A *NEW YORK TIMES* BUSINESS BESTSELLER

"As entertaining and thought-provoking as *The Tipping Point* by Malcolm Gladwell. . . . *The Wisdom of Crowds* ranges far and wide."
—*The Boston Globe*

THE WISDOM OF CROWDS

JAMES SUROWIECKI

WITH A NEW AFTERWORD BY THE AUTHOR

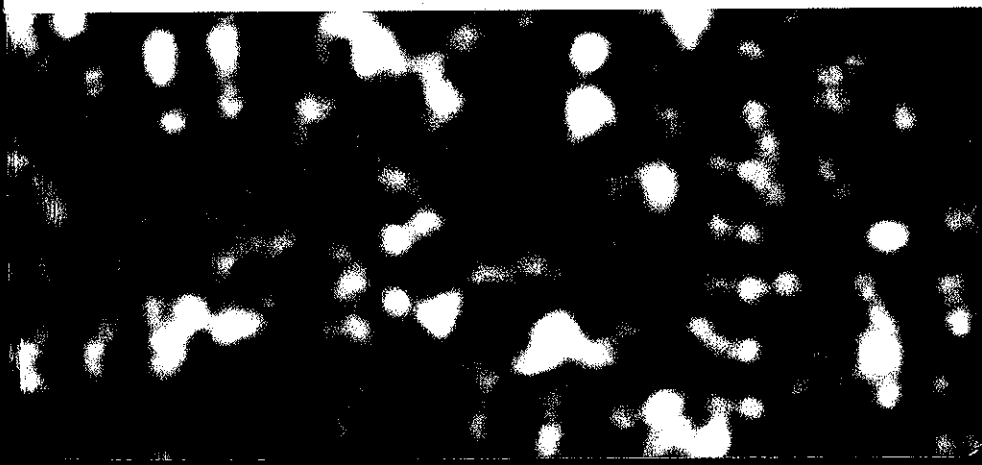

FIRST ANCHOR BOOKS EDITION, AUGUST 2005

Published in the United States by Anchor Books, a division of Random House, Inc., New York, and in Canada by Random House of Canada Limited, Toronto.
Originally published in hardcover in the United States in slightly different form by Doubleday, a division of Random House, Inc., New York, in 2004.

Anchor Books and colophon are registered trademarks of Random House, Inc.

Some of the material in this book was originally published in different form in *The New Yorker*.

The Library of Congress has cataloged the Doubleday edition as follows:
Surowiecki, James, 1967–
The wisdom of crowds : why the many are smarter than the few and how collective wisdom shapes business, economies, societies, and nations / James Surowiecki.
p. cm.
Includes bibliographical references.
1. Consensus (Social sciences) 2. Common good. I. Title.
JC328.2.S87 2003
303.3'8—dc22
2003070095

Anchor ISBN: 0-385-72170-6

www.anchorbooks.com

Printed in the United States of America
20 19 18 17 16 15 14 13 12 11

To Mom and Dad

INTRODUCTION

I

One day in the fall of 1906, the British scientist Francis Galton left his home in the town of Plymouth and headed for a country fair. Galton was eighty-five years old and beginning to feel his age, but he was still brimming with the curiosity that had won him renown—and notoriety—for his work on statistics and the science of heredity. And on that particular day, what Galton was curious about was livestock.

Galton's destination was the annual West of England Fat Stock and Poultry Exhibition, a regional fair where the local farmers and townspeople gathered to appraise the quality of each other's cattle, sheep, chickens, horses, and pigs. Wandering through rows of stalls examining workhorses and prize hogs may seem to have been a strange way for a scientist (especially an elderly one) to spend an afternoon, but there was a certain logic to it. Galton was a man obsessed with two things: the measurement of physical and mental qualities, and breeding. And what, after all, is a livestock show but a big showcase for the effects of good and bad breeding?

Breeding mattered to Galton because he believed that only a very few people had the characteristics necessary to keep societies healthy. He had devoted much of his career to measuring those characteristics, in fact, in order to prove that the vast majority of

people did not have them. At the International Exhibition of 1884 in London, for instance, he set up an "Anthropometric Laboratory," where he used devices of his own making to test exhibition-goers on, among other things, their "Keenness of Sight and of Hearing, Colour Sense, Judgment of Eye, [and] Reaction Time." His experiments left him with little faith in the intelligence of the average person, "the stupidity and wrong-headedness of many men and women being so great as to be scarcely credible." Only if power and control stayed in the hands of the select, well-bred few, Galton believed, could a society remain healthy and strong.

As he walked through the exhibition that day, Galton came across a weight-judging competition. A fat ox had been selected and placed on display, and members of a gathering crowd were lining up to place wagers on the weight of the ox. (Or rather, they were placing wagers on what the weight of the ox would be *after* it had been "slaughtered and dressed.") For sixpence, you could buy a stamped and numbered ticket, where you filled in your name, your address, and your estimate. The best guesses would receive prizes.

Eight hundred people tried their luck. They were a diverse lot. Many of them were butchers and farmers, who were presumably expert at judging the weight of livestock, but there were also quite a few people who had, as it were, no insider knowledge of cattle. "Many non-experts competed," Galton wrote later in the scientific journal *Nature*, "like those clerks and others who have no expert knowledge of horses, but who bet on races, guided by newspapers, friends, and their own fancies." The analogy to a democracy, in which people of radically different abilities and interests each get one vote, had suggested itself to Galton immediately. "The average competitor was probably as well fitted for making a just estimate of the dressed weight of the ox, as an average voter is of judging the merits of most political issues on which he votes," he wrote.

Galton was interested in figuring out what the "average voter" was capable of because he wanted to prove that the average voter was capable of very little. So he turned the competition into an impromptu experiment. When the contest was over and the prizes had been awarded, Galton borrowed the tickets from the organizers and ran a series of statistical tests on them. Galton arranged the guesses (which totaled 787 in all, after he had to discard thirteen because they were illegible) in order from highest to lowest and graphed them to see if they would form a bell curve. Then, among other things, he added all the contestants' estimates, and calculated the mean of the group's guesses. That number represented, you could say, the collective wisdom of the Plymouth crowd. If the crowd were a single person, that was how much it would have guessed the ox weighed.

Galton undoubtedly thought that the average guess of the group would be way off the mark. After all, mix a few very smart people with some mediocre people and a lot of dumb people, and it seems likely you'd end up with a dumb answer. But Galton was wrong. The crowd had guessed that the ox, after it had been slaughtered and dressed, would weigh 1,197 pounds. After it had been slaughtered and dressed, the ox weighed 1,198 pounds. In other words, the crowd's judgment was essentially perfect. Perhaps breeding did not mean so much after all. Galton wrote later: "The result seems more creditable to the trustworthiness of a democratic judgment than might have been expected." That was, to say the least, an understatement.

II

What Francis Galton stumbled on that day in Plymouth was the simple, but powerful, truth that is at the heart of this book: under the right circumstances, groups are remarkably intelligent, and are often smarter than the smartest people in them. Groups do not need to be dominated by exceptionally intelligent people in order to be smart. Even if most of the people within a group are not especially well-informed or rational, it can still reach a collectively wise

decision. This is a good thing, since human beings are not perfectly designed decision makers. Instead, we are what the economist Herbert Simon called "boundedly rational." We generally have less information than we'd like. We have limited foresight into the future. Most of us lack the ability—and the desire—to make sophisticated cost-benefit calculations. Instead of insisting on finding the best possible decision, we will often accept one that seems good enough. And we often let emotion affect our judgment. Yet despite all these limitations, when our imperfect judgments are aggregated in the right way, our collective intelligence is often excellent.

This intelligence, or what I'll call "the wisdom of crowds," is at work in the world in many different guises. It's the reason the Internet search engine Google can scan a billion Web pages and find the one page that has the exact piece of information you were looking for. It's the reason it's so hard to make money betting on NFL games, and it helps explain why, for the past fifteen years, a few hundred amateur traders in the middle of Iowa have done a better job of predicting election results than Gallup polls have. The wisdom of crowds has something to tell us about why the stock market works (and about why, every so often, it stops working). The idea of collective intelligence helps explain why, when you go to the convenience store in search of milk at two in the morning, there is a carton of milk waiting there for you, and it even tells us something important about why people pay their taxes and help coach Little League. It's essential to good science. And it has the potential to make a profound difference in the way companies do business.

In one sense, this book tries to describe the world as it is, looking at things that at first glance may not seem similar but that are ultimately very much alike. But this book is also about the world as it might be. One of the striking things about the wisdom of crowds is that even though its effects are all around us, it's easy to miss, and, even when it's seen, it can be hard to accept. Most of us, whether as voters or investors or consumers or managers, believe that valuable knowledge is concentrated in a very few hands

(or, rather, in a very few heads). We assume that the key to solving problems or making good decisions is finding that one right person who will have the answer. Even when we see a large crowd of people, many of them not especially well-informed, do something amazing like, say, predict the outcomes of horse races, we are more likely to attribute that success to a few smart people in the crowd than to the crowd itself. As sociologists Jack B. Soll and Richard Larrick put it, we feel the need to "chase the expert." The argument of this book is that chasing the expert is a mistake, and a costly one at that. We should stop hunting and ask the crowd (which, of course, includes the geniuses as well as everyone else) instead. Chances are, it knows.

III

Charles Mackay would have scoffed at the idea that a crowd of people could know anything at all. Mackay was the Scottish journalist who, in 1841, published *Extraordinary Popular Delusions and the Madness of Crowds,* an endlessly entertaining chronicle of mass manias and collective follies, to which the title of my book pays homage. For Mackay, crowds were never wise. They were never even reasonable. Collective judgments were doomed to be extreme. "Men, it has been well said, think in herds," he wrote. "It will be seen that they go mad in herds, while they only recover their senses slowly, and one by one." Mackay's take on collective madness is not an unusual one. In the popular imagination, groups tend to make people either dumb or crazy, or both. The speculator Bernard Baruch, for instance, famously said: "Anyone taken as an individual is tolerably sensible and reasonable—as a member of a crowd, he at once becomes a blockhead." Henry David Thoreau lamented: "The mass never comes up to the standard of its best member, but on the contrary degrades itself to a level with the lowest." Friedrich Nietzsche wrote, "Madness is the exception in indi-

viduals but the rule in groups," while the English historian Thomas Carlyle put it succinctly: "I do not believe in the collective wisdom of individual ignorance."

Perhaps the most severe critic of the stupidity of groups was the French writer Gustave Le Bon, who in 1895 published the polemical classic *The Crowd: A Study of the Popular Mind*. Le Bon was appalled by the rise of democracy in the West in the nineteenth century, and dismayed by the idea that ordinary people had come to wield political and cultural power. But his disdain for groups went deeper than that. A crowd, Le Bon argued, was more than just the sum of its members. Instead, it was a kind of independent organism. It had an identity and a will of its own, and it often acted in ways that no one within the crowd intended. When the crowd did act, Le Bon argued, it invariably acted foolishly. A crowd might be brave or cowardly or cruel, but it could never be smart. As he wrote, "In crowds it is stupidity and not mother wit that is accumulated." Crowds "can never accomplish acts demanding a high degree of intelligence," and they are "always intellectually inferior to the isolated individual." Strikingly, for Le Bon, the idea of "the crowd" included not just obvious examples of collective wildness, like lynch mobs or rioters. It also included just about any kind of group that could make decisions.

So Le Bon lambasted juries, which "deliver verdicts of which each individual juror would disapprove." Parliaments, he argued, adopt laws that each of their members would normally reject. In fact, if you assembled smart people who were specialists in a host of different fields and asked them to "make decisions affecting matters of general interest," the decisions they would reach would be no better, on the whole, than those "adopted by a gathering of imbeciles."

Over the course of this book, I follow Le Bon's lead in giving the words "group" and "crowd" broad definitions, using the words to refer to everything from game-show audiences to multibillion-dollar corporations to a crowd of sports gamblers. Some of the groups in

this book, like the management teams in Chapter 9, are tightly organized and very much aware of their identities as groups. Other crowds, like the herds of cars caught in traffic that I write about in Chapter 7, have no formal organization at all. And still others, like the stock market, exist mainly as an ever-changing collection of numbers and dollars. These groups are all different, but they have in common the ability to act collectively to make decisions and solve problems—even if the people in the groups aren't always aware that's what they're doing. And what is demonstrably true of some of these groups—namely, that they are smart and good at problem solving—is potentially true of most, if not all, of them. In that sense, Gustave Le Bon had things exactly backward. If you put together a big enough and diverse enough group of people and ask them to "make decisions affecting matters of general interest," that group's decisions will, over time, be "intellectually [superior] to the isolated individual," no matter how smart or well-informed he is.

