

Hard to Believe

Robert Herritt

You are enjoying after-work drinks with friends when, two rounds in, the conversation turns to a contentious policy issue. Maybe it is the effects of raising the minimum wage or the best way to organize the healthcare sector. An informal debate takes shape. Predictably, smartphones are drawn, as the combatants search for fresh ammunition. One cites a decorated economist writing in the *New York Times*. Another reads from some think tank's factsheet on the subject. Others point to the personal blog of a well-credentialed policy analyst, or a piece of journalism that claims to provide "everything you need to know" about the topic at hand.

Such disputes are rarely settled. Instead, for non-experts, disagreements over technical topics often devolve into claims that "my source is better than yours"—the *Wall Street Journal* is hopelessly biased, the *New York Times* a model of objectivity; my preferred Nobel-laureate economist is a disinterested advocate for the truth, yours a partisan obscurant. Of course, it is likely that one of the views being advanced is closer to the truth than the others. But by the time the conflict becomes a competition between arguments from authority, the chances of a conclusive victory are slim. After all, what would such a victory even look like?

For those who relish having their beliefs vindicated, such happy-hour stalemates can be dismaying in their own right. But the scenario described above points to a more fundamental problem: when it comes to forming beliefs about specialized subjects, good strategies are thin on the ground—at least for laypeople, a group of which I often consider myself a member. In the end, on any number of subjects, most of us must rely on our own hopelessly flawed judgments when deciding which views to endorse.

At worst, the necessity of choosing beliefs in such an imprecise way can lead us to see well-established truths as mere matters of taste, with more than a little help from our latest information technologies. Those seeking to reject the expert consensus—whether on the health risks of vaccines or the validity of evolution by natural selection—find themselves equipped not only with easy access to information that conveniently reinforces their favored views but with unprecedented power to spread those views.

Robert Herritt is an American writer living in Paris.

WINTER 2016 ~ 79

This state of affairs has provoked no shortage of hand-wringing in the commentariat. For example, in a March 2015 article in *National Geographic*, Joel Achenbach lamented the supposed rise of science skepticism in American culture. “Empowered by their own sources of information and their own interpretations of research,” he writes somewhat dramatically, “doubters have declared war on the consensus of experts.” A few months later, Lee McIntyre of Boston University offered a similar analysis in the *Chronicle of Higher Education*. Explaining what he sees as a growing disrespect for truth in American culture, McIntyre points to the Internet as a likely culprit. After all, he argues, “outright lies can survive on the Internet. Worse, those who embrace willful ignorance are now much more likely to find an electronic home where their marginal views are embraced.”

Complaints of this kind are not without merit. Consider a recent survey from the Pew Research Center’s Initiative on Science and Society showing a significant gap between the views of laypeople and those of scientists (a sample from the American Association for the Advancement of Science) on a wide range of scientific issues. To take one notable example, 88 percent of the polled AAAS scientists believe genetically modified foods to be safe, compared to only 37 percent of the respondents from the general public.

The discussions surrounding this situation often focus on the same basic question: Why is there such a gap between those in the know and everybody else? Or, as Yale Law School’s Dan Kahan and his coauthors put it in a 2011 paper, “Why do members of the public disagree—sharply and persistently—about facts on which expert scientists largely agree?” Kahan and his colleagues have identified several cultural forces and cognitive tendencies that help explain the discrepancy between expert consensus and lay opinion. For instance, the authors write that “individuals systematically overestimate the degree of scientific support for positions they are culturally predisposed to accept.” On especially divisive issues, such as climate change or gun regulation or nuclear waste disposal, there is a strong correlation between people’s own cultural values and their perceptions of the consensus among scientists. Not surprisingly, studies such as this one are frequently discussed in articles like Achenbach’s.

But as worthwhile as such research may be, it has little to say about a closely related question: What ought we to believe? How should non-experts go about seeking reliable knowledge about complex matters? Absent a granular understanding of the theories underpinning a given area of knowledge, how should laypeople weigh rival claims, choose

between conflicting interpretations, and sort the dependable expert positions from the dubious or controversial ones? This is not a new question, of course, but it has become more urgent thanks to our glut of instant information, not to mention the proliferation of expert opinion.

