

The promise of **prediction markets:** A roundtable

Although they draw together widely dispersed information, prediction markets face organizational and legal challenges.

Renée Dye

Every senior executive knows that business decisions are seldom better than the information behind them. Yet although it is usually lower-level employees who interact directly with the customer, decision makers rarely ask them how, for example, new products will fare. Leaders therefore deprive themselves of information that could enrich their analysis and reduce the risk of ivory tower decision making.

Some executives understand that valuable information lies scattered around the organization but don't know how to retrieve it. Others don't even try, perhaps for hierarchical reasons or because they suspect they might get answers colored by the desire to second their real or assumed viewpoint.

Prediction markets might solve these problems. Initially a field of research, true prediction markets in essence are small-scale electronic markets, frequently open to any employee, that tie payoffs to measurable future events, such as sales data for a computer workstation, the number of bugs in an application, or a product's usage patterns.¹ Some companies, particularly in the high-tech sector, have adopted them in earnest, and a few major companies elsewhere are experimenting with them.

¹ Prediction markets can also assume simpler forms. E-mail surveys tracked through spreadsheets enable a company to collect and average single-point estimates from people throughout the organization. Surveys like these have been shown to outperform expert forecasts significantly.

At the table:



Bo Cowgill, a product manager at Google, has managed the company's prediction markets for two and a half years. He is also a writer and technologist who spends much of his time researching the validity and efficacy of prediction markets.



Todd Henderson, an assistant professor at the University of Chicago Law School, was an engagement manager at McKinsey (2001–04), where he specialized in telecommunications and high tech. In addition to conducting research on prediction markets, he wrote "Prediction markets for corporate governance," with coauthor Michael Abramowicz, for the *Notre Dame Law Review*.

These markets yield prices on prediction contracts—prices that can be interpreted as market-aggregated forecasts. Their “collective wisdom” is usually at least as accurate as expert opinion. Proponents say that prediction markets work by rapidly aggregating information dispersed across an organization while freeing participants from constraints: for instance, employees can share unwelcome information about a project’s launch date or a new product’s performance anonymously, without fear their careers might suffer. What’s more, advocates say, competition among colleagues and the prospect of winning a prize create incentives for seeking information and making the best-informed bets.

To assess the potential of prediction markets and the organizational and legal challenges they must surmount to become a more widely used business tool, a roundtable was convened at a recent McKinsey conference in Dubai. The panelists were Bo Cowgill, who manages Google’s prediction markets; Todd Henderson, an assistant professor at the University of Chicago Law School; Jeff Severts, general manager of Geek Squad, the



Jeff Severts is the vice president and general manager of Geek Squad, the services arm of the US consumer electronics retailer Best Buy. His work also involves forecasting models, prediction markets, and alternative-capital-allocation processes.



James Surowiecki is a staff writer at the *New Yorker*, where he writes a business column. He is the author of *The Wisdom of Crowds*, a national bestseller about the virtues of collective intelligence, including prediction markets.

services arm of US electronics retailer Best Buy; and James Surowiecki, author of *The Wisdom of Crowds*, a book about prediction markets and other forms of collective intelligence. Renée Dye, a consultant based in McKinsey's Atlanta office, moderated the discussion. What follows is an edited and abridged version of it.

The Quarterly: *Let's start the discussion by talking about the basic premise of prediction markets and why they often seem to work so well.*

James Surowiecki: The premise is that under the right circumstances, the collective judgment of a large group of people will generally provide a better picture of what the future might look like than anything one expert or even a small group of experts will come up with.

In most organizations, there's a lot of knowledge that the people who are making the decisions don't have access to. They may not know that it's out there or they don't know whom to ask. Even if they do, the respondents

might not want to disclose this knowledge, because they're worried about what the boss would say. Prediction markets are a way to get at, and aggregate, information dispersed across the organization. They work much like a futures market, in which the price of a contract reflects the collective day-to-day judgment either on a straight number—for instance, what level sales will reach over a certain period—or a probability—for example, the likelihood, measured as a percentage, that a product will hit a certain milestone by a certain date.

Jeff Severts: Our first experiments at Best Buy were inspired by James's book, and the results suggest that even a rudimentary survey strips away the filters that typically distort information as it moves higher in an organization.


At the time, I was managing the gift card business, which is a relatively small part of our portfolio, but I had a particular interest in it. We sent e-mails to hundreds of people throughout the company and asked them what they thought our gift card sales would be in February 2005. The only information we gave them as the context for their predictions was simple, readily available data. We got some 190 responses and ran a simple average. It turned out to be 99.5 percent accurate, whereas the people who got paid to forecast this were five percentage points off.

We ran a similar experiment later that year, when 350 random people predicted our holiday sales. Once again, the nonexperts, off by just one-tenth of 1 percent, were more accurate than the experts, who were off by 7 percent. The participants were surprised by the outcome when we shared it with them well after the actual results were in and reported. These early experiments encouraged us to get into prediction contracts, and we have to date seen over 2,000 traders make a total of 70,000 trades on 147 contracts.

Todd Henderson: It's important to note, of course, that a real prediction market, as opposed to the survey Jeff has described, adds a powerful dimension by allowing not just one but many continuous bets until the market closes. This creates a constant feedback loop of updated information.