IV

Judging the weight of an ox is hardly a complex task. But, as I suggested above, collective intelligence can be brought to bear on a wide variety of problems, and complexity is no bar. In this book, I concentrate on three kinds of problems. The first are what I'll call *cognition* problems. These are problems that have or will have definitive solutions. For example, "Who will win the Super Bowl this year?" and "How many copies of this new ink-jet printer will we sell in the next three months?" are cognition problems. So, too, is "How likely is it that this drug will be approved by the FDA?" Questions to which there may not be a single right answer, but to which some answers are certainly better than others—such as, "What would be the best place to build this new public swimming pool?"—are cognition problems, too.

The second kind of problem is what's usually called a *coordi-*

nation problem. Coordination problems require members of a group (market, subway riders, college students looking for a party) to figure out how to coordinate their behavior with each other, knowing that everyone else is trying to do the same. How do buyers and sellers find each other and trade at a fair price? How do companies organize their operations? How can you drive safely in heavy traffic? These are all problems of coordination.

The final kind of problem is a *cooperation* problem. As their name suggests, cooperation problems involve the challenge of getting self-interested, distrustful people to work together, even when narrow self-interest would seem to dictate that no individual should take part. Paying taxes, dealing with pollution, and agreeing on definitions of what counts as reasonable pay are all examples of cooperation problems.

A word about structure. The first half of this book is, you might say, theory, although leavened by practical examples. There's a chapter for each of the three problems (cognition, coordination, and cooperation), and there are chapters covering the conditions that are necessary for the crowd to be wise: diversity, independence, and a particular kind of decentralization. The first half begins with the wisdom of crowds, and then explores the three conditions that make it possible, before moving on to deal with coordination and cooperation.

The second part of the book consists of what are essentially case studies. Each of the chapters is devoted to a different way of organizing people toward a common (or at least loosely common) goal, and each chapter is about the way collective intelligence either flourishes or flounders. In the chapter about corporations, for instance, the tension is between a system in which only a few people exercise power and a system in which many have a voice. The chapter about markets starts with the question of whether markets can be collectively intelligent, and ends with a look at the dynamics of a stock-market bubble.

There are many stories in this book of groups making bad

decisions, as well as groups making good ones. Why? Well, one reason is that this is the way the world works. The wisdom of crowds has a far more important and beneficial impact on our everyday lives than we recognize, and its implications for the future are immense. But in the present, many groups struggle to make even mediocre decisions, while others wreak havoc with their bad judgment. Groups work well under certain circumstances, and less well under others. Groups generally need rules to maintain order and coherence, and when they're missing or malfunctioning, the result is trouble. Groups benefit from members talking to and learning from each other, but too much communication, paradoxically, can actually make the group as a whole less intelligent. While big groups are often good for solving certain kinds of problems, big groups can also be unmanageable and inefficient. Conversely, small groups have the virtue of being easy to run, but they risk having too little diversity of thought and too much consensus. Finally, Mackay was right about the extremes of collective behavior: there are times—think of a riot, or a stock-market bubble—when aggregating individual decisions produces a collective decision that is utterly irrational. The stories of these kinds of mistakes are negative proofs of this book's argument, underscoring the importance to good decision making of diversity and independence by demonstrating what happens when they're missing.

Diversity and independence are important because the best collective decisions are the product of disagreement and contest, not consensus or compromise. An intelligent group, especially when confronted with cognition problems, does not ask its members to modify their positions in order to let the group reach a decision everyone can be happy with. Instead, it figures out how to use mechanisms—like market prices, or intelligent voting systems—to aggregate and produce collective judgments that represent not what any one person in the group thinks but rather, in some sense, what they all think. Paradoxically, the best way for a group to be

smart is for each person in it to think and act as independently as possible.

V

I began this Introduction with an example of a group solving a simple problem: figuring out the weight of an ox. I'll end it with an example of a group solving an incredibly complex problem: locating a lost submarine. The differences between the two cases are immense. But the principle in each is the same.

In May 1968, the U.S. submarine *Scorpion* disappeared on its way back to Newport News after a tour of duty in the North Atlantic. Although the navy knew the sub's last reported location, it had no idea what had happened to the *Scorpion*, and only the vaguest sense of how far it might have traveled after it had last made radio contact. As a result, the area where the navy began searching for the *Scorpion* was a circle twenty miles wide and many thousands of feet deep. You could not imagine a more hopeless task. The only possible solution, one might have thought, was to track down three or four top experts on submarines and ocean currents, ask them where they thought the *Scorpion* was, and search there. But, as Sherry Sontag and Christopher Drew recount in their book *Blind Man's Bluff*, a naval officer named John Craven had a different plan.

First, Craven concocted a series of scenarios—alternative explanations for what might have happened to the *Scorpion*. Then he assembled a team of men with a wide range of knowledge, including mathematicians, submarine specialists, and salvage men. Instead of asking them to consult with each other to come up with an answer, he asked each of them to offer his best guess about how likely each of the scenarios was. To keep things interesting, the guesses were in the form of wagers, with bottles of Chivas Regal as prizes. And so Craven's men bet on why the submarine ran into

trouble, on its speed as it headed to the ocean bottom, on the steepness of its descent, and so forth.

Needless to say, no one of these pieces of information could tell Craven where the *Scorpion* was. But Craven believed that if he put all the answers together, building a composite picture of how the *Scorpion* died, he'd end up with a pretty good idea of where it was. And that's exactly what he did. He took all the guesses, and used a formula called Bayes's theorem to estimate the *Scorpion*'s final location. (Bayes's theorem is a way of calculating how new information about an event changes your preexisting expectations of how likely the event was.) When he was done, Craven had what was, roughly speaking, the group's collective estimate of where the submarine was.

The location that Craven came up with was not a spot that any individual member of the group had picked. In other words, not one of the members of the group had a picture in his head that matched the one Craven had constructed using the information gathered from all of them. The final estimate was a genuinely collective judgment that the group as a whole had made, as opposed to representing the individual judgment of the smartest people in it. It was also a genuinely brilliant judgment. Five months after the *Scorpion* disappeared, a navy ship found it. It was 220 yards from where Craven's group had said it would be.

What's astonishing about this story is that the evidence that the group was relying on in this case amounted to almost nothing. It was really just tiny scraps of data. No one knew why the submarine sank, no one had any idea how fast it was traveling or how steeply it fell to the ocean floor. And yet even though no one in the group knew any of these things, the group as a whole knew them all.

1.

THE WISDOM OF CROWDS

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 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?

II

At 11:38 AM on January 28, 1986, the space shuttle *Challenger* lifted off from its launch pad at Cape Canaveral. Seventy-four seconds later, it was ten miles high and rising. Then it blew up. The launch was televised, so news of the accident spread quickly. Eight minutes after the explosion, the first story hit the Dow Jones News Wire.

The stock market did not pause to mourn. Within minutes, investors started dumping the stocks of the four major contractors who had participated in the *Challenger* launch: Rockwell International, which built the shuttle and its main engines; Lockheed, which managed ground support; Martin Marietta, which manufactured the ship's external fuel tank; and Morton Thiokol, which built the solid-fuel booster rocket. Twenty-one minutes after the explosion, Lockheed's stock was down 5 percent, Martin Marietta's was down 3 percent, and Rockwell was down 6 percent.

Morton Thiokol's stock was hit hardest of all. As the finance professors Michael T. Maloney and J. Harold Mulherin report in their fascinating study of the market's reaction to the *Challenger* disaster, so many investors were trying to sell Thiokol stock and so few people were interested in buying it that a trading halt was called almost immediately. When the stock started trading again, almost an hour after the explosion, it was down 6 percent. By the end of the day, its decline had almost doubled, so that at market close, Thiokol's stock was down nearly 12 percent. By contrast, the stocks of the three other firms started to creep back up, and by the end of the day their value had fallen only around 3 percent.

What this means is that the stock market had, almost immediately, labeled Morton Thiokol as the company that was responsible for the *Challenger* disaster. The stock market is, at least in theory, a machine for calculating the present value of all the "free cash flow" a company will earn in the future. (Free cash flow is the money that's left over after a company has paid all its bills and its taxes, has accounted for depreciation, and has invested in the business. It's the money you'd get to take home and put in the bank if you were the sole owner of the company.) The steep decline in Thiokol's stock price—especially compared with the slight declines in the stock prices of its competitors—was an unmistakable sign that investors believed that Thiokol was responsible, and that the consequences for its bottom line would be severe.

As Maloney and Mulherin point out, though, on the day of the disaster there were no public comments singling out Thiokol as the guilty party. While the *New York Times* article on the disaster that appeared the next morning did mention two rumors that had been making the rounds, neither of the rumors implicated Thiokol, and the *Times* declared, "There are no clues to the cause of the accident."

Regardless, the market was right. Six months after the explosion, the Presidential Commission on the *Challenger* revealed that the O-ring seals on the booster rockets made by Thiokol—seals that were supposed to prevent hot exhaust gases from escaping—became less resilient in cold weather, creating gaps that allowed the gases to leak out. (The physicist Richard Feynman famously demonstrated this at a congressional hearing by dropping an O-ring in a glass of ice water. When he pulled it out, the drop in temperature had made it brittle.) In the case of the *Challenger*, the hot gases had escaped and burned into the main fuel tank, causing the cataclysmic explosion. Thiokol was held liable for the accident. The other companies were exonerated.

In other words, within a half hour of the shuttle blowing up, the stock market *knew* what company was responsible. To be sure,

this was a single event, and it's possible that the market's singling out of Thiokol was just luck. Or perhaps the company's business seemed especially susceptible to a downturn in the space program. Possibly the trading halt had sent a signal to investors to be wary. These all are important cautions, but there is still something eerie about what the market did. That's especially true because in this case the stock market was working as a pure weighing machine, undistorted by the factors—media speculation, momentum trading, and Wall Street hype—that make it a peculiarly erratic mechanism for aggregating the collective wisdom of investors. That day, it was just buyers and sellers trying to figure out what happened and getting it right.

How did they get it right? That's the question that Maloney and Mulherin found so vexing. First, they looked at the records of insider trades to see if Thiokol executives, who might have known that their company was responsible, had dumped stock on January 28. They hadn't. Nor had executives at Thiokol's competitors, who might have heard about the O-rings and sold Thiokol's stock short. There was no evidence that anyone had dumped Thiokol stock while buying the stocks of the other three contractors (which would have been the logical trade for someone with inside information). Savvy insiders alone did not cause that first-day drop in Thiokol's price. It was all those investors—most of them relatively uninformed—who simply refused to buy the stock.

But why did they not want Thiokol's stock? Maloney and Mulherin were finally unable to come up with a convincing answer to that question. In the end, they assumed that insider information was responsible for the fall in Thiokol's price, but they could not explain how. Tellingly, they quoted the Cornell economist Maureen O'Hara, who has said, "While markets appear to work in practice, we are not sure how they work in theory."

Maybe. But it depends on what you mean by "theory." If you strip the story down to its basics, after all, what happened that January day was this: a large group of individuals (the actual and po-

tential shareholders of Thiokol's stock, and the stocks of its competitors) was asked a question—"How much less are these four companies worth now that the *Challenger* has exploded?"—that had an objectively correct answer. Those are conditions under which a crowd's average estimate—which is, dollar weighted, what a stock price is—is likely to be accurate. Perhaps someone did, in fact, have inside knowledge of what had happened to the O-rings. But even if no one did, it's plausible that once you aggregated all the bits of information about the explosion that all the traders in the market had in their heads that day, it added up to something close to the truth. As was true of those who helped John Craven find the *Scorpion*, even if none of the traders was sure that Thiokol was responsible, collectively they were certain it was.

The market was smart that day because it satisfied the four conditions that characterize wise crowds: diversity of opinion (each person should have some private information, even if it's just an eccentric interpretation of the known facts), independence (people's opinions are not determined by the opinions of those around them), decentralization (people are able to specialize and draw on local knowledge), and aggregation (some mechanism exists for turning private judgments into a collective decision). If a group satisfies those conditions, its judgment is likely to be accurate. Why? At heart, the answer rests on a mathematical truism. If you ask a large enough group of diverse, independent people to make a prediction or estimate a probability, and then average those estimates, the errors each of them makes in coming up with an answer will cancel themselves out. Each person's guess, you might say, has two components: information and error. Subtract the error, and you're left with the information.

Now, even with the errors canceled out, it's possible that a group's judgment will be bad. For the group to be smart, there has to be at least some information in the "information" part of the "information minus error" equation. (If you'd asked a large group of

children to buy and sell stocks in the wake of the *Challenger* disaster, it's unlikely they would have picked out Thiokol as the culprit.) What is striking, though—and what makes a phrase like "the wisdom of crowds" meaningful—is just how much information a group's collective verdict so often contains. In cases like Francis Galton's experiment or the *Challenger* explosion, the crowd is holding a nearly complete picture of the world in its collective brain.

