



powells.com

REVIEW-A-DAY

The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations

by [James Surowiecki](#)

THE
NEWREPUBLIC
ONLINE

Thursday, June 24th,
2004

Mobbed up

A Review by Cass R. Sunstein

In the summer of 2003, analysts at the Department of Defense had an unusual idea. To predict important events in the world, including terrorist attacks, they would create a kind of market in which ordinary people could actually place bets. The proposed Policy Analysis Market would allow each of us to invest in our predictions about such matters as the growth of the Egyptian economy, the death of Yasir Arafat, and the likelihood of terrorist attacks in the United States. Investors would win or lose money on the basis of the accuracy of their predictions. Predictably, the Policy Analysis Market produced a storm of criticism. Ridiculed as "offensive" and "useless," the proposal was abandoned.

Amid the war on terrorism, why was the Defense Department so interested in the Policy Analysis Market? The answer is simple: it wanted to have some help in predicting geopolitical events, including those that would endanger American interests, and it believed that a market would provide that help. It speculated that if a large number of people could be given an incentive to aggregate their private information, in the way that the Policy Analysis Market would do, government officials would learn a great deal.

Does this idea seem ludicrous? Since 1988, the University of Iowa has run the Iowa Electronic Markets, which allow people to bet on the outcome of presidential elections. As a predictor, the Iowa Electronic Markets have produced extraordinarily accurate judgments, often doing better than professional polling organizations. In the week before each of the last four elections, the predictions in the Iowa market have shown an average absolute error of just 1.5 percentage points, a significant improvement over the 2.1 percentage point error in the final Gallup Polls. Or consider the Hollywood Stock Exchange, in which people predict Oscar nominees and winners, as well as opening weekend box-office successes. Here, too, the level of accuracy has been exceptionally impressive, with (for example) correct predictions of thirty-five out of forty Oscar nominees in 2002.

In fact, prediction markets are springing up all over the Internet, allowing people to make bets on the likely outcomes of sports, entertainment, finance, and political events. On tradesports.com, people have been betting on whether Donald Rumsfeld will resign soon (extremely unlikely), whether Osama bin Laden will be captured by June 2004 (extremely unlikely), whether John Edwards will be selected as John Kerry's running mate (a good chance, but probably not), and whether George W. Bush will be re-elected (more likely than not). One can imagine prediction markets on any number of questions: Will gas prices reach \$3

per gallon? Will cellular life be found on Mars? Will smallpox return to the United States? Will there be a sequel to *Master and Commander*? Will the Federal Communications Commission be abolished? (I didn't make these up; they are actual or proposed questions on existing markets.)

James Surowiecki is fascinated by prediction markets. In his opinion, they demonstrate that crowds are often wise. He rejects the widespread view that groups of ordinary people are usually wrong — and that we do better to ignore them and follow experts instead. Even when individuals blunder, he believes, groups can excel: "Under the right circumstances, groups are remarkably intelligent, and are often smarter than the smartest people in them." This is so even when "most of the people within the group are not especially well-informed or rational." What is wonderful, and surprising, is that "when our imperfect judgments are aggregated in the right way, our collective intelligence is often excellent." Instead of chasing experts, we should consult that collective intelligence.

As an example, Surowiecki points to the television game show *Who Wants to Be a Millionaire?*, on which contestants, if stumped, are permitted either to consult the studio audience or to place a call to a trusted friend or family member (selected in advance precisely because of his or her knowledge and intelligence). As it happens, the trusted allies perform well, producing correct answers 65 percent of the time. But the studio audience performs much better, picking right answers a remarkable 91 percent of the time. Surowiecki also invokes an astonishing finding by the British scientist Francis Galton, who tried to draw lessons about collective intelligence by examining a competition in which contestants guessed the weight of a fat ox at a regional fair in England. The ox weighed 1,198 pounds; the average guess, from the 787 contestants, was 1,197 pounds. Or consider Google, the astonishingly successful Internet search engine. Why does Google work so well? Surowiecki contends that its technology "is built on the wisdom of crowds." The company's founders, Sergey Brin and Lawrence Page, explain that the system "capitalizes on the uniquely democratic character of the Web." Google is good at telling you which site you are likely to want for one reason: it uses the collective "votes" of many other people.