The closest thing to an answer one hears is simply to trust the experts. And, indeed, when it comes to the charge of the electron or the oral-health benefits of fluoride, this response is hard to quarrel with. The wisdom of trusting experts is also a primary assumption behind the work of scholars like Kahan. But once we dispense with the easy cases, a reflexive trust in specialist judgment doesn't get us very far. On all manner of consequential questions an average citizen faces—including whether to support a hike in the minimum wage or a new health regulation—expert opinion is often conflicting, speculative, and difficult to decipher. What then? In so many cases, laypeople are left to choose for themselves which views to accept—precisely the kind of haphazard process that the critics of “willful ignorance” condemn and that leaves us subject to our own whims. The concern is that, if we doubt the experts, many people will draw on cherry-picked facts and self-serving anecdotes to furnish their own versions of reality.

This is certainly the case. But, in fixating on this danger, we neglect an important truth: it is simply not feasible to outsource to experts all of our epistemological work—nor would it be desirable. We frequently have no alternative but to choose for ourselves which beliefs to accept. The failure to come to grips with this fact has left us without the kinds of strategies and tools that would enable non-experts to make more effective use of the increasingly opaque theories that explain our world. We need, in other words, something more to appeal to once disagreements reach the “my-source-versus-your-source” phase.

Developing approaches that fit this description will require an examination of our everyday assumptions about knowledge—that is, about which beliefs are worth adopting and why. Not surprisingly, those assumptions have been significantly shaped by our era's information and communication technologies, and not always for the better.

Anonymous Sources

During an early scene in the 2010 film *The Social Network*, Harvard twins Tyler and Cameron Winklevoss—both portrayed by actor Armie Hammer—participate in an early-morning crew workout on the Charles River, their double scull well ahead of the rest of the team. Cameron asks

if there is “any way to make this a fair fight?” Tyler suggests that “you could row forward and I could row backward.” To this obviously absurd idea Cameron responds, tongue-in-cheek, “We’re genetically identical. Science says we’d stay in one place.”

Cameron’s last comment features a telling construction: “science says.” The phrase wouldn’t give most English-speakers pause, as expressions like this are now commonplace. A 2014 article in *Scientific American*, for instance, features the headline “What Science Tells Us about Why We Lie.” Another 2014 article, in the magazine *The Week*, promises to explain “What economics tells us about the trustworthiness of movie reviews.” A *National Journal* piece from 2015 relays to readers “What Science Says about ‘Sounding Presidential.’” One recent Wired.com item reveals that “Physics Says Tiny Ant-Man Should Be Running Weirder,” while a creatively punctuated Inc.com headline reads: “It’s Not Just ‘Star Wars:?’ Psychology Says There Really Is a Dark Side.”

These formulations are, of course, understood to be shorthand for more elaborate and rigorous discovery processes. And one should not make too much of the tricks headline writers deploy for the sake of brevity. But the implication of this particular turn of phrase is that methods of inquiry such as science and economics are akin to blind mechanisms for the delivery of good beliefs about the world. We can say we know something after reading it off from the list of truths that “science tells us.” Journalism operates in a similar way. When confirming a particular fact one needs a credible source—a government insider, say, or an eyewitness, depending on the story—a source that is sometimes hidden from the reader.

Of course, anyone with an Internet connection has likely accepted this model of knowledge-seeking to some extent. A high schooler looking to know the capital of Azerbaijan or the president of Fiji might end his or her investigation once consulting a reputable website such as the Central Intelligence Agency’s World Fact Book. Debates over movie trivia are settled by the Internet Movie Database. Information about a prospective client, meanwhile, can be retrieved and skimmed moments before an unexpected meeting thanks to social networks like LinkedIn.

In all of these examples, knowing a piece of information is a matter of obtaining it from the right source: peer-reviewed journals, certain beyond-reproach websites, the federal government, social networks. And once an individual has formed a belief in such a manner, it is assumed that his or her epistemological responsibilities have been discharged. Knowledge, for all practical purposes, has been achieved.

The technology sector seems more than willing to capitalize on this view of knowledge as a kind of acquiescence to pre-established facts. Google has compiled what it calls its Knowledge Vault, potentially the world's largest repository of facts extracted from the web, using a system that "computes calibrated probabilities of fact correctness," as a Google research team wrote in a 2014 paper. Another Google paper further outlines the method for how a web source's trustworthiness gets quantified: Google's algorithm calculates a "Knowledge-Based Trust" score, which could then be used to rank webpages by their level of veracity. We may soon be using the phrase "Google tells us" with the same confidence as "science tells us."