Perhaps this isn't surprising. After all, we are the products of evolution, and presumably we have been equipped to make sense of the world around us. But who knew that, given the chance, we can collectively make so *much* sense of the world. After all, think about what happens if you ask a hundred people to run a 100-meter race, and then average their times. The average time will not be better than the time of the fastest runners. It will be worse. It will be a mediocre time. But ask a hundred people to answer a question or solve a problem, and the average answer will often be at least as good as the answer of the smartest member. With most things, the average is mediocrity. With decision making, it's often excellence. You could say it's as if we've been programmed to be collectively smart.

III

Truly successful decision making, of course, demands more than just a picture of the world as it is. It demands in addition a picture of the world as it will (or at least as it may) be. Any decision-making mechanism therefore has to be good under conditions of uncertainty. And what's more uncertain than the future? Group intelligence may be good at telling how many jelly beans are in a jar or remembering the year Nirvana released *Nevermind*. But how does it perform under conditions of true uncertainty, when the right answer is seemingly unknowable—because it hasn't happened yet?

Robert Walker's entire career depends on the answer to that

question. Walker is the sports book director at the Mirage Hotel and Casino in Las Vegas, which means that every week he fields thousands of bets in sports ranging from pro football to Ivy League basketball. For all those games, Walker has to offer a line (or point spread), which lets bettors know which team is favored to win and by how many points. The way the line works is simple. Say the Giants are favored this week by three and a half points over the Rams. If you bet on the Giants, they have to win by four points or more for you to win the bet. Conversely, if you bet on the Rams, they have to lose by three points or less (or win), for you to walk away with the casino's money. In other sports, bets are framed in terms of odds: if you bet on the favorite, you might have to put down $150 to get $100 back, while if you bet on the underdog, you'd have to lay down $75 to win $100.

As a bookmaker, Walker's job is not to try to pick what team will win. He leaves that to the gamblers, at least in theory. Instead, his job is to make sure that the gamblers bet roughly the same amount of money on one team as on the other. If he does that, then he knows that he will win half the bets he's taken in and lose the other half. Why would Walker be satisfied with just breaking even? Because bookies make more money on every bet they win than they lose on every bet they get wrong. If you place a point-spread bet with a bookie, you have to put up $11 to win $10. Imagine there are only two bettors, one who bets on the favorite and the other who bets on the underdog. Walker takes in $22 ($11 from each of them). He pays out $21 to the winner. The $1 he keeps is his profit. That slim advantage, which is known as the vigorish, or the vig, is what pays the bookie's bills. And the bookie keeps that advantage only when he avoids having too much money riding on one side of a bet.

To keep that from happening, Walker needs to massage the point spread so that bets keep coming in for both teams. "The line we want is the line that'll split the public, because that's when you

start earning that vig," he said. In the week before the 2001 Super Bowl, for instance, the Mirage's opening line had the Baltimore Ravens favored by two and a half points. But soon after the line was posted, the Mirage booked a couple of early $3,000 bets on Baltimore. That's not much money, but it was enough to convince Walker to raise the point spread to three. If everyone wanted to bet on Baltimore, chances were the line wasn't right. So the line moved. The opening line is set by the bookmaker, but it shifts largely in response to what bettors do—much as stock prices rise and fall with investor demand.

In theory, you could set the opening line wherever, and simply allow it to adjust from there automatically, so that the point spread would rise or fall anytime there was a significant imbalance between the amounts wagered on each side. The Mirage would have no problem doing this; its computerized database tracks the bets as they come in. But bookies place a premium on making the opening line as accurate as possible, because if they set it badly they're going to get stuck taking a lot of bad bets. Once a line opens, though, it's out of the bookie's hands, and a game's point spread ends up representing bettors' collective judgment of what the final outcome of that game will be. As Bob Martin, who was essentially the country's oddsmaker in the 1970s, said, "Once you put a number on the board, it becomes public property."

The public, it turns out, is pretty smart. It does not have a crystal ball: point spreads only weakly predict the final scores of most NFL games, for instance. But it is very hard for even well-informed gamblers to beat the final spread consistently. In about half the games, favorites cover the spread, while in the other half underdogs beat the spread. This is exactly what a bookie wants to have happen. And there are no obvious mistakes in the market's judgment—like, say, home teams winning more than the crowd predicts they will, or road underdogs being consistently undervalued. Flaws in the crowd's judgment are found occasionally, but

when they are they're typically like the one documented in a recent paper that found that in weeks fifteen, sixteen, and seventeen of the NFL season, home underdogs have historically been a good bet. So you have to search hard to outperform the betting crowd. Roughly three-quarters of the time, the Mirage's final line will be the most reliable forecast of the outcomes of NFL games that you can find.

The same is true in many other sports. Because sports betting is a kind of ready-made laboratory to study predictions and their outcomes, a host of academics have perused gambling markets to see how efficient—that is, how good at capturing all the available information—they are. The results of their studies are consistent: in general, in most major sports the market is relatively efficient. In some cases, the crowd's performance is especially good: in horse racing, for instance, the final odds reliably predict the race's order of finish (that is, the favorite wins most often, the horse with the second-lowest odds is the second-most-often winner, and so on) and also provide, in economist Raymond D. Sauer's words, "reasonably good estimates of the probability of winning." In other words, a three-to-one horse will win roughly a quarter of the time. There are exceptions: odds are less accurate in those sports and games where the betting market is smaller and less liquid (meaning that the odds can change dramatically thanks to only a few bets), like hockey or golf or small-college basketball games. These are often the sports where professional gamblers can make real money, which makes sense given that we know the bigger the group, the more accurate it becomes. And there are also some interesting quirks: in horse racing, for instance, people tend to bet on long shots slightly more often than they should and bet on favorites slightly less often than they should. (This seems to be a case of risk-seeking behavior: bettors, especially bettors who have been losing, would rather take a flyer on a long shot that offers the possibility of big returns than grind it out by betting on short-odds favorites.) But on the whole, if bettors aren't collectively foreseeing the future, they're doing the next best thing.

IV

Recently I decided I needed—*this minute!*—the exact text of Bill Murray's *Caddyshack* riff about toting the Dalai Lama's golf bag. The punch line of the riff is "So I got that going for me, which is nice" and the Dalai Lama, in Murray's telling, likes to say "Gunga galunga." So I went to Google, the Internet search engine, typed in "going for me" and "gunga," and hit the search button. A list of 695 Web pages came back. First on the list was an article from GolfOnline, which included the second half of the riff. That was okay, but third on the list was a Web site for something called the Penn State Soccer Club. The goalie, a guy named David Feist, had posted the entire monologue. The search took 0.18 seconds.

Then I needed to check out the Mulherin paper on the *Challenger* that I discuss above. I couldn't remember the author's name, so I typed in "'stock market' challenger reaction": 2,370 pages came back. The first one was an article by *Slate*'s Daniel Gross about the Mulherin paper. The third was Mulherin's own Web site, with a link to his paper. That search—which, remember, did not include Mulherin's name—took 0.10 seconds. A few minutes later, my search for the lyrics to a Ramones song about Ronald Reagan visiting the Bitburg cemetery took 0.23 seconds, and the first item on the list had what I needed.

If you use the Internet regularly, these examples of Google's performance will not surprise you. This is what we have come to expect from Google: instantaneous responses with the exact page we need up high in the rankings. But if possible, it's worth letting yourself be a little amazed at what happened during those routine searches. Each time, Google surveyed billions of Web pages and

picked exactly the pages that I would find most useful. The cumulative time for all the searches: about a minute and a half.

Google started in 1998, at a time when Yahoo! seemed to have a stranglehold on the search business—and if Yahoo! stumbled, then AltaVista or Lycos looked certain to be the last man standing. But within a couple of years, Google had become the default search engine for anyone who used the Internet regularly, simply because it was able to do a better job of finding the right page quickly. And the way it does that—and does it while surveying three billion Web pages—is built on the wisdom of crowds.

Google keeps the details of its technology to itself, but the core of the Google system is the PageRank algorithm, which was first defined by the company's founders, Sergey Brin and Lawrence Page, in a now-legendary 1998 paper called "The Anatomy of a Large-Scale Hypertextual Web Search Engine." PageRank is an algorithm—a calculating method—that attempts to let all the Web pages on the Internet decide which pages are most relevant to a particular search. Here's how Google puts it:

> PageRank capitalizes on the uniquely democratic characteristic of the web by using its vast link structure as an organizational tool. In essence, Google interprets a link from page A to page B as a vote, by page A, for page B. Google assesses a page's importance by the votes it receives. But Google looks at more than sheer volume of votes, or links; it also analyzes the page that casts the vote. Votes cast by pages that are themselves "important" weigh more heavily and help to make other pages "important."

In that 0.12 seconds, what Google is doing is asking the entire Web to decide which page contains the most useful information, and the page that gets the most votes goes first on the list. And that page, or the one immediately beneath it, more often than not is in fact the one with the most useful information.

Now, Google is a republic, not a perfect democracy. As the description says, the more people that have linked to a page, the more influence that page has on the final decision. The final vote is a "weighted average"—just as a stock price or an NFL point spread is—rather than a simple average like the ox-weighers' estimate. Nonetheless, the big sites that have more influence over the crowd's final verdict have that influence only because of all the votes that smaller sites have given them. If the smaller sites were giving the wrong sites too much influence, Google's search results would not be accurate. In the end, the crowd still rules. To be smart at the top, the system has to be smart all the way through.

V

If allowing people to bet on sporting events effectively creates a kind of machine that's good at predicting the outcome of those events, an obvious question follows: Wouldn't people betting on other kinds of events be equally good, as a group, at predicting them? Why confine ourselves to forecasting the results of basketball games if we could also come up with accurate predictions of, say, presidential elections?

Of course, we already have a well-established way of predicting presidential elections: the poll. If you want to know how people are going to vote, you just ask them. Polling is, relatively speaking, accurate. It has a solid methodology behind it, and is statistically rigorous. But there's reason to wonder if a market such as the betting market—one that allowed the people participating in it to rely on many different kinds of information, including but not limited to polls—might at the very least offer a competitive alternative to Gallup. That's why the Iowa Electronic Markets (IEM) project was created.

Founded in 1988 and run by the College of Business at the University of Iowa, the IEM features a host of markets designed to predict the outcomes of elections—presidential, congressional, gu-

bernatorial, and foreign. Open to anyone who wants to participate, the IEM allows people to buy and sell futures "contracts" based on how they think a given candidate will do in an upcoming election. While the IEM offers many different types of contracts, two are most common. One is designed to predict the winner of an election. In the case of the California recall in 2003, for instance, you could have bought an "Arnold Schwarzenegger to win" contract, which would have paid you $1 when Schwarzenegger won. Had he lost, you would have gotten nothing. The price you pay for this kind of contract reflects the market's judgment of a candidate's chances of victory. If a candidate's contract costs 50 cents, it means, roughly speaking, that the market thinks he has a 50 percent chance of winning. If it costs 80 cents, he has an 80 percent chance of winning, and so on.

The other major kind of IEM contract is set up to predict what percentage of the final popular vote a candidate will get. In this case, the payoffs are determined by the vote percentage: if you'd bought a George W. Bush contract in 2004, you would have received 51 cents (he got 51 percent of the vote) when the election was over.

If the IEM's predictions are accurate, the prices of these different contracts will be close to their true values. In the market to predict election winners, the favorite should win more often, and bigger favorites should win by bigger margins. Similarly, in the vote-share market, if a candidate ends up getting 49 percent of the vote in an election, then the price of his contract in the run-up to election day should be close to 49 cents.

So how has the IEM done? Well, a study of the IEM's performance in forty-nine different elections between 1988 and 2000 found that the election-eve prices in the IEM were, on average, off by just 1.37 percent in presidential elections, 3.43 percent in other U.S. elections, and 2.12 percent in foreign elections. (Those numbers are in absolute terms, meaning that the market would have been off by 1.37 percent if, say, it had predicted that Al Gore would get 48.63 percent of the vote when in reality he got 50 percent.)

The IEM has generally outperformed the major national polls, and has been more accurate than those polls even months in advance of the actual election. Over the course of the presidential elections between 1988 and 2000, for instance, 596 different polls were released. Three-fourths of the time, the IEM's market price on the day each of those polls was released was more accurate. Polls tend to be very volatile, with vote shares swinging wildly up and down. But the IEM forecasts, though ever-changing, are considerably less volatile, and tend to change dramatically only in response to new information. That makes them more reliable as forecasts.

What's especially interesting about this is that the IEM isn't very big—there have never been more than a few thousand traders in the market—and it doesn't, in any way, reflect the makeup of the electorate as a whole. The vast majority of traders are men, and a disproportionate—though shrinking—number of them are from Iowa. So the people in the market aren't predicting their own behavior. But their predictions of what the voters of the country will do are better than the predictions you get when you ask the voters themselves what they're going to do. And while the IEM traders undoubtedly get information from the polls, the superior accuracy of their collective forecasts suggests that the traders are also adding information to what's in the polls.

