CH 8

Science: Collaboration, Competition, and Reputation

pp. 159–172

THE WISDOM OF CROWDS

WHY THE MANY ARE SMARTER THAN THE FEW AND HOW COLLECTIVE WISDOM SHAPES BUSINESS, ECONOMIES, SOCIETIES, AND NATIONS

JAMES SUROWIECKI

DOUBLEDAY
New York London Toronto Sydney Auckland

8.

SCIENCE: COLLABORATION, COMPETITION, AND REPUTATION

In early February of 2003, the Ministry of Health of the People's Republic of China notified the World Health Organization that since November of 2002, 305 people in Guangdong Province had been stricken with a severe respiratory disease, which had killed five of them. Although the disease's symptoms resembled the flu, laboratory tests had come back negative for influenza viruses. A couple of weeks after the WHO got this news, a man returning from a trip to China and Hong Kong fell ill with a severe respiratory disease in Hanoi and was hospitalized, even as a number of workers at a Hong Kong hospital came down with similar symptoms. Reports of new outbreaks continued to arrive, and by early March, it seemed clear that SARS—as the illness had been dubbed—was not a new kind of flu but an entirely new disease. In response, the WHO issued a global warning about SARS, cautioning travelers about journeying to southern Asia and activating a global surveillance system that was meant to alert the organization to any new outbreaks of the disease.

While tracking the disease was important—since it was already clear that SARS was transmitted from person to person, and that therefore quarantining might be an important strategy in fighting the disease—it was just as important to discover the cause of the disease, which would open the door for testing and, perhaps,

an eventual vaccine. And so even as it issued its global alert, the WHO set in motion a global effort to uncover the source of SARS. On March 15 and 16, the organization contacted eleven research laboratories from countries around the world—including France, Germany, the Netherlands, Japan, the United States, Hong Kong, Singapore, Canada, the United Kingdom, and China—and asked them to work together to find and analyze the SARS virus. All of them agreed, and on March 17 embarked on what the WHO called a "collaborative multicenter research project." Every day the labs took part in daily teleconferences, where they shared their work, discussed avenues for future investigation, and debated current results. On a WHO Web site, the labs posted electron-microscope photographs of viruses isolated from SARS victims (any one of which might have been the cause of the disease), virus analyses, and test results. The labs regularly traded virus samples, allowing them to both check on and learn from each other's work.

Because of the way the collaboration functioned, different labs were able to work at the same time on the same samples, multiplying their speed and effectiveness. In the first few days of the effort, the labs considered and then dismissed a host of possible causes of the disease, including a series of viruses that were found in samples from some SARS patients but not others. By March 21, scientists at Hong Kong University had already isolated a virus that seemed like a likely candidate. That same day, scientists at the Centers for Disease Control in the United States separately isolated a virus that, under the electron microscope, looked like what's called a coronavirus. This was something of a surprise. Coronaviruses make animals very sick, but in humans their effects tend to be rather mild. But over the next week, labs in the network detected the coronavirus in a wide variety of samples from people who had been diagnosed with SARS. Labs in Germany, the Netherlands, and Hong Kong began sequencing the virus. In early April, monkeys in the Netherlands laboratory who had been infected with the coronavirus came down with full-blown cases of

SARS. By April 16, a mere month after their collaboration had begun, the labs were confident enough to announce that the coronavirus did, in fact, cause SARS.

The discovery of the SARS virus was, by any measure, a remarkable feat. And when we're faced with a remarkable feat, our natural inclination is to ask: Who did it? Who actually discovered the cause of SARS? But the truth is, that's an impossible question to answer. We know the name of the person who first spotted the coronavirus. She was an electron microscopist named Cynthia Goldsmith, who worked in the Centers for Disease Control and Prevention lab in Atlanta. But you can't say she discovered what caused SARS, since it took weeks of work by labs all over the world to prove that the coronavirus actually made people sick. For that matter, all the work that proved that other viruses didn't cause SARS was instrumental as well, since it narrowed the field of possible candidates. Ultimately, no one person discovered the cause of SARS. Instead, as the WHO's own account of the search for the virus argues, it was the group of labs that "collectively . . . discovered" the coronavirus. Working on their own, any one of those labs might very well have taken months or years to isolate the virus. Together it took them just a matter of weeks.