Surowiecki is concerned with how crowds can solve three kinds of problems. The first are cognitive. These are factual questions with definite solutions, identifiable now or in the future. Who will win the World Series? How far from the Sun is the Earth? Will a certain surgery be successful? A second set of problems involve coordination. Individuals often need to select a shared course of action — driving on the same side of the road or meeting at a certain place. A third set involve cooperation. If people follow their self-interest, they might fail to cooperate with one another, and hence they will lose their opportunities for mutual advantage. Surowiecki contends that groups of people show far more cooperation than we might predict.

Surowiecki does not make the implausible suggestion that all crowds are wise. To qualify as such, a crowd needs to satisfy three conditions. It must be diverse; its members must be independent; and it must have a "particular kind of decentralization." Each of these conditions is designed to ensure what most interests Surowiecki, which is the emergence and the aggregation of information that group members have. Diversity is important simply to ensure that the group has a lot of information. If a crowd consists of nearly identical people, it is unlikely to be wise, because the group will not know more than the individuals of whom it is composed. Independence is necessary to ensure that people say what they know rather than hide it. Surowiecki is alert to the fact that groups often go wrong if members simply follow one another without pooling individually held information. Hence he notes, correctly, that organizations often do best if each individual behaves independently and does not pay a great deal of attention to the acts and the statements of others. "The smartest groups," he writes, "are made up of people with diverse perspectives who are able to stay independent of each other." The worst-performing investment clubs in the United States consist of people who like one another, socialize together, and show a great deal of consensus. The best performers

consist of people who do not see each other much and welcome dissent.

In calling for independence, Surowiecki emphasizes the serious risks associated with "information cascades," which occur when people neglect what they know and pay attention instead to the signals given by others. (In social science, such cascades have been found to arise not only among ordinary people choosing restaurants, sneakers, and political candidates, but also among doctors making diagnoses and even federal judges deciding cases.) The problem with information cascades is that group members are likely to do far worse than they would if everyone disclosed his or her private information. By pointing to the dangers of bad cascades, Surowiecki signals the importance of starting with a "wide array of options and information" and of having at least a few people who are willing "to put their own judgment ahead of the group's, even when it's not sensible to do so." Much of the time, Surowiecki writes, groups do best if their members pay little "attention to what everyone else is saying."

What about decentralization? Of Surowiecki's three conditions, this is the least intuitive. He attempts to clarify it by focusing on the war against terrorism. To wage that war successfully, of course, a great deal of information must be assembled. Surowiecki is critical of the widespread idea that what is needed is more centralization. Good solutions are far more likely to follow, he argues, "if you set a crowd of self-interested, independent people to work in a decentralized way on the same problem." Surowiecki seeks processes in which independent people, all armed with their own knowledge, are able to attend to problems "while also being able to aggregate that local knowledge and private information into a collective whole." The Iraq war is Surowiecki's example. Local American commanders had considerable latitude to act on their own, but they were also able to communicate rapidly, thus allowing successful overall strategies to develop from a multitude of local judgments. Surowiecki concludes that successful wars "may depend as much on the fast aggregation of information from the field as on preexisting, top-down strategies." (The problems that have arisen since the end of formal hostilities raise obvious difficulties for Surowiecki's claims; perhaps information on the ground is not being properly aggregated, or perhaps American officials don't have enough information on the ground to stop continuing attacks.) For intelligence relating to terrorism, Surowiecki argues that what is needed is aggregation, not centralization. And here Surowiecki returns to the ill-fated and roundly condemned Policy Analysis Market, suggesting that it "was potentially a very good idea."