This view of epistemology is a version of what philosopher John Dewey criticized as the "spectator theory of knowledge." According to Dewey, this approach "ascribe[s] the ultimate test of knowledge to impressions passively received, forced upon us whether we will or no." And if, in a growing number of circumstances, forming beliefs is simply a matter of taking in pre-digested information, of "impressions passively received," it gets easier and easier to see "knowing" as something that happens to us. There are costs to this way of thinking.

Covert Operations

One consequence of this view of knowledge is that it has become largely unnecessary to consider how a given piece of information was discovered when determining its trustworthiness. The research, experiments, mathematical models, or—in the case of Google—algorithms that went into establishing a given fact are invisible. Ask scientists why their enterprise produces reliable knowledge and you will likely be told "the scientific method." And this is correct—more or less. But it is rare that one gets anything but a crude schematic of what this process entails. How is it, a reasonable person might ask, that a single method involving hypothesis, prediction, experimentation, and revision is applied to fields as disparate as theoretical physics, geology, and evolutionary biology—or, for that matter, social-scientific disciplines such as economics and sociology?

Even among practitioners this question is rarely asked in earnest. Science writer and former *Nature* editorial staffer Philip Ball has condemned "the simplistic view of the fictitious 'scientific method' that many scientists hold, in which they simply test their theories to destruction against the unrelenting candor of experiment. Needless to say, that's rarely how it really works."

Like the algorithms behind Google’s proposed “truth” rankings, the processes that go into establishing a given empirical finding are often out of view. All the lay reader gets is conclusions such as “the universe is fundamentally composed of vibrating strings of energy,” or “eye color is an inherited trait.” By failing to explain—or sometimes even to acknowledge—how, exactly, “the scientific method” generates reliable knowledge about the world in various domains, scientists and science communicators are asking laypeople to accept the supremacy of science on authority.

Far from bolstering the status of experts who engage in rigorous scientific inquiry, this way of thinking actually gives them short shrift. Science, broadly construed, is not a fact-generating machine. It is an activity carried out by people and requiring the very human capacities of reason, intuition, and creativity. Scientific explanations are not the inevitable result of a purely mechanical process called “the scientific method” but the product of imaginative attempts to make empirical data more intelligible and coherent, and to make accurate predictions. Put another way, science doesn’t tell us anything; scientists do.

Failure to recognize the processes involved in adding to our stores of knowledge creates a problem for those of us genuinely interested in getting our beliefs right, as it denies us relevant information for understanding why a given finding deserves our acceptance. If the results of a single, unreplicated neuroscience study are to be considered just as much an instance of good science as the rigorously tested Standard Model of particle physics, then we laypeople have little choice but to give them equal weight. But, as any scientist will tell you, not all findings deserve the same credibility; determining which ones merit attention requires at least a basic grasp of methodology.

To understand the potential costs of failing to engage at the level of method, consider the Innocence Project’s recent investigation of 268 criminal trials in which evidence from hair analysis had been used to convict defendants. In 257 of those cases, the organization found forensic testimony by FBI scientists to be flawed—a conclusion the FBI does not dispute. What is more, each inaccurate analysis overstated the strength of hair evidence in favor of the prosecution. Thirty-two defendants in those cases were eventually sentenced to death, of whom fourteen have either died in prison or have been executed. This is an extreme example of how straightforwardly deferring to expert opinion—without considering how those opinions were arrived at—is not only an inadequate truth-seeking strategy, but a potentially harmful one.

Reacting to the discoveries of forensic malpractice at the FBI, the co-chairman of the President's Council of Advisors on Science and Technology, biologist Eric S. Lander, suggested a single rule that would make such lapses far less common. As he wrote in the *New York Times*, "No expert should be permitted to testify without showing three things: a public database of patterns from many representative samples; precise and objective criteria for declaring matches; and peer-reviewed published studies that validate the methods."

Lander's suggestion amounts to the demand that forensic experts "show their work," so to speak, instead of handing down their conclusions from on high. And it is an institutional arrangement that could, with a few adjustments, be applied to other instances where expert analyses carry significant weight. It might be too optimistic to assume that such information will be widely used by the average person on the street. But, at least in theory, efforts to make the method by which certain facts are established more available and better understood will leave each of us more able to decide which claims to believe. And these sorts of procedural norms would help create the expectation that, when choosing what to believe, we laypeople have responsibilities extending beyond just trusting the most credentialed person in the room.

Conclusions on Demand

Contemporary computer and information technologies only strengthen the temptation to ignore the processes used to establish various facts about the world. A number of recent findings in psychology suggest that certain tools, particularly search engines, make it easy to mistake information obtained online with knowledge we have achieved and internalized for ourselves.