The intriguing thing about the success of the laboratories' collaboration is that no one, strictly speaking, was in charge of it. Although the WHO orchestrated the creation of the network of labs, there was no one at the top dictating what different labs would do, what viruses or samples they would work on, or how information would be exchanged. The labs agreed that they would share all the relevant data they had, and they agreed to talk every morning, but other than that it was really up to them to make the collaboration work. The guiding assumption of the search for SARS was that on their own, the labs would figure out the most efficient way to divide up the work. Part of this, of course, was simple necessity: the WHO has no real authority to make academic or government laboratories do anything. But in this case, necessity became virtue. In the absence of top-down direction, the laboratories did a remarkably good job of organizing themselves. The collaborative nature of the project gave each lab the freedom to focus on what it believed to be the most promising lines of investigation, and to play to its particular analytical strengths, while also allowing the labs to reap the benefits—in real time—of each other's data and analyses. And the result was that this cobbled-together multinational alliance found an answer to its problem as quickly and efficiently as any top-down organization could have.

THE SCOPE AND SPEED of the SARS research effort made it unique. But in one sense the successful collaboration between the labs was simply an exemplary case of the way much modern science gets done. Although in the popular imagination science remains the province of the lone genius working alone in his lab, in fact it is, in more ways than one, a profoundly collective enterprise. Before World War I, collaboration was relatively rare for scientists. But that began to change in the decades before World War II, and in the postwar years teamwork and group projects proliferated rapidly. Researchers, particularly experimental researchers, routinely work in large groups, and it's no longer strange to see scientific papers that are co-authored by ten or twenty people. (This is in sharp contrast to the humanities, where single authorship remains the norm.) A classic example of this phenomenon was the discovery, in 1994, of the quantum particle called the "top quark." When the discovery was announced, it was credited to 450 different physicists.

Why do scientists collaborate? Part of it is a result of what's often called the "division of cognitive labor." As science has become ever more specialized and as the number of subfields within each discipline has proliferated, it's become difficult for a single person to know everything he needs to know. This is especially true in experimental science, where sophisticated machinery demands unique skills. Collaboration allows scientists to incorporate many different kinds of knowledge, and to do so in an active way (rather than simply learning the information from a book). Collaboration also makes it easier for scientists to work on interdisciplinary problems—which happen to be among today's most important and interesting scientific problems. Small groups do face tremendous challenges in solving problems and making decisions, and they can waste a great deal of time dividing up the labor, discussing results, and debating conclusions. But those potential costs are clearly, for most scientists, outweighed by the benefits.

Collaboration also works because, when it works well, it guarantees a diversity of perspectives. In the case of the search for the SARS virus, for instance, the fact that different labs had different initial ideas about the possible origin of the virus meant that a wide range of possibilities would be considered. And the fact that different laboratories were doing parallel work on the same samples, while it ran the risk of producing too much duplicated effort, also produced rich results in the form of unique data.

Ultimately, for a collaboration to be successful it has to make each individual scientist more productive. A wide array of studies have found that, more often than not, collaboration seems to do just that. Economist Paula Stephan has argued, "Scientists who collaborate with each other are more productive, often times producing 'better' science, than are individual investigators." And social scientist Etienne Wenger adds: "Today's complex problem solving requires multiple perspectives. The days of Leonardo da Vinci are over."