Surowiecki is also fascinated by the very different phenomenon of social coordination. He points to the behavior of pedestrians on streets and sidewalks, where individuals are able to coordinate their movement so as not to bump into one another. Surowiecki pays tribute to "the beauty of a well-coordinated crowd, in which lots of small, subtle adjustments in pace and stride and direction add up to a relatively smooth and efficient flow," as people "are constantly anticipating each other's behavior." He thinks that pedestrian behavior helps to explain a great deal about the human ability to understand and to follow norms or conventions that other people follow at the same time.

Consider a little experiment by Thomas Schelling, who put the following puzzle to law students at Yale in 1958: You are going to meet someone in New York City. You do not know when or where, and you are unable to talk to the other person ahead of time. What time and place do you choose? Almost all the students said that they would meet at noon, and more than half said that they would meet at the information booth at Grand Central Station. As a more practical example, consider the universally accepted rule of first-come, first-served seating in buses, subways, and movie theaters. In Surowiecki's account, people are extremely good at generating conventions by which they organize their relationships.

Solutions to coordination problems are stable; once we hit upon a shared approach, we are likely to stick to it. Unfortunately, social cooperation is much more fragile, simply because each cooperator has an incentive

to defect. Suppose that everyone in a certain community thinks that the community will be better off if people engage in a recycling program. Even if everyone agrees, some people will refuse to participate, thinking that for them as individuals the costs of recycling exceed the benefits, even if the reverse is true for the group as a whole. Self-interested human beings try to "free ride" on the cooperation of others. And positing that people are self-interested, many economists expect cooperation to be rare. What interests Surowiecki is that cooperation is not rare at all. He emphasizes the enormous importance of reciprocity to human endeavors. Usually people will participate in a cooperative endeavor as long as they believe that other people are doing so too. Borrowing a claim by the political scientist Margaret Levi, Surowiecki concludes that people are "contingent cooperators." Most people don't want to be selfish jerks, but they also don't want to be dupes or fools. They will contribute to the common good if they believe that this is the general practice.

Surowiecki uses these points to explore a wide range of social phenomena, including scientific collaboration, stock prices, and corporate performance. One of his most interesting discussions involves the *Columbia* disaster and less-than-wise group deliberations at NASA. In Surowiecki's account, NASA emphasized consensus over dissent, and so it failed to take advantage of the information held by its engineers, who were perfectly aware of the underlying uncertainties. Stressing "the utter absence of debate and minority findings" in pre-launch discussions about the *Columbia*, Surowiecki argues for the need to counteract the risks associated with "group polarization." When group polarization occurs, people engaged in deliberation with one another end up thinking a more extreme version of what they thought before they started to talk. For example, those who believe that global warming is a serious problem are likely, as a result of internal discussions, to come to believe that global warming is an extremely serious problem; people who think that the Department of Justice is compromising civil liberties are likely to think, after they talk with one another, that the Department of Justice has no respect for civil liberties at all. So too officials at NASA, thinking that space shuttles are essentially safe, might well end up believing that safety is not a problem — even if several of them have private information suggesting otherwise.

Surowiecki knows that the phenomenon of group polarization raises problems for his thesis. If group members predictably end up thinking a more extreme version of what they thought before, what makes them likely to be wise? His answer is that groups need to contain safeguards to ensure that individual judgments are genuinely independent. NASA would have done far better if it had promoted a diversity of opinions and asked people to say what they really thought, rather than allowing internal pressures to lead people to squelch their doubts. The lesson here extends to many private and public institutions.