For instance, in a study published in June 2015 in the *Journal of Experimental Psychology*, researchers asked participants to report their self-assessed level of explanatory knowledge about various topics (for instance, "How do tornadoes form?"), after having searched the Internet for answers to *unrelated* questions ("How does a zipper work?"). The authors concluded that "searching for answers online leads to an illusion such that externally accessible information is conflated with knowledge 'in the head.'" Further, they suggested that "searching the Internet may cause a systematic failure to recognize the extent to which we rely on outsourced knowledge."

Earlier experiments conducted by Adrian F. Ward of the University of Texas for his doctoral dissertation found a similar tendency regarding

fact-based, as opposed to explanatory, knowledge. Ward suspects that because the Internet is so fast and unobtrusive, for instance when we use a “memory partner” like Google, people often get the false sense that they “know what they never knew,” while the means by which they received the information “quickly fades from awareness.”

One imagines that technologies that make information retrieval even less effortful will only facilitate these sorts of errors in self-assessment. Apple, for instance, announced in September 2015 that its latest Apple TV will include the voice-recognizing virtual assistant Siri. People watching a film through the device will need only say the words “Hey Siri, who directed this?” to have an answer spoken back to them by a disembodied voice. When any factual itch can be instantly scratched in this way, the illusion of knowing more than we do will be that much more powerful.

It is worth noting that fears about the intellectual dangers that accompany easy access to ready-made knowledge are not unique to today’s information technologies. In his 1851 essay “On Reading and Books,” the philosopher Arthur Schopenhauer expressed concerns about, of all things, print. “When we read,” he explains, “another person thinks for us: we merely repeat his mental process. In learning to write the pupil goes over with his pen what the teacher has outlined in pencil; so in reading, the greater part of the work of thought is already done for us.” Simply ingesting the conclusions of others is not what it means to know something. None of this is to suggest that empirical disciplines, whether in the natural or social sciences, do not deserve the authority they currently enjoy, nor that the argument from authority is not a satisfactory way of acquiring information in many circumstances. It is often unavoidable. As the economic historian Deirdre McCloskey has written, the appeal to authority “is a common and often legitimate argument.... No science would advance without it, because no scientist can redo every previous argument.” But for non-experts to accept such authority responsibly, they must first have an accurate understanding of why certain modes of inquiry are better than others.

Getting to Know Better

So far, I have only alluded to an alternative conception of knowing as an activity of sorts. As mentioned, this view can be found in the work of philosopher John Dewey. In his Gifford Lectures, published in 1929 as *The Quest for Certainty*, Dewey argues that “knowing is itself a kind of action, the only one which progressively and securely clothes natural existence with realized meanings.”

Related ideas can be found in the later work of Ludwig Wittgenstein. One well-known example from his posthumously published *Philosophical Investigations* involves the so-called duck-rabbit illusion—a drawing that can be viewed as an image of either animal, but not both simultaneously. It is clear that, when one stops seeing the picture as a rabbit and starts seeing it as a duck, something changes. “But what is different?,” Wittgenstein asks. He shows us how something as seemingly passive as visual experience is more than a mere imposition of the outside world on our senses; our experience is in part determined by how we respond to stimuli—how we act. And if our most direct sensory experiences are the product of our own actions, it is easy to see how the far more cognitive task of forming beliefs is as well. (It is worth noting that Dewey saw the mistaken “spectator theory of knowing” as being “modeled after what was supposed to take place in the act of vision.”)

A more recent example of this way of thinking can be found in the work of Oxford philosopher of information Luciano Floridi. As he puts it, “we do not and cannot gain knowledge by passively recording reality in declarative sentences, as if we were baskets ready to be filled; instead, we must handle it interactively.”

If knowing is a kind of activity, it follows that forming beliefs—in any domain—is something we can do with varying degrees of proficiency. It is, in a sense, a skill, not unlike oil painting or poker—an ability that we may be able to improve.

Research from psychologist Philip Tetlock and colleagues lends support to this idea. Tetlock is co-creator of The Good Judgment Project, an initiative that won a multi-year forecasting tournament conducted by the Intelligence Advanced Research Projects Activity, a U.S. government research agency. Beginning in 2011, participants in the competition were asked a range of specific questions regarding future geopolitical events, such as, “Will the United Nations General Assembly recognize a Palestinian state by Sept. 30, 2011?,” or “Before March 1, 2014, will North Korea conduct another successful nuclear detonation?” Tetlock’s forecasters, mind you, were not career analysts, but volunteers from various backgrounds. In fact, a pharmacist and a retired irrigation specialist were among the top performers—so-called “superforecasters.”