Saying that the days of Leonardo da Vinci are over, though, is not the same as saying that collaboration waters down or squelches individual creativity. In fact, one of the more intriguing aspects of scientific collaboration is that the more productive and better known a scientist is, the more frequently he or she works with others. This has been the case for decades. In a 1966 study of 592 scientists' publications and collaborative activities, for instance, D. J. de Solla Price and Donald B. Beaver found that "the most prolific man is also by far the most collaborating, and three of the four next most prolific are also among the next most frequently collaborating." A similar study by Harriet Zuckerman, which compared forty-one Nobel laureates with a sample of similarly placed scientists, found that the laureates collaborated more often than regular scientists. Of course, it's easier for well-known scientists to collaborate because everyone wants to work with them. But the fact that they are committed to working with others, when you might expect them to assume that they have nothing to gain from it, testifies to the centrality of cooperative efforts to modern science.

Still, the kind of global collaboration that we witnessed in the search for the SARS virus remains unusual. Although the scientific community clearly is global in nature, most collaboration takes place, even today, with people in a scientist's immediate vicinity. Barry Bozeman, for instance, found that academic researchers spend only a third of their time working with people who are not in their immediate work group, and only a quarter of their time working with people who are outside their university. That's not too surprising. For all the talk of the "death of distance," people still prefer to work in close physical proximity to their colleagues. But as the SARS example suggests, this may be changing. Technology is now making global collaboration not just possible but easy and productive. And the value of working across not only universities but nations is clearly immense, while limiting yourself to the skill set found in your immediate department or working group seems self-defeating. It's perhaps not surprising, then, that researchers who spend a lot of time working with researchers in other nations are significantly more productive than researchers who don't. Again, it's possible that the correlation here runs in the opposite direction: that it's easier for more productive—which generally means better known—scholars to collaborate internationally. But regardless of why it's true, what's telling is that it is.

EXPLICIT COLLABORATION ON ACADEMIC papers and research projects is not the only thing that makes science a collective enterprise. Science is collective because it depends on and has tried to institutionalize the free and open exchange of information. When scientists make an important new discovery or experimentally prove some hypothesis, they do not, in general, keep that information to themselves so that they alone can ponder its meaning and derive additional theories from it. Instead, they publish their results and make their data available for inspection. This makes it possible for other scientists to reconsider their data and possibly refute their conclusions. More important, though, it makes it possible for other scientists to use that data to construct new hypotheses and perform new experiments. The assumption is that society as a whole will end up knowing more if information is diffused as widely as possible, rather than being limited to a few people. In a strict sense, every scientist depends on the work of other scientists.

Newton pointed to something like this when he spoke of "standing on the shoulders of giants." But Newton, who did most of his theoretical work alone and who was obsessed with being sui generis, was suggesting only that his insights depended on the work of those who had come before him. He was making the point that scientific knowledge is, in some sense, cumulative. (Of course, Newton used the phrase in a letter to his rival Robert Hooke, who happened to be a dwarf, so it's possible that the phrase was intended only as a cruel joke.) But that knowledge is more than cumulative. It's collective. Scientists depend not just on the work of their predecessors, but also on the work of their contemporaries, who are in turn dependent on them. Even scientists whose hypotheses fail are helping their peers, by letting them know where they do not need to go.

Although the effect of the work of individual scientists is to accumulate scientific knowledge for the community as a whole, that's not really the point of scientific endeavor. Scientists want to solve particular problems. And they want to be recognized, to earn the attention of their contemporaries, to transform the way other scientists think. The coin of the realm, for most scientists, is not cash but rather recognition. Even so, scientists are undoubtedly as self-seeking and as self-interested as the rest of us. The genius of the way science is organized, though, makes their self-interested behavior redound to the benefit of all of us. In the process of winning notoriety for themselves, they make the group—that is, the scientific community and then, indirectly, the rest of us—smarter.