Surowiecki is aware that his celebration of wise crowds has implications for democracy. In light of his general argument, Surowiecki is suspicious of rule by a "technocratic elite," insisting that insulated officials lack the information to produce good decisions. But he does not think that democracies are really solving cognition problems; that's not their business. The reason is that unlike in cases involving simple facts, we "have no standard that allows us to judge a political decision to be 'right' or 'wrong.'" For all the public talk about the "common good," that idea is too disputed to provide objective solutions to political disputes. Surowiecki concludes that democracy should be seen as a way not to produce correct answers to particular questions, but to deal with "the most fundamental problems of cooperation and coordination: How do we live together? How can living together work to our mutual benefit?" On that count, democracy has crucial advantages.

The performance of groups is a wonderful subject, and Surowiecki has a remarkable eye for the telling anecdote, illustrating abstract claims with vivid examples. His central point is convincing. Groups, and even crowds, can be wiser than most and sometimes even all of their members, at least if they aggregate information. But there is a serious problem with Surowiecki's discussion: he does not provide an adequate

account of the circumstances that make crowds wise or stupid. Note first that the "conditions" that he identifies (diversity, independence, and decentralization) are neither necessary nor sufficient for the wisdom of crowds. On his own analysis, those are the conditions for the solution of problems of cognition, not problems of coordination or cooperation. People do not have to be diverse, or independent, to choose Grand Central Station as a meeting place in New York. If we want people to coordinate or to cooperate, it might well be best if they are similar and if they follow one another. To solve Surowiecki's three kinds of problems, quite different conditions come into play. In any case, coordination and cooperation problems don't come in neat boxes; life turns up all sorts of mixtures (consider marriage) that Surowiecki neglects.

Even for cognition problems, some groups sometimes perform best if their members are not independent and if they listen closely to one another. Groups can benefit when error-prone people silence themselves and follow the views expressed by their most sensible members. If the group contains authorities on the question at hand, members ought to listen closely — and possibly to shut up. Diversity is usually good, above all because it allows groups to acquire more information. But what is needed is not diversity as such, but diversity of the right kind. NASA's judgment would not have been improved if the relevant officials had included members of the Flat Earth Society, or people who believed that aliens are among us or that space flight is simply impossible.

All of these points suggest that the key question is how much information is held by various group members. Most generally, groups are wise only if their members actually know something about the relevant questions. Suppose that the studio audience in *Who Wants to Be a Millionaire?* were asked not about popular culture but instead about the number of decisions made by the Supreme Court every year. Is there any reason to expect that the majority or even the plurality would be right? Galton's crowd was good at judging the weight of a fat ox. But if its members were asked about the number of atoms in that ox, the median guess wouldn't be very reliable. (To have a reliable average response, the answers have to be better than random, and there cannot be a systematic bias in one or another direction.) Or imagine that a group of law professors is making decisions about how to build a space shuttle. They are unlikely to decide well, simply because law professors tend to know nothing about space shuttles. (I undertook a little experiment, asking law professors to guess the weight of the fuel used on a space shuttle; the right answer is 4.6 million pounds, and I won't embarrass my colleagues by announcing their answers, except to say that the average was way off.)

Surowiecki thinks that the "simplest way to get reliably good answers is just to ask the group each time." Judging the numbers of beans in a jar, groups almost always outperform most of their individual members. (Try it and you'll see.) Asking two hundred students to rank items by weight, one experimenter found that the group's estimate was 94 percent accurate — a figure excelled by only five individuals in that group. But it doesn't follow that groups will always, or generally, produce good answers. Everything depends on what the relevant people know. If you ask a group of randomly selected people about how to perform heart surgery, you will probably do better than if you asked a randomly selected individual; but you would do better still if you asked someone who actually knew how to perform heart surgery. Surowiecki loads the dice by pointing to areas in which good answers come from properly aggregating information that is held by many. In many areas, it is far more sensible to consult specialists.

