.

3 The media is, or is part of, a properly constituted epistemic authority. Why did it let us down so badly with regard to Iraq? Part of the reason is that the evidence upon which we went to war was *not* publicly available, either within or beyond the network. Journalists and the experts they rely upon were unable to scrutinize the evidence to test the claims made by government; in addition, journalists may have been seduced by the glamour of secret intelligence into credulity. In addition, it should be obvious that the media is only a very rough approximation to a properly constituted epistemic authority. For a range of reasons, some to do with the nature of democracy, some less creditable, the network is in some ways more open and in others more closed than a properly constituted epistemic authority should be, and the rewards for journalists are distributed on the basis of a number of criteria, with the production of knowledge ranking less high than one might want.

4 Scholars for 9/11 Truth is typical here. Its membership includes philosophers, attorneys, professors of English and French, of cultural studies and politics, but few engineers and no structural engineers, even though the most infamous claim associated with the group is that the Twin Towers *couldn't have* been brought down by the impact of the jets. Similarly, academics prominent in Holocaust revisionism include Robert Faurisson, who specializes in twentieth-century French literature, and Arthur R. Butz, a professor of electrical engineering, but very few academic historians.

5 I would like to thank David Coady for criticisms which greatly improved this paper.

Neil Levy is a Senior Research Fellow at the Centre for Applied Philosophy and Public Ethics, University of Melbourne, and a James Martin 21st Century School Research Fellow at the University of Oxford. He is the author of five books and many articles on applied ethics, free will, philosophical psychology and social epistemology.

https://doi.org/10.3366/epi.2007.4.2.181 Published online by Cambridge University Press