What's striking about the organization of modern science is that—like the SARS network of labs—no one is in charge. Obviously, there have been massive and important top-down research projects—think of the Manhattan Project or the Atlas missile project—in which scientists worked under explicit direction to solve particular problems, and these projects, most of them government sponsored, have often been successful. At the same time, since the late nineteenth century, a good deal of scientific work has taken place in corporate research labs, where there has often—though not always—been a more systematized, command-and-control approach to research. But in the history of science and technology, top-down organization has always been more of an anomaly than the ordinary way of doing business. For the most part, scientists (at least established ones) have been left to their own devices to choose what they were interested in, how they would work on it, and what they would do with their results.

That's not to say that the choices that scientists make are innocent. A scientist does not enter his lab as a blank slate, waiting to hear what the data will tell him. Instead he enters it as someone whose understanding of what problems are interesting, what problems can be solved, and what problems should be solved has been shaped by the interests (in both senses of the word) of his community. And since a hefty chunk of scientific research has been and is still today funded by the government, with grants handed out by peer review boards, the interests of a scientist's peers often have a direct and concrete impact on the kind of work he chooses to do. Even so, the important point is that there is no Science Czar telling researchers what they should do. We trust that allowing individuals to pursue their own self-interest will produce collectively better results than dictating orders.

Pursuing their own self-interest is more complicated for scientists than it might sound. While scientists are fundamentally competing for recognition and attention, that recognition and attention can only be afforded them by the very people they're competing against. So science presents us with the curious paradox of an enterprise that is simultaneously intensely competitive and intensely cooperative. The quest for recognition ensures a steady infusion of diverse thought, since no one becomes famous for restating what's already known. (This makes it less important that scientists tend to be interested in what other scientists are interested in, since the quest for originality forces researchers to think past convention.) And the competition also works to provide an inherent check on flawed ideas, since, as the philosopher David Hull has argued, showing the flaws in other people's work is one way to make a name for yourself. But all that competition depends on a given level of cooperation, because it's the rare scientist who can flourish in isolation from the work of his peers.

What allows this strange blend of collaboration and competition to flourish is the scientific ethos that demands open access to information. This ethos dates back to the origins of the scientific revolution in the seventeenth century. In 1665, the Royal Society—one of the first institutions, and certainly the most important, formed to foster the growth of scientific knowledge—published the first issue of its *Philosophical Transactions.* It was a seminal moment in the history of science, because of the journal's fierce commitment to the idea that all new discoveries should be disseminated as widely and freely as possible. Henry Oldenburg, the first secretary of the Royal Society and the editor of the *Transactions,* pioneered the idea that secrecy was inimical to scientific progress, and convinced scientists that they should give up their sole ownership of their ideas in exchange for the recognition they would receive as the creator or discoverer of those ideas. What Oldenburg grasped was the peculiar character of knowledge, which does not, unlike other commodities, get used up as it is consumed and which can be there-

fore spread widely without losing its value. If anything, in fact, the more a piece of knowledge becomes available, the more valuable it potentially becomes, because of the wider array of possible uses for it. As a result, the historian Joel Mokyr writes, the scientific revolution became the period "in which 'open science' emerged, when knowledge about the natural world became increasingly nonproprietary and scientific advances and discoveries were freely shared with the public at large. Thus scientific knowledge became a public good, communicated freely rather than confined to a secretive exclusive few as had been the custom in medieval Europe."

This tradition of open publication and communication of insights was, of course, central to the success of Western science. It's open science that made the self-interested behavior of scientists collectively beneficial. Scientists were willing to publish their insights because that was the route to public recognition and influence. If one wanted to think about this process in market terms—as some have tried to do—you could say that scientists were paid by other people's attention. As the sociologist of science Robert K. Merton famously put it, "In science, one's private property is established by giving its substance away."

