

THE WISDOM OF CROWDS

JAMES SUROWIECKI

Excerpt

- [Home](#) 
- [Meet the Author](#) 
- [About the Author](#) 
- [Q & A](#) 
- [Excerpt](#) 
- [Audio](#) 
- [Press Room](#) 
- [Buy the Book](#) 


I

If, years hence, people remember anything about the TV game show "Who Wants to Be a Millionaire?," they will probably remember the contestants' panicked phone calls to friends and relatives. Or they may have a faint memory of that short-lived moment when Regis Philbin became a fashion icon for his willingness to wear a dark blue tie with a dark blue shirt. What people probably won't remember is that every week "Who Wants to Be a Millionaire?" pitted group intelligence against individual intelligence, and that every week, group intelligence won.

"Who Wants to Be a Millionaire?" was a simple show in terms of structure: a contestant was asked multiple-choice questions, which got successively more difficult, and if she answered fifteen questions in a row correctly, she walked away with \$1 million. The show's gimmick was that if a contestant got stumped by a question, she could pursue three avenues of assistance. First, she could have two of the four multiple-choice answers removed (so she'd have at least a fifty-fifty shot at the right response). Second, she could place a call to a friend or relative, a person whom, before the show, she had singled out as one of the smartest people she knew, and ask him or her for the answer. And third, she could poll the studio audience, which would immediately cast its votes by computer. Everything we think we know about intelligence suggests that the smart individual would offer the most help. And, in fact, the "experts" did okay, offering the right answer—under pressure—almost 65 percent of the time. But they paled in comparison to the audiences. Those random crowds of people with nothing better to do on a weekday afternoon than sit in a TV studio picked the right answer 91 percent of the time.

Now, the results of "Who Wants to Be a Millionaire?" would never stand up to scientific scrutiny. We don't know how smart the experts were, so we don't know how impressive outperforming them was. And since the experts and the audiences didn't always answer the same questions, it's possible, though not likely, that the audiences were asked easier questions. Even so, it's hard to resist the thought that the success of the Millionaire audience was a modern example of the same phenomenon that Francis Galton caught a glimpse of a century ago.

As it happens, the possibilities of group intelligence, at least when it came to judging questions of fact, were demonstrated by a host of experiments conducted by American sociologists and psychologists between 1920 and the mid-1950s, the heyday of research into group dynamics. Although in general, as we'll see, the bigger the crowd the better, the groups in most of these early experiments—which for some reason remained relatively unknown outside of academia—were relatively small. Yet they nonetheless performed very well. The Columbia sociologist Hazel Knight kicked things off with a series of studies in the early 1920s, the first of which had the virtue of simplicity. In that study Knight asked the students in her class to estimate the room's temperature, and then took a simple average of the estimates. The group guessed 72.4 degrees, while the actual temperature was 72 degrees. This was not, to be sure, the most auspicious beginning, since classroom temperatures are so stable that it's hard to imagine a class's estimate being too far off base. But in the years that followed, far more convincing evidence emerged, as students and soldiers across America were subjected to a barrage of puzzles, intelligence tests, and word games. The sociologist Kate H. Gordon asked two hundred students to rank items by weight, and found that the group's "estimate" was 94 percent accurate, which was better than all but five of the individual guesses. In another experiment students were asked to look at ten piles of buckshot—each a slightly different size than the rest—that had been glued



The Wisdom of Crowds
0-385-72170-6
\$14.00 (Can. \$21.00)

to a piece of white cardboard, and rank them by size. This time, the group's guess was 94.5 percent accurate. A classic demonstration of group intelligence is the jelly-beans-in-the-jar experiment, in which invariably the group's estimate is superior to the vast majority of the individual guesses. When finance professor Jack Treynor ran the experiment in his class with a jar that held 850 beans, the group estimate was 871. Only one of the fifty-six people in the class made a better guess.

There are two lessons to draw from these experiments. First, in most of them the members of the group were not talking to each other or working on a problem together. They were making individual guesses, which were aggregated and then averaged. This is exactly what Galton did, and it is likely to produce excellent results. (In a later chapter, we'll see how having members interact changes things, sometimes for the better, sometimes for the worse.) Second, the group's guess will not be better than that of every single person in the group each time. In many (perhaps most) cases, there will be a few people who do better than the group. This is, in some sense, a good thing, since especially in situations where there is an incentive for doing well (like, say, the stock market) it gives people reason to keep participating. But there is no evidence in these studies that certain people consistently outperform the group. In other words, if you run ten different jelly-bean-counting experiments, it's likely that each time one or two students will outperform the group. But they will not be the same students each time. Over the ten experiments, the group's performance will almost certainly be the best possible. The simplest way to get reliably good answers is just to ask the group each time.

A similarly blunt approach also seems to work when wrestling with other kinds of problems. The theoretical physicist Norman L. Johnson has demonstrated this using computer simulations of individual "agents" making their way through a maze. Johnson, who does his work at the Los Alamos National Laboratory, was interested in understanding how groups might be able to solve problems that individuals on their own found difficult. So he built a maze—one that could be navigated via many different paths, some shorter, and some longer—and sent a group of agents into the maze one by one. The first time through, they just wandered around, the way you would if you were looking for a particular café in a city where you'd never been before. Whenever they came to a turning point—what Johnson called a "node"—they would randomly choose to go right or left. Therefore some people found their way, by chance, to the exit quickly, others more slowly. Then Johnson sent the agents back into the maze, but this time he allowed them to use the information they'd learned on their first trip, as if they'd dropped bread crumbs behind them the first time around. Johnson wanted to know how well his agents would use their new information. Predictably enough, they used it well, and were much smarter the second time through. The average agent took 34.3 steps to find the exit the first time, and just 12.8 steps to find it the second.

The key to the experiment, though, was this: Johnson took the results of all the trips through the maze and used them to calculate what he called the group's "collective solution." He figured out what a majority of the group did at each node of the maze, and then plotted a path through the maze based on the majority's decisions. (If more people turned left than right at a given node, that was the direction he assumed the group took. Tie votes were broken randomly.) The group's path was just nine steps long, which was not only shorter than the path of the average individual (12.8 steps), but as short as the path that even the smartest individual had been able to come up with. It was also as good an answer as you could find. There was no way to get through the maze in fewer than nine steps, so the group had discovered the optimal solution. The obvious question that follows, though, is: The judgment of crowds may be good in laboratory settings and classrooms, but what happens in the real world?

Excerpted from **The Wisdom of Crowds** by James Surowiecki Copyright © 2004 by James Surowiecki. Excerpted by permission of Anchor, a division of Random House, Inc. All rights reserved. No part of this excerpt may be reproduced or reprinted without permission in writing from the publisher.

[Home](#) | [About the Author](#) | [Meet the Author](#) | [Author Q&A](#) | [Excerpt](#) | [Audio](#)
| [Press Room](#) | [Buy the Book](#)

Copyright © 2004, Random House, Inc. | [Privacy Policy](#)