The IEM's success has helped inspire other similar markets, including the Hollywood Stock Exchange (HSX), which allows people to wager on box-office returns, opening-weekend performance, and the Oscars. The HSX enjoyed its most notable success in March of 2000. That was when a team of twelve reporters from *The Wall Street Journal* assiduously canvassed members of the Academy of Motion Pictures Arts and Sciences in order to find out how they had voted. The Academy was not happy about this. The organization's president publicly attacked the *Journal* for trying to "scoop us before Oscar night," and the Academy urged members not to talk to reporters. But with the *Journal* promising anonymity, more than a few people—356, or about 6 percent of all members—

disclosed how they had filled out their ballots. The Friday before the ceremony, the *Journal* published its results, forecasting the winners in the six major Oscar categories—Best Picture, Best Director, Best Actor and Best Actress, Best Supporting Actor and Best Supporting Actress. And when the envelopes were opened, the *Journal*'s predictions—much to the Academy's dismay—turned out to be pretty much on target, with the paper picking five of the six winners. The HSX, though, had done even better, getting all six of the six right. In 2002, the exchange, perhaps even more impressively, picked thirty-five of the eventual forty Oscar nominees.

The HSX's box-office forecasts are not as impressive or as accurate as the IEM's election forecasts. But Anita Elberse, a professor of marketing at Harvard Business School, has compared the HSX's forecasts to other Hollywood prediction tools, and found that the HSX's closing price the night before a movie opens is the single best available forecast of its weekend box office. As a result, the HSX's owner, Cantor Index Holdings, is now marketing its data to Hollywood studios.

One of the interesting things about markets like the IEM and the HSX is that they work fairly well without much—or any—money at stake. The IEM is a real-money market, but the most you can invest is $500, and the average trader has only $50 at stake. In the HSX, the wagering is done entirely with play money. All the evidence we have suggests that people focus better on a decision when there are financial rewards attached to it (which may help explain why the IEM's forecasts tend to be more accurate). But David Pennock—a researcher at Overture who has studied these markets closely—found that, especially for active traders in these markets, status and reputation provided incentive enough to encourage a serious investment of time and energy in what is, after all, a game.

As the potential virtues of these decision markets have become obvious, the range of subjects they cover has grown rapidly. At the NewsFutures and TradeSports exchanges, people could bet,

in the fall of 2003, on whether or not Kobe Bryant would be convicted of sexual assault, on whether and when weapons of mass destruction would be found in Iraq, and on whether Ariel Sharon would remain in power longer than Yasir Arafat. Ely Dahan, a professor at UCLA, has experimented with a classroom-decision market in which students bought and sold securities representing a variety of consumer goods and services, including SUVs, ski resorts, and personal digital assistants. (In a real-life market of this kind, the value of a security might depend on the first-year sales of a particular SUV.) The market's forecasts were eerily similar to the predictions that conventional market research had made (but the classroom research was much cheaper). In the fall of 2003, meanwhile, MIT's *Technology Review* set up a site called Innovation Futures, where people could wager on future technological developments. And Robin Hanson, an economics professor at George Mason University who was one of the first to write about the possibility of using decision markets in myriad contexts, has suggested that such markets could be used to guide scientific research and even as a tool to help governments adopt better policies.

Some of these markets will undoubtedly end up being of little use, either because they'll fail to attract enough participants to make intelligent forecasts or because they'll be trying to predict the unpredictable. But given the right conditions and the right problems, a decision market's fundamental characteristics—diversity, independence, and decentralization—are guaranteed to make for good group decisions. And because such markets represent a relatively simple and quick means of transforming many diverse opinions into a single collective judgment, they have the chance to improve dramatically the way organizations make decisions and think about the future.

In that sense, the most mystifying thing about decision markets is how little interest corporate America has shown in them. Corporate strategy is all about collecting information from many different sources, evaluating the probabilities of potential out-

comes, and making decisions in the face of an uncertain future. These are tasks for which decision markets are tailor-made. Yet companies have remained, for the most part, indifferent to this source of potentially excellent information, and have been surprisingly unwilling to improve their decision making by tapping into the collective wisdom of their employees. We'll look more closely at people's discomfort with the idea of the wisdom of crowds, but the problem is simple enough: just because collective intelligence is real doesn't mean that it will be put to good use.

A DECISION MARKET is an elegant and well-designed method for capturing the collective wisdom. But the truth is that the specific method that one uses probably doesn't matter very much. In this chapter, we've looked at a host of different ways of tapping into what a group knows: stock prices, votes, point spreads, pari-mutuel odds, computer algorithms, and futures contracts. Some of these methods seem to work better than others, but in the end there's nothing about a futures market that makes it inherently smarter than, say, Google or a pari-mutuel pool. These are all attempts to tap into the wisdom of the crowd, and that's the reason they work. The real key, it turns out, is not so much perfecting a particular method, but satisfying the conditions—diversity, independence, and decentralization—that a group needs to be smart. As we'll see in the chapters that follow, that's the hardest, but also perhaps the most interesting, part of the story.

2.

THE DIFFERENCE DIFFERENCE MAKES: WAGGLE DANCES, THE BAY OF PIGS, AND THE VALUE OF DIVERSITY

I

In 1899, Ransom E. Olds opened the Olds Motor Works in Detroit, Michigan. Olds had been in the automobile business since the mid-1880s, when he built his first car, a steam-powered vehicle with three wheels. But success had remained elusive. After moving on to gasoline-powered cars, Olds started his own company in the early 1890s, but it floundered, leaving him nearly destitute. He was only able to start the Motor Works, in fact, by convincing a financier named Samuel Smith to put up nearly all the money. Olds got his company, but he also got a boss to whom he had to answer. This was a problem, because the two did not agree on what the Olds Motor Works should be making. Smith thought the company should cater to the high end of the market, building large, expensive cars with all the trimmings. Olds, though, was more intrigued by the possibility of building a car that could be marketed to the middle class. In 1900, the auto market was still minuscule—there were fewer than 15,000 cars on the road that year. But it seemed plausible that an invention as revolutionary as the car would be able to find a mass audience, if you could figure out a way to make one cheaply enough.

4.

PUTTING THE PIECES TOGETHER: THE CIA, LINUX, AND THE ART OF DECENTRALIZATION

I

In April 1946, at a forum organized by the *New York Herald-Tribune*, General Wild Bill Donovan gave a speech entitled "Our Foreign Policy Needs a Central Intelligence Agency." During World War II, Donovan had been the head of the Office of Strategic Services, the United States' chief wartime intelligence organization, and once the war ended he became a loud public advocate for the creation of a more powerful peacetime version of the OSS. Before the war, the United States had divided intelligence-gathering responsibilities among the different military services. But the failure of any of those services to anticipate the attack on Pearl Harbor—despite what seemed, in retrospect, to be ample evidence that a major Japanese strike was in the works—had pointed up the system's limitations and suggested the need for a more comprehensive approach to intelligence gathering. So, too, did the prospect of conflict with the Soviet Union, which even in 1946 loomed as a real possibility, and the advent of new technologies—Donovan cited "the rocket, the atomic bomb, bacteriological warfare"—that made America's borders seem far from impregnable. In his April speech,

Donovan hit on all of these themes, arguing that what the United States needed was "a centralized, impartial, independent agency" to take charge of all of the country's intelligence operations.

Donovan's public speaking didn't do much for his own career, since his sharp criticisms alienated the intelligence community and probably doomed his chances of returning to government service. Nonetheless, in 1947, Congress passed the National Security Act and created the Central Intelligence Agency. As historian Michael Warner has put it, the goal of the law was to "implement the principles of unity of command and unity of intelligence." Fragmentation and division had left the United States vulnerable to surprise attack. Centralization and unity would keep it safe in the future.

In fact, though, the centralization of intelligence never happened. Although the CIA was initially the key player in the postwar period, as time passed the intelligence community became more fragmented than ever, divided into a kind of alphabet soup of agencies with overlapping responsibilities and missions, including not just the CIA but also the National Security Agency, the National Imagery and Mapping Agency, the National Reconnaissance Office, the Defense Intelligence Agency, and the intelligence arms of each of the three major military services. In theory, the director of the CIA was in charge of the U.S. intelligence community as a whole, but in practice he exercised very little supervision over these agencies, and most of the money for intelligence operations came from the Department of Defense. In addition, the FBI—which was responsible for domestic law enforcement—operated almost completely outside the orbit of this intelligence community, even though information about foreign terrorists operating inside the United States would obviously be of interest to the CIA. In place of the centralized repository of information and analysis that Donovan had envisioned, the U.S. intelligence community evolved into a collection of virtually autonomous, decentralized groups, all working toward the same broad goal—keeping the United States safe from attack—but in very different ways.

Until September 11, 2001, the flaws of this system were overlooked. The intelligence community had failed to anticipate the 1993 bombing of the World Trade Center and the 1998 bombings of the U.S. embassy in Kenya and the USS *Cole* in Yemen. But not until September 11 did the failure of U.S. intelligence gathering come to seem undeniable. The Congressional Joint Inquiry into the attacks found that the U.S. intelligence community had "failed to capitalize on both the individual and collective significance of available information that appears relevant to the events of September 11." Intelligence agencies "missed opportunities to disrupt the September 11th plot," and allowed information to pass by unnoticed that, if appreciated, would have "greatly enhanced its chances of uncovering and preventing" the attacks. It was, in other words, Pearl Harbor all over again.

The congressional inquiry was unquestionably a classic example of Monday-morning quarterbacking. Given the sheer volume of information that intelligence agencies process, it's hardly surprising that a retrospective look at the data they had on hand at the time of the attack would uncover material that seemed relevant to what happened on September 11. That doesn't necessarily mean the agencies could have been realistically expected to recognize the relevance of the material beforehand. In her classic account of the intelligence failures at Pearl Harbor, *Warning and Decision*, Roberta Wohlstetter shows how many signals there were of an impending Japanese attack, but suggests that it was still unreasonable to expect human beings to have picked the right signals out from "the buzzing and blooming confusion" that accompanied them. Strategic surprise, Wohlstetter suggests, is an intractable problem to solve. And if a massive Japanese naval attack comprising hundreds of planes and ships and thousands of men was difficult to foresee, how much harder would it have been to predict a terrorist attack involving just nineteen men?

And yet one has to wonder. Given the almost complete failure of the intelligence community to anticipate any of four major

terrorist attacks from 1993 through 2001, is it not possible that organizing the intelligence community differently would have, at the very least, improved its chances of recognizing what the Joint Inquiry called "the collective significance" of the data it had on hand? Predicting the actual attacks on the World Trade Center and the Pentagon may have been impossible. But coming up with a reasonable, concrete estimate of the likelihood of such an attack may not have been.

That, at least, was the conclusion that Congress reached: better processes would have produced a better result. In particular, it stressed the lack of "information sharing" between the various agencies. Instead of producing a coherent picture of the threats the United States faced, the various agencies produced a lot of localized snapshots. The sharpest critic of the agencies' work, Senator Richard Shelby, argued that the FBI in particular was crippled by its "decentralized organizational structure," which "left information-holdings fragmented into largely independent fiefdoms." And the intelligence community as a whole was hurt by a failure to put the right information in the hands of the right people. What needed to be done, Shelby suggested, was to abolish the fiefdoms and return to the idea for which Bill Donovan had argued half a century ago. One agency, which could stand "above and independent from the disputatious bureaucracies," needed to be put in charge of U.S. intelligence. Decentralization had led the United States astray. Centralization would put things right.

II

In challenging the virtues of decentralization, Shelby was challenging an idea that in the past fifteen years has seized the imagination of businessmen, academics, scientists, and technologists everywhere. In business, management theories like reengineering advocated replacing supervisors and managers with self-managed teams

that were responsible for solving most problems on their own, while more utopian thinkers deemed the corporation itself outmoded. In physics and biology, scientists paid increasing attention to self-organizing, decentralized systems—like ant colonies or beehives—which, even without a center, proved robust and adaptable. And social scientists placed renewed emphasis on the importance of social networks, which allow people to connect and coordinate with each other without a single person being in charge. Most important, of course, was the rise of the Internet—in some respects, the most visible decentralized system in the world—and of corollary technologies like peer-to-peer file sharing (exemplified by Napster), which offered a clear demonstration of the possibilities (economic, organizational, and more) that decentralization had to offer.

The idea of the wisdom of crowds also takes decentralization as a given and a good, since it implies that if you set a crowd of self-interested, independent people to work in a decentralized way on the same problem, instead of trying to direct their efforts from the top down, their collective solution is likely to be better than any other solution you could come up with. American intelligence agents and analysts were self-interested, independent people working in a decentralized way on roughly the same problem (keeping the country safe). So what went wrong? Why did those agents not produce a better forecast? Was decentralization really the problem?