The challenge the scientific community faces today is whether the success of Western science can survive the growing commercialization of scientific endeavors. Science and commerce have, of course, been intertwined for centuries. But as an increasing share of scientific research and development is funded by corporations, which see themselves as having an economic interest in protecting information rather than in disseminating it widely, the nature of scientific exchange may change. The sociologist Warren Hagstrom talked about science as a "gift economy" rather than an exchange economy. And the idea of science as made up of "invisible colleges" of researchers bound by their common interest in expanding knowledge, if perhaps naïve, still has a powerful hold not just on laypeople but on scientists themselves. Corporations, on the other hand, are generally not gift givers nor do they thrive on collegiality. The fact that public funding is still instrumental to science, and particularly to basic research, insulates scientists to some extent from commercial pressures. And although the patent system limits what others can do with a given invention, it also—by requiring the inventor to publish the details of his invention in order to get a patent—plays a role in continuing to fuel the free flow of information. But the conflict between science and business is not imaginary. The spectacle of companies funding studies and then demanding that they be suppressed when the results do not come back to their satisfaction is not something that would have pleased Henry Oldenburg.

TALKING ABOUT SCIENTIFIC ENDEAVOR in terms of the quest for recognition may make it sound as if scientists were simply fame hounds (which, of course, some of them are). But recognition is not, at least in theory, about celebrity or fashion. Recognition is instead the proper reward for genuinely new and interesting discoveries. Scientists want to be recognized because it's nice to be recognized. But they also want to be recognized because recognition is what allows new ideas to be incorporated into the general body of scientific knowledge. What's intriguing about science from the perspective of collective problem solving is that it is the community as a whole that bestows the recognition, which is to say that it's the community as a whole that decides whether or not a scientific hypothesis is true and whether it's original. This doesn't mean that scientific truth is in the eye of the beholder. The coronavirus caused SARS before the WHO announced that the coronavirus caused SARS. But in scientific terms, the coronavirus only became the cause of SARS once other scientists had scrutinized the work of the labs and accepted it as proving what they said it proved. Academic labs and corporate research labs across the world are now busy working on possible tests and vaccines for SARS, all predicated on the idea that the SARS virus is a coronavirus. They are doing so only because the scientific community has reached—in an indirect way—a consensus on the issue. As Robert K. Merton wrote, "There is no such thing as a scientific truth

believed by one person and disbelieved by the rest of the scientific community; an idea becomes a truth only when a vast majority of scientists accept it without question. That is, after all, what we mean by the expression 'scientific contribution': an offering that is accepted, however provisionally, into the common fund of knowledge."

This seems so obvious to us that it's easy to miss how much faith this places in the good judgment of the scientific community as a whole. Instead of relying on an elite group of scientists to pronounce on the validity of new ideas, scientists simply toss their ideas out into the world, trusting that the ones that survive are the ones that deserve to. The process is dramatically different from the way markets or democracies work. There are no literal votes taken, and ideas do not carry a price tag. But at the core of the process of accepting new ideas into the common fund of knowledge is a kind of unexpressed faith in the collective wisdom of scientists.

It's true, of course, that since scientific results should be replicable, you don't in theory have to trust anyone's judgment. If an experiment works, it will work whether or not the vast majority of scientists say it does or not. But the picture is more complicated than this. Most scientists are never going to replicate other experiments. They're going to trust that the data is correct and that the experiments worked as the scientist who performed them said they did. A successful hypothesis is a hypothesis that most scientists find credible, not a hypothesis that most scientists have tested for themselves and found to be true. In fact, once a theory has been accepted, simply failing to replicate the data on which it's based isn't enough. As the Hungarian scientist and philosopher Michael Polanyi argued, if you tried to reproduce a well-known experiment and failed, your initial response wouldn't be to doubt the experiment. You would doubt, and rightly so, your own lack of skill. This is best for science, since if researchers were constantly testing each other's results, they'd spend all their time retracing old ground instead of breaking new ground. And in any case, even to test another scientist's data requires you to rely on a host of other things that you almost certainly haven't tested yourself. Of an experiment in which he extracted DNA from an animal, the historian of science Steve Shapin wrote, "My extraction of DNA took on trust the identity of the animal tissue supplied, the speed of the centrifuge, the reliability of thermometric readings, the qualitative and quantitative makeup of various solvents, the rules of arithmetic."