The uses and the limits of Surowiecki's argument are helpfully approached via the Condorcet Jury Theorem, a significant omission from Surowiecki's presentation. Suppose that people are answering a common question with two possible answers, one false and one true, and that the average probability that each voter will answer correctly exceeds 50 percent. The Condorcet Jury Theorem holds that if each member of the group is answering independently, the probability of a correct answer, by a majority of the group, increases toward certainty as the size of the group increases. The theorem is based on some simple

arithmetic, the details of which are irrelevant here. Its importance lies in the demonstration that groups are likely to do better than individuals, and large groups better than small ones, if majority rule is used and if each person is more likely than not to be correct. The crucial proviso is the last one. If each person is more likely than not to err, then the theorem's prediction is reversed: the probability of a correct answer, by a majority of the group, decreases toward zero as the size of the group increases! It follows that groups are error-prone if most of their members are likely to blunder.

Surowiecki might object that some crowds can be wise even when ignorance is widespread. Consider the astonishing accuracy of the Iowa Electronic Markets (and other prediction markets), in which good judgments come from groups of investors that include many people who know little and are perhaps more likely to be wrong than to be right. But we cannot easily generalize from prediction markets, because they have several distinctive features. Most important, they do not simply rely on the median or average judgment of a randomly selected group of people. They are genuine markets, in which people voluntarily choose to participate, presumably because they think they know something. In addition, people are permitted to buy and to sell shares on a continuing basis. In these circumstances, accurate answers can emerge even if only a small percentage of participants have good information.

In the Iowa Electronic Markets, it turns out that 85 percent of the traders aren't so smart. They hold onto their shares for a long period and then just accept someone else's prices. The market's predictions appear to be driven by the other 15 percent — frequent traders who post their offers rather than accepting those made by other people. The broader point is that to work well, prediction markets do not require accurate judgments by anything like the majority of participants. In this sense, prediction markets are very different from judgments by ordinary crowds. Surowiecki's claims about group wisdom don't adequately emphasize the unique characteristics of these markets.

Surowiecki might have been expected to celebrate the cognitive virtues of democratic judgments — to suggest that a system that allows a voice for heterogeneous people and that encourages dissent is likely to come to sensible decisions, simply because it heeds the wisdom of the crowd. If 51 percent of voters support George W. Bush, maybe we have reason to think that they are right. But Surowiecki does not make this argument. As I have said, he insists that in the democratic domain we lack standards permitting us to distinguish decisions that are right from those that are wrong. But political decisions depend crucially on predictions. Will large deficits significantly increase interest rates? Will pre-emptive wars increase or decrease the threat of terrorist attacks? Will tax cuts spur economic growth? Problems of cognition are absolutely central to democratic governance. Of course democratic judgments often involve disputed judgments of value, for which demonstrably objective evidence is hard to find. Are pre-emptive wars just? Is economic growth more important than generous social safety nets? Still, in many matters government's performance is improved, or undermined, because of how it deals with cognition problems, and Surowiecki does not acknowledge this.

Return here to NASA, whose failures have been partly a product of a culture that disfavors dissent. In fact, group polarization is a pervasive problem in government circles, where like-minded officials often end up holding a more extreme version of the view with which they began. Surowiecki offers the example of the Bay of Pigs disaster, in which President Kennedy's advisers squelched their private doubts and developed unjustified enthusiasm for a ludicrous invasion plan predicated on the absurd thought that twelve hundred people could unseat Castro and take over Cuba. Is it too speculative to suggest that the current problems in Iraq are partly a product of group polarization within the executive branch — and that those problems could have been anticipated if the White House had had a better process for aggregating privately held information?

Franklin Delano Roosevelt's famously disorganized and much-criticized White House, with confusing lines of authority and multiple people working on similar tasks, was ideally suited to the production of a wide range of views and information. In this light, Surowiecki's dismissal of the idea that sometimes democracy faces cognition problems prevents him from exploring, or even seeing, some possible lessons for how to structure democratic institutions. In war and in peace, such institutions could take much more aggressive steps to elicit and to use existing information, above all by creating mechanisms to aggregate what people know.

Copyright © 1994-2003 Powells.com · [Terms of Use](#)