BEFORE WE ANSWER THAT question, we need to answer a simpler one first: What do we mean by "decentralization," anyway? It's a capacious term, and in the past few years it's been tossed around more freely than ever. Flocks of birds, free-market economies, cities, peer-to-peer computer networks: these are all considered examples of decentralization. Yet so, too, in other contexts, are the American public-school system and the modern corporation. These systems are dramatically different from each other, but they do have this in common: in each, power does not fully reside in one

central location, and many of the important decisions are made by individuals based on their own local and specific knowledge rather than by an omniscient or farseeing planner.

In terms of decision making and problem solving, there are a couple of things about decentralization that really matter. It fosters, and in turn is fed by, specialization—of labor, interest, attention, or what have you. Specialization, as we've known since Adam Smith, tends to make people more productive and efficient. And it increases the scope and the diversity of the opinions and information in the system (even if each individual person's interests become more narrow).

Decentralization is also crucial to what the economist Friedrich Hayek described as tacit knowledge. Tacit knowledge is knowledge that can't be easily summarized or conveyed to others, because it is specific to a particular place or job or experience, but it is nonetheless tremendously valuable. (In fact, figuring out how to take advantage of individuals' tacit knowledge is a central challenge for any group or organization.) Connected with this is the assumption that is at the heart of decentralization, namely that the closer a person is to a problem, the more likely he or she is to have a good solution to it. This practice dates back to ancient Athens, where decisions about local festivals were left up to the *demes*, as opposed to the Athenian assembly, and regional magistrates handled most nonserious crimes. It can also be seen in Exodus, where Moses' father-in-law counseled him to judge only in "great matter[s]" and to leave all other decisions to local rulers.

Decentralization's great strength is that it encourages independence and specialization on the one hand while still allowing people to coordinate their activities and solve difficult problems on the other. Decentralization's great weakness is that there's no guarantee that valuable information which is uncovered in one part of the system will find its way through the rest of the system. Sometimes valuable information never gets disseminated, making it less useful than it otherwise would be. What you'd like is a way for in-

dividuals to specialize and to acquire local knowledge—which increases the total amount of information available in the system—while also being able to aggregate that local knowledge and private information into a collective whole, much as Google relies on the local knowledge of millions of Web-page operators to make Google searches ever-smarter and ever-quicker. To accomplish this, any "crowd"—whether it be a market, a corporation, or an intelligence agency—needs to find the right balance between the two imperatives: making individual knowledge globally and collectively useful (as we know it can be), while still allowing it to remain resolutely specific and local.

III

In 1991, Finnish hacker Linus Torvalds created his own version of the Unix operating system, dubbing it Linux. He then released the source code he had written to the public, so everyone out there—well, everyone who understood computer code—could see what he had done. More important, he attached a note that read, "If your efforts are freely distributable, I'd like to hear from you, so I can add them to the system." It was a propitious decision. As one history of Linux points out: "Of the first ten people to download Linux, five sent back bug fixes, code improvements, and new features." Over time, this improvement process became institutionalized, as thousands of programmers, working for free, contributed thousands of minor and major fixes to the operating system, making Linux ever-more reliable and robust.

Unlike Windows, which is owned by Microsoft and worked on only by Microsoft employees, Linux is owned by no one. When a problem arises with the way Linux works, it only gets fixed if someone, on his own, offers a good solution. There are no bosses ordering people around, no organizational charts dictating people's responsibilities. Instead, people work on what they're interested in

and ignore the rest. This seems like—in fact, it is—a rather haphazard way to solve problems. But so far, at least, it has been remarkably effective, making Linux the single most important challenger to Microsoft.

Linux is clearly a decentralized system, since it has no formal organization and its contributors come from all over the world. What decentralization offers Linux is diversity. In the traditional corporate model, top management hires the best employees it can, pays them to work full-time, generally gives them some direction about what problems to work on, and hopes for the best. That is not a bad model. It has the great virtue of making it easy to mobilize people to work on a particular problem, and it also allows companies to get very good at doing the things they know how to do. But it also necessarily limits the number of possible solutions that a corporation can come up with, both because of mathematical reality (a company has only so many workers, and they have only so much time) and because of the reality of organizational and bureaucratic politics. Linux, practically speaking, doesn't worry much about either. Surprisingly, there seems to be a huge supply of programmers willing to contribute their efforts to make the system better. That guarantees that the field of possible solutions will be immense. There's enough variety among programmers, and there are enough programmers, that no matter what the bug is, someone is going to come up with a fix for it. And there's enough diversity that someone will recognize bugs when they appear. In the words of open-source guru Eric Raymond, "Given enough eyeballs, all bugs are shallow."

In the way it operates, in fact, Linux is not all that different from a market, as we saw in Chapter 2 on diversity. Like a bee colony, it sends out lots of foragers and assumes that one of them will find the best route to the flower fields. This is, without a doubt, less efficient than simply trying to define the best route to the field or even picking the smartest forager and letting him go. After all, if hundreds or thousands of programmers are spending

their time trying to come up with a solution that only a few of them are going to find, that's many hours wasted that could be spent doing something else. And yet, just as the free market's ability to generate lots of alternatives and then winnow them down is central to its continued growth, Linux's seeming wastefulness is a kind of strength (a kind of strength that for-profit companies cannot, fortunately or unfortunately, rely on). You can let a thousand flowers bloom and then pick the one that smells the sweetest.

IV

So who picks the sweetest-smelling one? Ideally, the crowd would. But here's where striking a balance between the local and the global is essential: a decentralized system can only produce genuinely intelligent results if there's a means of aggregating the information of everyone in the system. Without such a means, there's no reason to think that decentralization will produce a smart result. In the case of the experiment with which this book opened, that aggregating mechanism was just Frances Galton counting the votes. In the case of the free market, that aggregating mechanism is obviously price. The price of a good reflects, imperfectly but effectively, the actions of buyers and sellers everywhere, and provides the necessary incentive to push the economy where the buyers and sellers want it to go. The price of a stock reflects, imperfectly but effectively, investors' judgment of how much a company is worth. In the case of Linux, it is the small number of coders, including Torvalds himself, who vet every potential change to the operating-system source code. There are would-be Linux programmers all over the world, but eventually all roads lead to Linus.

Now, it's not clear that the decision about what goes into Linux's code needs to be or should be in the hands of such a small group of people. If my argument in this book is right, a large group of programmers, even if they weren't as skilled as Torvalds and his

lieutenants, would do an excellent job of evaluating which code was worth keeping. But set that aside. The important point here is that if the decision were not being made by someone, Linux itself would not be as successful as it is. If a group of autonomous individuals tries to solve a problem without any means of putting their judgments together, then the best solution they can hope for is the solution that the smartest person in the group produces, and there's no guarantee they'll get that. If that same group, though, has a means of aggregating all those different opinions, the group's collective solution may well be smarter than even the smartest person's solution. Aggregation—which could be seen as a curious form of centralization—is therefore paradoxically important to the success of decentralization. If this seems dubious, it may be because when we hear centralization we think "central planners," as in the old Soviet Union, and imagine a small group of men—or perhaps just a single man—deciding how many shoes will be made today. But in fact there's no reason to confuse the two. It's possible, and desirable, to have collective decisions made by decentralized agents.

Understanding when decentralization is a recipe for collective wisdom matters because in recent years the fetish for decentralization has sometimes made it seem like the ideal solution for every problem. Obviously, given the premise of this book, I think decentralized ways of organizing human effort are, more often than not, likely to produce better results than centralized ways. But decentralization works well under some conditions and not very well under others. In the past decade, it's been easy to believe that if a system is decentralized, then it must work well. But all you need to do is look at a traffic jam—or, for that matter, at the U.S. intelligence community—to recognize that getting rid of a central authority is not a panacea. Similarly, people have become enamored of the idea that decentralization is somehow *natural* or *automatic*, perhaps because so many of our pictures of what decentralization looks like come from biology. Ants, after all, don't need to do any-

thing special to form an ant colony. Forming ant colonies is inherent in their biology. The same is not, however, true of human beings. It's hard to make real decentralization work, and hard to keep it going, and easy for decentralization to become disorganization.

A good example of this was the performance of the Iraqi military during the U.S.-Iraq war in 2003. In the early days of the war, when Iraqi fedayeen paramilitaries had surprised U.S. and British troops with the intensity of their resistance, the fedayeen were held up as an example of a successful decentralized group, which was able to flourish in the absence of any top-down control. In fact, one newspaper columnist compared the fedayeen to ants in an ant colony, finding their way to a "good" solution while communicating only with the soldiers right next to them. But after a few days, the idea that the fedayeen were mounting a meaningful, organized resistance vanished, as it became clear that their attacks were little more than random, uncoordinated assaults that had no connection to what was happening elsewhere in the country. As one British commander remarked, it was all tactics and no strategy. To put it differently, the individual actions of the fedayeen fighters never added up to anything bigger, precisely because there was no method of aggregating their local wisdom. The fedayeen were much like ants—following local rules. But where ants who follow their local rules actually end up fostering the well-being of the colony, soldiers who followed their local rules ended up dead. (It may be, though, that once the actual war was over, and the conflict shifted to a clash between the occupying U.S. military and guerrillas using hit-and-run terrorist tactics, the absence of aggregation became less important, since the goal was not to defeat the United States in battle, but simply to inflict enough damage to make staying seem no longer worth it. In that context, tactics may have been enough.)

The irony is that the true decentralized military in the U.S.-Iraq war was the U.S. Army. American troops have always been given significantly more initiative in the field than other armies, as

the military has run itself on the "local knowledge is good" theory. But in recent years, the army has dramatically reinvented itself. Today, local commanders have considerably greater latitude to act, and sophisticated communications systems mean that collectively wise strategies can emerge from local tactics. Commanders at the top are not isolated from what's happening in the field, and their decisions will inevitably reflect, in a deep sense, the local knowledge that field commanders are acquiring. In the case of the invasion of Baghdad for instance, the U.S. strategy adapted quickly to the reality of Iraq's lack of strength, once local commanders reported little or no resistance. This is not to say, as some have suggested, that the military has become a true bottom-up organization. The chain of command remains essential to the way the military works, and all battlefield action takes place within a framework defined by what's known as the Commander's Intent, which essentially lays out a campaign's objectives. But increasingly, successful campaigns may depend as much on the fast aggregation of information from the field as on preexisting, top-down strategies.

V

When it comes to the problems of the U.S. intelligence community before September 11, the problem was not decentralization. The problem was the *kind* of decentralization that the intelligence community was practicing. On the face of it, the division of labor between the different agencies makes a good deal of sense. Specialization allows for a more fine-grained appreciation of information and greater expertise in analysis. And everything we know about decision making suggests that the more diverse the available perspectives on a problem, the more likely it is that the final decision will be smart. Acting Defense Intelligence Agency director Lowell Jacoby suggested precisely this in written testimony before Congress, writing, "Information considered irrelevant noise by one

set of analysts may provide critical clues or reveal significant relationships when subjected to analytic scrutiny by another."

What was missing in the intelligence community, though, was any real means of aggregating not just information but also judgments. In other words, there was no mechanism to tap into the collective wisdom of National Security Agency nerds, CIA spooks, and FBI agents. There was decentralization but no aggregation, and therefore no organization. Richard Shelby's solution to the problem—creating a truly central intelligence agency—would solve the organization problem, and would make it easier for at least one agency to be in charge of all the information. But it would also forgo all the benefits—diversity, local knowledge, independence—that decentralization brings. Shelby was right that information needed to be shared. But he assumed that someone—or a small group of someones—needed to be at the center, sifting through the information, figuring out what was important and what was not. But everything we know about cognition suggests that a small group of people, no matter how intelligent, simply will not be smarter than the larger group. And the best tool for appreciating the collective significance of the information that the intelligence community had gathered was the collective wisdom of the intelligence community. Centralization is not the answer. But aggregation is.

There were and are a number of paths the intelligence community could follow to aggregate information without adopting a traditional top-down organization. To begin with, simply linking the computer databases of the various agencies would facilitate the flow of information while still allowing the agencies to retain their autonomy. Remarkably, two years after September 11, the government still did not have a single unified "watch list" that drew on data from all parts of the intelligence community. In some sense, quite simple, almost mechanical steps would have allowed the intelligence community to be significantly smarter.

Other, more far-reaching possibilities were available, too, and

in fact some within the intelligence community tried to investigate them. The most important of these, arguably, was the FutureMAP program, an abortive plan to set up decision markets—much like those of the IEM—that would have, in theory, allowed analysts from different agencies and bureaucracies to buy and sell futures contracts based on their expectations of what might happen in the Middle East and elsewhere. FutureMAP, which got its funding from the Defense Advanced Research Projects Agency (DARPA), had two elements. The first was a set of internal markets, which would have been quite small (perhaps limited to twenty or thirty people), and open only to intelligence analysts and perhaps a small number of outside experts. These markets might actually have tried to predict the probability of specific events (like, presumably, terrorist attacks), since the traders in them would have been able to rely on, among other things, classified information and hard intelligence data in reaching their conclusions. The hope was that an internal market would help circumvent the internal politics and bureaucratic wrangling that have indisputably had a negative effect on American intelligence gathering, in no small part by shaping the kinds of conclusions analysts feel comfortable reaching. In theory, at least, an internal market would have placed a premium not on keeping one's boss or one's agency happy (or on satisfying the White House) but rather on offering the most accurate forecast. And since it would have been open to people from different agencies, it might have offered the kind of collective judgment that the intelligence community has found difficult to make in the past decade.