Of course, experiments can be, and are, replicated. And scientific fraud is revealed. So the point is not that all truths are relative. Instead, the fact that what scientists know depends on the communications of others has two important consequences. First, good science requires a degree of trust among scientists that even as they compete, they will also cooperate by playing fair with their data. Second, and more important, science depends not only on an ever-replenishing pool of common knowledge, but also on an implicit faith in the collective wisdom of the scientific community to distinguish between those hypotheses that are trustworthy and those that are not.

UNFORTUNATELY, THERE IS SOMETHING of a flaw in this idealized picture of the way the scientific community discovers truth. And the flaw is that most scientific work never gets noticed. Study after study has shown that most scientific papers are read by almost no one, while a small number of papers are read by many people. Famous scientists find their work cited vastly more often than scientists who are less well known. When famous scientists collaborate with others they're given a disproportionate share of the credit for the work. And when two scientists—or two teams of scientists—independently make the same discovery, it's the famous scientists who end up getting the credit for that, as well. Merton dubbed this "the Matthew effect," after the Gospel lines "From unto every one that hath shall be given, and he shall have abundance: but from him that hath not shall be taken away even that which he hath." The rich get richer and the poor get poorer.

The Matthew effect can be seen in part as a kind of heuristic

device, a way for other scientists to filter the torrent of information that they are confronted with every day. And since there is a great deal of redundancy in scientific effort—that is, scientists often come up with the same hypotheses or run the same experiments—the Matthew effect does have the virtue of ensuring that some attention gets paid to work that otherwise might just disappear. Even so, the power of name recognition is startling. The geneticist Richard Lewontin, for instance, tells a story of publishing two papers, which he had co-authored with the biochemist John Hubby, back-to-back in the same issue of a scientific journal in 1966. The two papers, Lewontin writes, "were a genuinely collaborative effort in conception, execution, and writing and clearly form an indivisible pair." For the first paper, the biochemist Hubby's name was listed first. For the second, the geneticist Lewontin's name was listed first. There seemed to be no obvious reason why people should be more interested in one paper than in the other. Yet the paper that listed Lewontin's name first was cited 50 percent more than the other. The only answer, Lewontin suggested, was that he was at that point fairly well known as a geneticist while Hubby was still relatively unknown. When Lewontin's name came first, scientists assumed the paper was more of his work and that it was, therefore, more valuable.

The problem, of course, is that the reverence for the well known tends to be accompanied by a disdain for the not so well known. The physicist Luis Alvarez summed up this point of view decades ago when he said: "There is no democracy in physics. We can't say that some second-rate guy has as much right to opinion as Fermi." While this approach makes sense in terms of economizing on your attention—you can't listen to or read everyone, so you only listen to the best—it has a number of dubious assumptions built into it, including the idea that we automatically know who the second-rate are, even before hearing them, as well as the idea that everything Fermi had to say was inherently valuable. The obvious peril is that important work will be ignored because the person who produced it does not have the right brand name. Perhaps the classic example of this is Gregor Mendel, who found his work on heredity ignored, at least in part, because he was an unknown monk and who, as a result, simply stopped publishing his results.

The point is not that reputation should be irrelevant. A proven record of achievement does—and should—confer credibility on a person's ideas. The point instead is that reputation should not become the basis of a scientific hierarchy. The genius of the scientific ethos, at least in theory, is its resolute commitment to meritocracy. As Merton wrote in a famous essay on scientific norms, "The acceptance or rejection of claims entering the lists of science is not to depend on the personal or social attributes of the protagonist; his race, nationality, religion, class, and personal qualities are irrelevant." Ideas are meant to triumph not because of who is (or who is not) advocating them but because of their inherent value, because they seem to explain the data better than any of the others. This is perhaps just an illusion. But it's a valuable one.