In analyzing the results of the tournament, researchers at the Good Judgment Project found a number of characteristics common to the best forecasters. For instance, these individuals “had more open-minded cognitive styles” and “spent more time deliberating and updating their forecasts.” In a January 2015 article in the *Washington Post*, two of the

researchers further explained that the best forecasters showed “the tendency to look for information that goes against one’s favored views,” and they “viewed forecasting not as an innate ability, but rather as a skill that required deliberate practice, sustained effort and constant monitoring of current affairs.”

What these findings suggest is that, when it comes to reaching conclusions on complex matters in situations where information is limited and imperfect, certain habits of mind can provide a significant advantage. What is more, thinking about this task as a skill that to some extent can be learned might actually encourage the development of the relevant mental capacities.

These discoveries fit nicely with a set of views in academic philosophy that go under the banner “virtue epistemology”—an approach with roots stretching back at least to Aristotle and reintroduced in the Anglo-American philosophical tradition by Ernest Sosa in his 1980 paper “The Raft and the Pyramid.” Central to many of these theories is the notion that good beliefs are those that exhibit virtues. Different versions of this approach characterize intellectual virtues in different ways. According to one camp, they might include traits such as intellectual courage, attentiveness, tenacity, carefulness, fairness in evaluating the ideas of others, and, wouldn’t you know it, open-mindedness.

These epistemic virtues, you will notice, are analogous to moral virtues. Just as morally sound actions, according to views like Aristotle’s, are instances of moral habits, such as courage and justice, true beliefs are grounded in certain epistemic tendencies.

Virtue epistemology is, of course, a controversial school of thought. But one need not embrace it fully in order to see how such ideas may prove useful to laypeople on the hunt for good beliefs. Faced with a difficult question about nutrition, public policy, or child-rearing, we might set ourselves the goal of developing certain habits of mind instead of simply trusting our instincts. This might mean remaining open to opposing points of view, deliberately and continually challenging our own first-blush assessments, and taking care to reevaluate our beliefs in light of new information. It would mean paying closer attention to the processes we use when sizing up the world around us.

Territorial Disputes

Familiarizing ourselves with processes that lead to knowledge might also clarify our sense of what a given mode of investigation can and cannot

reveal. In the realm of public policy, for instance, there has long been a temptation to cast thorny social issues as problems conducive to straightforward, quantitative solutions. Issues like drug prohibition, health policy, and early childhood education are, on this view, for the most part best left to economists and their models and analyses.

The debate over health policy, for instance, is often concerned with purely empirical issues, such as the cost of expanding Medicaid or the feasibility of a single-payer system. But many of the questions underlying these discussions are inescapably value-laden. Should universal insurance coverage be our chief aim at all costs, or do our political obligations extend only to maximizing access to affordable care? A good grasp of the methodologies of social science would reveal that such matters cannot be settled through only empirical means.

By presenting normative issues as something akin to technical puzzles, the basic matters of value bound up in them become less noticeable. As this process repeats itself over time, the space for genuine political conversation by average citizens—conversations that appeal to less quantifiable aspects of life such as morality and tradition—gets smaller and smaller. The policies that affect our everyday lives are seen as the business of scientists and few others.

The Cambridge philosopher, economist, and mathematician Frank P. Ramsey expressed this sort of expert elitism rather unabashedly in a 1925 lecture: “Science, history, and politics are not suited for discussion except by experts. Others are simply in the position of requiring more information; and, till they have acquired all available information, cannot do anything but accept on authority the opinions of those better qualified.”

It may be a tempting line, but it comes at too high a cost, especially the part about politics. If we are ever to play anything but a symbolic role in the political decisions that shape our lives, there must be a place for informed non-experts to contribute meaningfully to the discussion. Staking out this territory—this middle ground between expertise and ignorance—will take work. To begin with, it will require us to reject the predominant idea of truth as something that arrives fully formed on our front porch each morning, or that is piped into our laptops, our phones, our crania. The alternative view I have sketched is one in which we take an active part in acquiring knowledge from the world, are responsible for our own beliefs, and in which our goal is continuously to improve our skills at apprehending reality. Absent such an alternative, we are all just barstool debaters, querying our phones for rhetorical ammunition, pretending to know what we are talking about.