The second part of FutureMAP was the so-called Policy Analysis Market (PAM), which in the summer of 2003 became the object of a firestorm of criticism from appalled politicians. The idea behind PAM was a simple one (and similar to the idea behind the internal markets): just as the IEM does a good job of forecasting election results and other markets seem to do a good job of forecasting the future, a market centered on the Middle East might provide intelligence that otherwise would be missed.

What distinguished PAM from the internal market was that it was going to be open to the public, and that it seemed to offer the possibility of ordinary people profiting from terrible things happening. Senators Ron Wyden and Byron Dorgan, who were the leaders of the effort to kill PAM, denounced it as "harebrained," "offensive," and "useless." The public, at least those who heard about PAM before it was unceremoniously killed, seemed equally appalled.

Given the thesis of this book, it will not surprise you to learn that I think PAM was potentially a very good idea. The fact that the market was going to be open to the public did not mean that its forecasts would be more inaccurate. On the contrary, we've seen that even when traders are not necessarily experts, their collective judgment is often remarkably good. More to the point, opening the market to the public was a way of getting people whom the American intelligence community might not normally hear from—whether because of patriotism, fear, or resentment—to offer up information they might have about conditions in the Middle East.

From the perspective of Shelby's attack on the intelligence community, PAM, like the internal markets, would have helped break down the institutional barriers that keep information from being aggregated in a single place. Again, since traders in a market have no incentive other than making the right prediction—that is, there are no bureaucratic or political factors influencing their decisions—and since they have that incentive to be right, they are more likely to offer honest evaluations instead of tailoring their opinions to fit the political climate or satisfy institutional demands.

Senator Wyden dismissed PAM as a "fairy tale" and suggested that DARPA would be better off putting its money into "real world" intelligence. But the dichotomy was a false one. No one suggested replacing traditional intelligence gathering with a market. PAM was intended to be simply another way of collecting information. And in any case, if PAM had, in fact, been a "fairy tale," we would have known it soon enough. Killing the project ensured only that

we would have no idea whether decision markets might have something to add to our current intelligence efforts.

The hostility toward PAM, in any case, had little to do with how effective it would or would not be. The real problem with it, Wyden and Dorgan made clear, was that it was "offensive" and "morally wrong" to wager on potential catastrophes. Let's admit there's something viscerally ghoulish about betting on an assassination attempt. But let's also admit that U.S. government analysts ask themselves every day the exact same questions that PAM traders would have been asking: How stable is the government of Jordan? How likely is it the House of Saud will fall? Who will be the head of the Palestinian Authority in 2005? If it isn't immoral for the U.S. government to ask these questions, it's hard to see how it's immoral for people outside the U.S. government to ask them.

Nor should we have shied from the prospect of people profiting from predicting catastrophe. CIA analysts, after all, don't volunteer their services. We pay them to predict catastrophes, as we pay informants for valuable information. Or consider our regular economy. The entire business of a life-insurance company is based on betting on when people are going to die (with a traditional life-insurance policy, the company is betting you'll die later than you think you will, while with an annuity it's betting you'll die sooner). There may be something viscerally unappealing about this, but most of us understand that it's necessary. This is, in some sense, what markets often do: harness amorality to improve the collective good. If the price of better intelligence was simply having our sensibilities bruised, that doesn't seem like too high a price to have paid. And surely letting people wager on the future was less morally problematic than many of the things our intelligence agencies have done and continue to do to get information. If PAM would actually have made America's national security stronger, it would have been morally wrong not to use it.

There were serious problems that the market would have had to overcome. Most notably, if the market was accurate, and the De-

partment of Defense acted on its predictions to stop, say, a coup in Jordan, that action would make the traders' predictions false and thereby destroy the incentives to make good predictions. A well-designed market would probably have to account for such U.S. interventions, presumably by making the wagers conditional on U.S. action (or, alternatively, traders would start to factor the possibility of U.S. action into their prices). But this would be a problem only if the market was in fact making good predictions. Had PAM ever become a fully liquid market, it would probably also have had the same problems other markets sometimes have, like bubbles and gaming. But it is not necessary to believe that markets work perfectly to believe that they work well.

More important, although most of the attention paid to PAM focused on the prospect of people betting on things like the assassination of Arafat, the vast majority of the "wagers" that PAM traders would have been making would have been on more mundane questions, such as the future economic growth of Jordan or how strong Syria's military was. At its core, PAM was not meant to tell us what Hamas was going to do next week or to stop the next September 11. Instead, it was meant to give us a better sense of the economic health, the civil stability, and the military readiness of Middle Eastern nations, with an eye on what that might mean for U.S. interests in the region. That seems like something about which the aggregated judgment of policy analysts, would-be Middle Eastern experts, and businessmen and academics from the Middle East itself (the kind of people who would likely have been trading on PAM) would have had something valuable to say.

We may yet find out if they do, because in the fall of 2003, NetExchange, the company that had been responsible for setting up PAM, announced that in 2004, a new, revised Policy Analysis Market (this one without government involvement of any sort) would be opened to the public. NetExchange was careful to make clear that the goal of the market would not be to predict terrorist incidents but rather to forecast broader economic, social, and military trends in the region. So perhaps the promise of PAM will actually get tested against reality, instead of being dismissed out of hand. It also seems plausible, and even likely, that the U.S. intelligence community will eventually return to the idea of using internal prediction markets—limited to analysts and experts—as a means of aggregating dispersed pieces of information and turning them into coherent forecasts and policy recommendations. Perhaps that would mean that the CIA would be running what Senators Wyden and Dorgan scornfully called "a betting parlor." But we know one thing about betting markets: they're very good at predicting the future.

5.

SHALL WE DANCE?: COORDINATION IN A COMPLEX WORLD

I

No one has ever paid more attention to the streets and sidewalks of New York City than William H. Whyte. In 1969, Whyte—the author of the sociological classic *The Organization Man*—got a grant to run what came to be known as the Street Life Project, and spent much of the next sixteen years simply watching what New Yorkers did as they moved through the city. Using time-lapse cameras and notebooks, Whyte and his group of young research assistants compiled a remarkable archive of material that helped explain how people used parks, how they walked on busy sidewalks, and how they handled heavy traffic. Whyte's work, which was eventually published in his book *City*, was full of fascinating ideas about architecture, urban design, and the importance to a city of keeping street life vibrant. It was also a paean to the urban pedestrian. "The pedestrian is a social being," Whyte wrote. "He is also a transportation unit, and a marvelously complex and efficient one." Pedestrians, Whyte showed, were able, even on crowded sidewalks, to move surprisingly fast without colliding with their neighbors. In fact, they were often at their best when the crowds were at their biggest. "The good pedestrian," Whyte wrote, "usually walks slightly to one side, so that he is looking over the shoulder of

the person ahead. In this position he has the maximum choice and the person ahead is in a sense running interference for him."

New Yorkers mastered arts like "the simple pass," which involved slowing ever so slightly in order to avoid a collision with an oncoming pedestrian. They platooned at crosswalks as a protection against traffic. In general, Whyte wrote, "They walk fast and they walk adroitly. They give and they take, at once aggressive and accommodating. With the subtlest of motions they signal their intentions to one another." The result was that "At eye level, the scene comes alive with movement and color—people walking quickly, walking slowly, skipping up steps, weaving in and out in crossing patterns, accelerating and retarding to match the moves of others. There is a beauty that is beguiling to watch."

What Whyte saw—and made us see—was 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. Pedestrians are constantly anticipating each other's behavior. No one tells them where or when or how to walk. Instead, they all decide for themselves what they'll do based on their best guess of what everyone else will do. And somehow it usually works out well. There is a kind of collective genius at play here.

It is, though, a different kind of genius from the one represented by the NFL point spread or Google. The problem that a crowd of pedestrians is "solving" is fundamentally different from a problem like "Who will win the Giants-Rams game, and by how much?" The pedestrian problem is an example of what are usually called *coordination problems*. Coordination problems are ubiquitous in everyday life. What time should you leave for work? Where do we want to eat tonight? How do we meet our friends? How do we allocate seats on the subway? These are all coordination problems. So, too, are many of the fundamental questions that any economic system has to answer: Who will work where? How much should my factory produce? How can we make sure that people get the goods and services they want? What defines a coordination problem is

6.

SOCIETY DOES EXIST: TAXES, TIPPING, TELEVISION, AND TRUST

I

In the summer of 2002, a great crime was perpetrated against the entire nation of Italy. Or so at least tens of millions of Italian soccer fans insisted after the country's national team was knocked out of the World Cup by upstart South Korea. The heavily favored Italians had scored an early goal against the Koreans and had clung to their 1–0 lead for most of the game, before yielding a late equalizer and then an overtime goal that sent them packing. The Italian performance had been mediocre at best. But the team was victimized by a couple of very bad officiating decisions, including one that disallowed a goal. Had those decisions gone the other way, it's likely Italy would have won.

The Italian fans, of course, blamed the referee, an Ecuadorean named Byron Moreno, for the defeat. Strikingly, though, they did not blame Moreno for being incompetent (which he was). Instead, they blamed him for being criminal. In the fans' minds, their team had been the victim of something more sinister than just bad officiating. Instead, the Italians had fallen prey to a global conspiracy—perhaps orchestrated by FIFA, soccer's governing body—designed to keep them from their just deserts. Moreno had been the point man for the conspiracy. And he had carried out his orders perfectly.

The Milan daily *Corriere della Sera,* for instance, protested against a system in which "referees . . . are used as hitmen." *La Gazzetta dello Sport* editorialized, "Italy counts for nothing in those places where they decide the results and put together million-dollar deals." A government minister declared, "It seemed as if they just sat around a table and decided to throw us out." And Francesco Totti, one of the stars of the Italian team, captured the conspiratorial mood best when he said, "This was a desired elimination. By who? I don't know—there are things greater than me but the feeling is that they wanted us out." In the weeks that followed the game, no proof of an anti-Italian cabal or of Moreno's supposed chicanery surfaced (despite the best efforts of the Italian papers). But the fans remained unwavering in their conviction that dark forces had united to destroy Italy's ambitions.

To an outside observer, the accusations of corruption seemed crazy. Honest referees make bad decisions all the time. What reason was there to believe that Moreno was any different? But to anyone familiar with Italian soccer the accusations were completely predictable. That's because in Italian soccer, corruption is assumed to be the natural state of affairs. Every year, the Italian soccer season is marred by weekly charges of criminality and skulduggery. Teams routinely claim that individual refs have been bought off, and request that particular referees not be assigned to their games. Refereeing is front-page news. Every Monday night, a TV show called *Biscardi's Trial* devotes two and a half hours to dissecting officiating mistakes and lambasting the officials for favoritism.

The effect of all this on actual Italian soccer games is not good. Although the players are among the very best in the world, the games are often halting, foul-ridden affairs repeatedly delayed by playacting, whining players more interested in working the refs than anything else. Defeat is never accepted as the outcome of a fair contest. And even victory is marred by the thought that perhaps backroom machinations were responsible for it.

So what does Italian soccer have to do with collective deci-

sion making and problem solving? Well, although the teams in a soccer game are trying to defeat each other, and therefore have competing interests, the teams also have a common interest: namely, making sure that the games are entertaining and compelling for the fans. The more interesting the games are, the more likely it is that people will come, the greater ticket sales and TV ratings will be, and the higher team profits and player salaries will be. When two soccer teams play each other, then, they're not just competing. They're also, at least in theory, working together—along with the officials—in order to produce an entertaining game. And this is precisely what the Italian teams are unable to do. Because neither side can be sure that its efforts will be fairly rewarded, the players devote an inordinate amount of time to protecting their own interests. Energy, time, and attention that would be better spent improving the quality of play instead goes into excoriating, monitoring, and trying to manipulate the referees. And the manipulation feeds on itself. Even if most players would rather be honest, they realize that they'd only be asking to be exploited. As Gennaro Gattuso, a winger for European champion AC Milan, said in October of 2003, "The system prevents you from telling the truth and being yourself." Hardly anyone likes the system the way it is, but no one can change it.

What Italian soccer is failing to do, then, is come up with a good solution to what I'll call here a *cooperation* problem. Cooperation problems often look something like coordination problems, because in both cases a good solution requires people to take into account what everyone else is doing. But if the mechanism is right, coordination problems can be solved even if each individual is single-mindedly pursuing his self-interest—in fact, in the case of price, that's what coordination seems to require. To solve cooperation problems—which include things like keeping the sidewalk free of snow, paying taxes, and curbing pollution—the members of a group or a society need to do more. They need to adopt a broader definition of self-interest than the myopic one that maximizing

profits in the short term demands. And they need to be able to trust those around them, because in the absence of trust the pursuit of myopic self-interest is the only strategy that makes sense. How does this happen? And does it make a difference when it does?

II

In September 2003, Richard Grasso, who was then the head of the New York Stock Exchange, became the first CEO in American history to get fired for making too much money. Grasso had run the NYSE since 1995, and by most accounts he had done a good job. He was aggressively self-promoting, but he did not appear to be incompetent or corrupt. But when the news broke that the NYSE was planning to give Grasso a lump-sum payment of $139.5 million—made up of retirement benefits, deferred pay, and bonuses—the public uproar was loud and immediate, and in the weeks that followed, the calls for Grasso's removal grew deafening. When the NYSE's board of directors (the very people, of course, who had agreed to pay him the $139.5 million in the first place) asked Grasso to step down, it was because the public's outrage had made it impossible to keep him around.

Why was the public so outraged? After all, they did not have to foot the bill for Grasso's millions. The NYSE was spending its own money. And complaining about Grasso's windfall didn't make anyone else any better off. He had already been paid, and the NYSE wasn't going to take the money it had promised him and give it to charity or invest it more wisely. From an economist's point of view, in fact, the public reaction seemed deeply irrational. Economists have traditionally assumed, reasonably, that human beings are basically self-interested. This means a couple of (perhaps obvious) things. First, faced with different choices (of products, services, or simply courses of action), a person will choose the one that benefits her personally. Second, her choices will not depend

on what anyone else does. But with the possible exception of business columnists, no one who expressed outrage over how much Dick Grasso made reaped any concrete benefits from their actions, making it irrational to invest time and energy complaining about him. And yet that's exactly what people did. So the question again is: Why?

The explanation for people's behavior might have something to do with an experiment called the "ultimatum game," which is perhaps the most-well-known experiment in behavioral economics. The rules of the game are simple. The experimenter pairs two people. (They can communicate with each other, but otherwise they're anonymous to each other.) They're given $10 to divide between them, according to this rule: One person (the proposer) decides, on his own, what the split should be (fifty-fifty, seventy-thirty, or whatever). He then makes a take-it-or-leave-it offer to the other person (the responder). The responder can either accept the offer, in which case both players pocket their respective shares of the cash, or reject it, in which case both players walk away empty-handed.

If both players are rational, the proposer will keep $9 for himself and offer the responder $1, and the responder will take it. After all, whatever the offer, the responder should accept it, since if he accepts he gets some money and if he rejects, he gets none. A rational proposer will realize this and therefore make a lowball offer.

In practice, though, this rarely happens. Instead, lowball offers—anything below $2—are routinely rejected. Think for a moment about what this means. People would rather have nothing than let their "partners" walk away with too much of the loot. They will give up free money to punish what they perceive as greedy or selfish behavior. And the interesting thing is that the proposers anticipate this—presumably because they know they would act the same way if they were in the responder's shoes. As a result, the proposers don't make many low offers in the first place. The most common offer in the ultimatum game, in fact, is $5.

Now, this is a long way from the "rational man" picture of human behavior. The players in the ultimatum game are not choosing what's materially best for them, and their choices are clearly completely dependent on what the other person does. People play the ultimatum game this way all across the developed world: cross-national studies of players in Japan, Russia, the United States, and France all document the same phenomenon. And increasing the size of the stakes doesn't seem to matter much either. Obviously, if the proposer were given the chance to divide $1 million, the responder wouldn't turn down $100,000 just to prove a point. But the game has been played in countries, like Indonesia, where the possible payoff was equal to three days' work, and responders still rejected lowball offers.

It isn't just humans who act this way, either. In a study that was fortuitously released the day Richard Grasso stepped down, primatologists Sarah F. Brosnan and Frans B. M. de Waal showed that female capuchin monkeys are also offended by unfair treatment. The capuchins had been trained to give Brosnan a granite pebble in exchange for food. The pay, as it were, was a slice of cucumber. The monkeys worked in pairs, and when they were both rewarded with cucumbers, they exchanged rock for food 95 percent of the time. This idyllic market economy was disrupted, though, when the scientists changed the rules, giving one capuchin a delicious grape as a reward while still giving the other a cucumber slice. Confronted with this injustice, the put-upon capuchins often refused to eat their cucumbers, and 40 percent of the time stopped trading entirely. Things only got worse when one monkey was given a grape in exchange for doing nothing at all. In that case, the other monkey often tossed away her pebble, and trades took place only 20 percent of the time. In other words, the capuchins were willing to give up cheap food—after all, a cucumber for a peb-

ble seems like a good deal—simply to express their displeasure at their comrades' unearned riches. Presumably if they'd been given the chance to stop their comrades from enjoying those riches—as the players in the ultimatum game were—the capuchins would have gladly taken it.

Capuchins and humans alike, then, seem to care whether rewards are, in some sense, "fair." That may seem like an obvious thing to worry about, but it's not. If the monkey thought a rock for a cucumber was a reasonable trade and was happy to make it before he saw his comrade get a grape, she should be happy to make the trade afterward, too. After all, her job hasn't gotten any harder, nor is the cucumber any less tasty. (Or if it is, that's because she's obsessed with what her neighbor's getting.) So her feelings about the deal should stay the same. Similarly, the responders in the ultimatum game are being offered money for what amounts to a few minutes of "work," which mostly consists of answering "yes" or "no." Turning down free money is not something that, in most circumstances, makes sense. But people are willing to do it in order to make sure that the distribution of resources is fair.

Does this mean people think that, in an ideal world, everyone would have the same amount of money? No. It means people think that, in an ideal world, everyone would end up with the amount of money they deserved. In the original version of the ultimatum game, only luck determines who gets to be the proposer and who gets to be the responder. So the split, people feel, should be fairly equal. But people's behavior in the game changes quite dramatically when the rules are changed. In the most interesting version of the ultimatum game, for instance, instead of assigning the proposer role randomly, the researchers made it seem as if the proposers had earned their positions by doing better on a test. In those experiments, proposers offered significantly less money, yet not a single offer was rejected. People apparently thought that a proposer who merited his position deserved to keep more of the wealth.

Put simply, people (and capuchins) want there to be a reasonable relationship between accomplishment and reward. That's what was missing in Grasso's case. He was getting too much for having done too little. Grasso seems to have been good at his job. But he was not irreplaceable: no one thought the NYSE would fall apart once he was gone. More to the point, the job was not a $140 million job. (What job is?) In terms of complexity and sophistication, it bore no resemblance to, say, running Merrill Lynch or Goldman Sachs. Yet Grasso was being paid as much as many Wall Street CEOs, who are themselves heftily overcompensated.

The impulse toward fairness that drove Grasso from office is a cross-cultural reality, but culture does have a major effect on what counts as fair. American CEOs, for instance, make significantly more money than European or Japanese CEOs, and salary packages that would send the Germans to the barricades barely merit a moment's notice in the United States. More generally, high incomes by themselves don't seem to bother Americans much—even though America has the most unequal distribution of income in the developed world, polls consistently show that Americans care much less about inequality than Europeans do. In fact, a 2001 study by economists Alberto Alesina, Rafael di Tella, and Robert MacCulloch found that in America the people whom inequality bothers most are the rich. One reason for this is that Americans are far more likely to believe that wealth is the result of initiative and skill, while Europeans are far more likely to attribute it to luck. Americans still think, perhaps inaccurately, of the United States as a relatively mobile society, in which it's possible for a working-class kid to become rich. The irony is that Grasso himself was a working-class kid who made good. But even for Americans, apparently, there is a limit to how good you can make it.

There's no doubt the indignation at Grasso's retirement package was, in an economic sense, irrational. But like the behavior of the ultimatum game responders, the indignation was an example of

what economists Samuel Bowles and Herbert Gintis call "strong reciprocity," which is the willingness to punish bad behavior (and reward good behavior) even when you get no personal material benefits from doing so. And, irrational or not, strong reciprocity is, as Bowles and Gintis term it, a "prosocial behavior," because it pushes people to transcend a narrow definition of self-interest and do things, intentionally or not, that end up serving the common good. Strong reciprocators are not altruists. They are not rejecting lowball offers, or hounding Dick Grasso, because they love humanity. They're rejecting lowball offers because the offers violate their individual sense of what a just exchange would be. But the effect is the same as if they loved humanity: the group benefits. Strong reciprocity works. Offers in the ultimatum game are usually quite equitable, which is what they should be given the way the resources are initially set up. And whenever the NYSE thinks about hiring a CEO, it will presumably be more rigorous in figuring out how much he's actually worth. Individually irrational acts, in other words, can produce a collectively rational outcome.

III

The mystery that the idea of prosocial behavior may help resolve is the mystery of why we cooperate at all. Societies and organizations work only if people cooperate. It's impossible for a society to rely on law alone to make sure citizens act honestly and responsibly. And it's impossible for any organization to rely on contracts alone to make sure that its managers and workers live up to their obligations. So cooperation typically makes everyone better off. But for each individual, it's rarely rational to cooperate. It always makes more sense to look after your own interests first and then live off everyone else's work if they are silly enough to cooperate. So why don't most of us do just that?

The classic and canonical explanation of why people cooperate was offered by political scientist Robert Axelrod, who argued in the 1980s that cooperation is the result of repeated interactions with the same people. As Axelrod put it in his classic *The Evolution of Cooperation*, "The foundation of cooperation is not really trust, but the durability of the relationship . . . Whether the players trust each other or not is less important in the long run than whether the conditions are ripe for them to build a stable pattern of cooperation with each other." People who repeatedly deal with each other over time recognize the benefits of cooperation, and they do not try to take advantage of each other, because they know if they do, the other person will be able to punish them. The key to cooperation is what Axelrod called "the shadow of the future." The promise of our continued interaction keeps us in line. Successful cooperation, Axelrod argued, required that people start off by being nice—that is, by being willing to cooperate—but that they had to be willing to punish noncooperative behavior as soon as it appeared. The best approach was to be "nice, forgiving, and retaliatory."

Those rules seem completely sensible, and are probably a good description of the way most people in a well-functioning society deal with those they know. But there's something unsatisfying, as Axelrod himself now seems to recognize, about the idea that cooperation is simply the product of repeated interactions with the same people. After all, we often act in a prosocial fashion even when there is no obvious payoff for ourselves. Look at the ultimatum game again. It is a one-shot game. You don't play it with the same person more than once. The responders who turned down lowball offers were therefore not doing so in order to teach the proposer to treat them better. And yet they still punished those who they thought were acting unfairly, which suggests that the "shadow of the future" alone cannot explain why we cooperate.

The interesting thing, ultimately, isn't that we cooperate with those we know and do business with regularly. The interesting thing is that we cooperate with strangers. We donate to charities. We buy things off eBay sight unseen. People sign on to Kazaa and

AFTERWORD

TO THE ANCHOR BOOKS EDITION

I

In the months after *The Wisdom of Crowds* came out, I spent a surprising amount of time running experiments that were a lot like the ox-weighing contest that Francis Galton stumbled upon back in 1906. On a regular basis, when I would be interviewed on radio or television, or when I'd speak to a group, I'd be asked to manage a contest that could demonstrate collective wisdom in action. So I went up to random people in Times Square and asked them how many jelly beans were in a plastic jar. In Grand Central Station, a group of tourists who had come to New York to embark on the *Queen Mary II* was asked to guess my weight. And the audience for the Diane Rehm radio show was asked how many books I had in my study—a room that I described on the air, but that no one in the audience had, obviously, ever seen.

I was glad to do all these things, because they were fun (well, approaching people in Times Square was daunting at first), and because they made the argument of the book seem more concrete. But when I first started these experiments, I realized something odd: each time, I had a flash of uncertainty that this time something was going to go wrong, and that the crowd's guess would be way off the mark.

Things never did go wrong. Each time, the crowd did just as expected: its collective guess was very accurate, and was better than the vast majority of individual guesses. And this was true even though in some cases the groups were relatively small, and even though sometimes the problems seemed close to impossible. (The Diane Rehm show listeners, for instance, did a remarkably good job of predicting the number of books in my study, when even I didn't know how many books were in there.) And over time, I got used to this, and those flashes of doubt came less frequently. But it's telling, I think, that they were there in the first place. Even after I'd written an entire book about collective wisdom, there was still a part of me that instinctively questioned whether it would really work.

This story is, I think, a testament to how uncanny—and counterintuitive—the wisdom of crowds really is. For most of us, our initial reaction to the idea of relying on the collective judgment of a diverse group of people is that it will not work. The well-informed people will be outweighed by the poorly informed, and the group's decision will be worse than that of even the average individual. While we may all nod at the slogan that "no one of us knows more than all of us," we're more likely to believe what Tommy Lee Jones's character said in *Men in Black*: "A person is smart. People are dumb." I wrote *The Wisdom of Crowds* in part to explain why this idea is wrong. But I know that it seems right.

That's why what I've been most surprised by is not the tenaciousness of the resistance to the idea that crowds might be wise, but rather the speed with which that idea is being embraced and put to good use. During the presidential election of 2004, for instance, there was an enormous amount of attention paid to the election forecasts that bettors were making on the Iowa Electronic Markets (IEM) and TradeSports.com. In previous elections, these markets were treated more as curiosities—intriguing experiments perhaps, but not necessarily anything that needed to be taken too seriously. In 2004, though, they were a constant part of election

discussions on the Internet, and they also made regular appearances in the mainstream media. When major events in the presidential race occurred, the first reaction of many was to see how these markets had reacted. The IEM and TradeSports were also looked to for reality checks when the polls seemed unusually volatile. And on the whole, people's confidence in these markets was borne out by their performance: the IEM's election-eve forecast of the vote share for both Bush and Kerry was very close to the actual outcome, while TradeSports did an exceptionally good job of predicting which states each candidate would carry. (On election day itself, though, these markets did poorly—they, like all the pundits, were deceived by the exit polls predicting victory for Kerry.)

The virtues of collective wisdom have also started to make inroads into corporate America, where more companies have begun to experiment with using internal prediction markets to tap the intelligence of their employees. When both *Time* magazine and *U.S. News & World Report* did big stories on information markets and corporations, it seemed to be evidence that things were changing. Similarly, various government agencies have begun trying out different forms of collective information gathering and decision making, and most strikingly, at least some people in the U.S. intelligence community have been discussing how collective mechanisms might help circumvent the well-documented problems that have hindered U.S. intelligence in recent years.

The growing interest in collective wisdom is the product of a host of different factors, but I think in many ways it's directly connected to the increased importance of the Internet. In part, that's because I think the ethos of the Net is fundamentally respectful of and invested in the idea of collective wisdom, and in some sense hostile to the idea that power and authority should belong to a select few. Many of the Net's most distinctive landmarks—Google, Slashdot, Wikipedia—are the products of the wisdom of crowds, and more generally, the Net, almost by its very structure, seems

antihierarchical. It provides a vivid demonstration every day that systems can work smoothly and intelligently without traditional hierarchies and without having any one person in charge.

Similarly important is the fact that the Internet simply makes it much easier to aggregate information from many different sources than ever before. Historically, one of the biggest objections to collective decision making was that it was necessarily inefficient, that once a group got beyond a certain size it would simply take too long and be too difficult to give members of the group a genuine say in problem solving. That no longer seems like a plausible objection (if it ever was), since it's now possible to solicit and aggregate information from people all over the world and arrive at a collective decision with a few clicks of a mouse. For instance, NASA recently concluded an online experiment that it called Clickworkers, which was designed to test whether the collective judgment of ordinary people would be of any use finding and classifying craters on Mars. You could go to the site, undergo a short training session, and then click away. The result, not surprisingly, was that the clickworkers' collective judgment was "virtually indistinguishable" from that of "a geologist with years of experience in identifying Mars craters."

For me, one of the key lessons of *The Wisdom of Crowds* is that we don't always know where good information is. That's why, in general, it's smarter to cast as wide a net as possible, rather than wasting time figuring out who should be in the group and who should not. This, too, is an idea that is well-suited to the Internet, since one of the common experiences on the Net is coming across information in a place that you never expected to find it. Instead of going only to well-established information filters, we recognize that the best information may be in nontraditional locations (and we try to find it using Google, relying there too on the wisdom of crowds).

The point is not that you need the Internet in order to take advantage of collective wisdom. Betting pools and the stock market and jelly bean–guessing contests existed long before the Net, and so therefore did the wisdom of crowds. But I do think that the ubiquity of Internet access and information technology, and the way in which people are now simply more connected (technologically if not socially) than ever before, have been important in making the virtues of collective wisdom seem less like a fantasy and more like a reality. Twenty-five years ago a company that wanted to open up its decision-making process to employees in different parts of the world and different parts of the organization would have had a very hard time doing so. Today, there are few practical hurdles. The only real obstacles are mental ones.

II

Perhaps the most important of those obstacles, of course, is the assumption that I alluded to at the beginning of the afterword, namely that you're better off putting your trust in one smart person rather than in many people. And certainly one of the questions that I was asked most after the book came out was, "Do you really think that a group of people off the street can know more about a complex question than an expert?"

It's an interesting question, but in some sense I think it doesn't really get to the important point. The truth is that *The Wisdom of Crowds* was never intended as a defense of laypeople against experts. (As I've said in the book, I assume that in most cases, "crowds" will include experts as well as amateurs.) Expertise is valuable; smart people are valuable. The more information a group has, the better its collective judgment will be, so you want as many people with good information in a group as possible.

Recognizing that expertise is valuable, though, is not the same as saying you should rely on a single expert to make decisions. *The Wisdom of Crowds* is not an argument against experts, but against our excessive faith in the single individual decision maker. I think there are two big problems with relying on a single

individual—no matter how well-informed. The first is that true experts—that is, the real titans—are surprisingly hard to identify. In fact, if a group is smart enough to know whether an individual is a genuine decision-making prodigy, then the group is smart enough not to need that individual.

The second, and more important, problem is that even brilliant experts have biases and blind spots, and so they make mistakes. And what's troubling is that, in general, they don't know when they're making those mistakes. Some of the most surprising material in *The Wisdom of Crowds* shows that expert judgments are very poorly calibrated—which means that there's little correlation between an expert's confidence in his judgment and the accuracy of it. In other words, experts don't know when they don't know something. That's why it's worthwhile to cast a wider net, and why relying on a crowd of decision makers improves (though it doesn't guarantee) your chances of reaching a good decision. Relying on a crowd rather than an individual improves your chances of finding information that you didn't know was out there. Just as important, it minimizes the impact of the mistakes that individuals make.

To me, that's one of the (and maybe the) great virtues of collective decision making: it doesn't matter when an individual makes a mistake. As long as the group is diverse and independent enough, the errors people make effectively cancel themselves out, leaving you with the knowledge that the group has. Now, I realize that to some people (who have told me so) this sounds either vaguely mystical or else overly simplified. But it just happens to be the way the world works. It's easy to imagine a different world, in which people all made the same mistakes, and therefore putting all their judgments together made them dumber, not smarter. (It's easy to imagine this world because sometimes—think of the stock-market bubble of the late nineties—it's the world we live in.) But if you're careful to keep the group diverse, and careful to prevent people from influencing one another too much, the individual mis-

takes people make will be irrelevant. And their collective judgment will be wise.

III

The question about experts versus groups is, of course, a question about the fundamental premise of *The Wisdom of Crowds*. But I've also been asked a lot of questions about the best way to tap into collective wisdom, and about what's required to make the wisdom of crowds work. The two most important ones are 1) does there need to be a reward for people to do their best? and 2) do you need, paradoxically, a leader who recognizes the wisdom of crowds to be able to tap that wisdom?

The first question is about incentives. Incentives almost certainly help, if only because the prospect of a gain (or a loss) helps concentrate people's minds, and for situations in which relevant information might not be obvious—it takes a little digging to uncover it—they can be very useful. But what's interesting is that those incentives don't need to be financial. I suggested this earlier, but we now have a much better sense that it's true, thanks to a remarkable study of NewsFutures.com and TradeSports.com, the two most important online trading sites.

As I mentioned in Chapter 1, both of these sites offer markets in a host of real-world events. Before the invasion of Iraq, for instance, you could buy and sell contracts at NewsFutures based on how likely it was that Saddam Hussein would still be in power by the end of April 2003. And during the presidential election of 2004, TradeSports offered contracts that allowed you to bet on the performance of George W. Bush and John Kerry in every state in the union, as well as a package of combination bets connected to whether Osama bin Laden would be caught before election day. But among the most popular markets on these sites are, predictably

enough, sports betting, and in particular, betting on NFL games, where the market is effectively trying to forecast the final point spread.

While NewsFutures and TradeSports are similar in many ways, they are different from each other in one crucial way: NewsFutures is a play-money market, while at TradeSports people are gambling with real cash. So a natural experiment suggested itself. The study's authors—Emile Servan-Schreiber, Justin Wolfers, David Pennock, and Brian Galebach—compared the performance of the two markets over the course of the 2003 NFL season. Not surprisingly, they found that the collective judgment of bettors in both markets was exceptionally good (and better than that of almost all the experts, for instance, in another online game). More important, though, they found that the predictions of the play-money market were as accurate as those of the real-money market. Apparently, the esteem that people in NewsFutures get from performing well against the competition was enough of an incentive.

The fact that financial incentives aren't necessary to reap the benefits of the wisdom of crowds is potentially very important in thinking about how to use this notion within organizations. One of the real challenges in putting things like internal markets in place inside, say, a corporation is figuring out exactly how rewards should work. You want people to care enough about the market to be serious in their forecasts, but not to care so much that they neglect their real work. Nor do you want the financial stakes in the market to be such that people are at risk of losing enormous amounts if their forecasts turn out to be wrong. (The higher the stakes, the greater the incentive for manipulation, too.) What the NewsFutures study suggests is that you don't need high stakes in order to get good results.

The second question is about the role of leaders in transforming organizations to embrace the idea of collective wisdom. In other words, do you need a figure like John Craven—the man who put together the team that found the *Scorpion* submarine—to take

charge before the crowd can actually become wise? This is an interesting problem, because it touches on the curious fact that groups tend not be wise about their own wisdom. And certainly within most established organizations, it has historically been unusual for change to bubble up from below on its own. So it is, in fact, more likely that someone will take it on himself to champion the idea of collective wisdom, and in that way create the conditions that allow it to flourish. This is paradoxical, but no more so than the fact that an individual, not a crowd, wrote *The Wisdom of Crowds*.

In any case, while it's certainly true that you often need a smart individual to recognize the intelligence of the group, in the future that may no longer be as necessary. As the value of collective wisdom becomes more widely recognized, people will be more likely to adopt, on their own, collective approaches to problem solving, and the Internet affords us any number of examples of wise crowds that are, for the most part, self-organized and self-managed. We're a long way from anything resembling bottom-up decision making, either in government or in corporate America, but certainly the potential for it now exists.

Will that potential be realized? Well, as it happens, that's the most important question I still have about the idea of collective wisdom. While I've been impressed by the way in which organizations have started to experiment with collective decision making, I have also been struck by how profoundly the wisdom of crowds challenges some of our most deeply held assumptions about leadership, power, and authority. Those assumptions aren't simply, or even primarily, about the value of expertise. Instead they're about the notion that power ultimately has to reside in a single place—a single person—if it's really to work. We want there to be one person we can point to and say, "He made the decision," and we fear that if we don't have that, nothing will get done.

There are obviously some cases, and some situations, in which this is true. But I'm more convinced than ever that decisions do

not have to be made in this way, and that collective judgments can be every bit as effective and authoritative as individual ones. To take a simple example, at the race track, every twelve minutes or so, a crowd of people collectively makes a judgment about the chances of every horse in a race. That judgment—that decision—determines how much each individual bettor wins. The crowd's judgment is, as it happens, remarkably intelligent, but it's also accepted as final by everyone in the group. We don't require a single individual to set the odds or render a final verdict on who gets paid what. The same is true, on a much bigger scale, of the stock and bond markets, which effectively render verdicts every day that have enormous financial and social consequences. If those decisions can be made collectively, it's hard to see why other decisions—like where to build a factory, or where to assign intelligence agents—can't be as well.

I don't want to underplay the power of the assumptions that underlie our traditional model of decision making. But while trusting the collective judgment of a group may be difficult, it's also smart. What I think we know now is that in the long run, the crowd's judgment is going to give us the best chance of making the right decision, and in the face of that knowledge, traditional notions of power and leadership should begin to pale. I am cautiously hopeful that they will, allowing us to begin to trust individual leaders less and ourselves more.

ACKNOWLEDGMENTS

This book began, in some sense, with a column I wrote for *The New Yorker* not long after I started there, in the spring of 2000. And much of what the book became was influenced by columns I've written for the magazine since. In more ways than one, if David Remnick hadn't taken a chance on hiring me, and a chance on publishing a business column in *The New Yorker*, this book would not exist. I owe him a lot. Thanks, David.

Henry Finder of *The New Yorker* edited the first pieces I wrote for the magazine, shepherded me through my first year as a columnist, and has remained a constant source of intellectual inspiration. Henry has the rare gift of making a writer feel as if what he does is important—not to the world, necessarily, but just in itself. This is a good thing for anyone to feel about what they do.

I want to thank Nick Paumgarten, who edits my columns, for putting so much time and energy into my work. Nick's made me a better writer and, I think, this book a better book than it otherwise would have been.

Michael Mauboussin, chief investment strategist at Credit Suisse First Boston, first got me thinking about the ideas that animate this book with his remarkable writings on the stock market and complex systems. Michael's work remains a touchstone for me, and my conversations with him over the last year have made me