The Quarterly: *Are there crowds that are wise and crowds that are not? Are crowds always wiser than experts?*

James Surowiecki: *The Wisdom of Crowds* is not an argument against experts. It is saying that you shouldn't rely wholly on the judgment of one person or even a very small group of people. But for a crowd to be



smart, it needs to satisfy certain criteria. It needs to be diverse, so that people are bringing different pieces of information to the table. It needs to be decentralized, so that no one at the top is dictating the crowd's answer. It needs to summarize people's opinions into one collective verdict. And the people in the crowd need to be independent, so that they pay attention mostly to their own information and don't worry about what everyone around them thinks.

The Quarterly: *Google appears to be into prediction markets in a relatively big way. What are the main applications, Bo?*

Bo Cowgill: We launched our prediction markets in April 2005, and since then we've asked about 275 different questions, and there've been some 80,000 trades. Around one-quarter of our markets have to do with demand forecasting—for instance, “How many people will use Gmail in the next three months?” Almost all Google products have had, or still have, a prediction market about their usage. Another 30 percent concern the company's performance—for example, will project deadlines be met? A smaller category concerns things that could happen in our industry, such as mergers and acquisitions that might impact Google significantly.

About 20 percent of our markets are about fun things like who will win the World Series. The fun markets tend to draw people in and make them familiar with how to trade in a prediction market. That's good, as many of these participants go on to provide liquidity in the serious markets.

Besides getting good answers to the questions we ask, we really try to use these markets to understand how our organization works. For example, we have been exploring the cognitive biases in different parts of the company and the way information moves inside it through different types of networks.

The Quarterly: *Which are the most common biases?*

Bo Cowgill: There were a few interesting ones. I did some work with two outside economists—Eric Zitzewitz and Justin Wolfers—to study them. Overall, the market performed quite well.

One important bias was optimism. Outcomes that would be good for Google—such as getting lots of users—were slightly overpriced. The market gave them a higher probability than it should have. The cause seems to be new employees, whose trades show that they are highly optimistic about our company. The external Google stock price also seems to play a role. When Google stock does well, the price of optimistic outcomes in the prediction markets also rises. People feel excited about the company when the stock performs well, so they're more likely to bet that good things will happen to Google.

We also noticed that traders underpriced extreme events, both good and bad. When we floated contracts with five different outcomes—for example, forecasts about the number of Gmail users—the highest and the lowest outcome happened more often than the market expected.

The trading yielded some useful insights about how information moves around the company. Our markets showed that beliefs are clustered, and these clusters are made up of individuals who physically sit and work close to each other, not only at the level of city and country, but at the micro-level of the office floor, measured in feet or meters between desks. Clusters also form around working together, socializing outside of work, and speaking a common language, even when this doesn't involve sitting close by. But these things seemed to matter less than geography.

Jeff Severts: At Best Buy, we've particularly noted a bias toward underestimating the ability of competitors or to think you know more about them than you do. Our prediction markets have not had a very respectable accuracy on anything related to our main competitor.

The Quarterly: *Do crowds learn over time?*

Bo Cowgill: Our experiences suggest that they do. The longer you work at Google and the longer you trade in the prediction markets, the more calibrated you become and the likelier you are to have a successful trading record. One important caveat is that being higher in the company, measured by distance from the CEO, actually seemed to place you at a disadvantage for trading profitably.

The market as a whole also got smarter. The trades we observed two and a half years after launch were better calibrated than the ones in the beginning, and we could observe a steady improvement over time.

The Quarterly: *We've talked about cognitive biases that may affect these markets. Do they have other shortcomings or limitations?*

James Surowiecki: One open question is whether they are good at forecasting genuinely discontinuous innovations or leaps. Or do they work very well only when the range of variables you are basing a forecast on is relatively well defined? We don't really have a good answer yet, because few experiments have been conducted to test this question.

One shortcoming is that a lot of people inside organizations don't find the market mechanism intuitive or easily understood. They find it very challenging to use, which limits the pool of people who participate.

Todd Henderson: They are becoming easier to use now that there are experts who can set up prediction markets for you and make the interface as accessible as possible. Some interesting findings are emerging. One is that information can reside in places that decision makers never would have thought possible—for instance, in what academics call “hidden profiles.” These are people within an organization who, because of their personality or position in the hierarchy, won't have the incentive or wherewithal to reveal information. An easily accessible prediction market provides a way for them to do so.

*'Prediction markets now have a track record in several fields, which has helped convince people that they can be useful in **real-world situations**'*

As to the shortcomings, prediction markets aren't a crystal ball. Sometimes, there isn't any dispersed information to aggregate, and in that case they won't really help.

The Quarterly: *Is there a saturation point when people in an organization get tired of prediction markets?*

Jeff Severts: You always have to be marketing them, just like everything else. This week, we had three big teach-outs at Best Buy where we taught employees how to use prediction markets. Every quarter, you have to refresh your list of prizes and try to come up with something at least as compelling as the last time. But these prizes are pretty modest. Generally, it's the innate sense of competition most people have that drives the markets.

Bo Cowgill: Actually, on a number of occasions, I've forgotten to pay out the small cash prizes we have at Google, and nobody noticed. But everyone notices when the T-shirts that show who won don't come. As for keeping up interest, there has been a slowdown in the rate of growth, but participation levels haven't dropped. We e-mail all new employees and introduce our prediction markets, explain how they work and the incentives for trading, and encourage people to take part, which many do. This

seems to offset any gradual decline in interest among existing participants. However, we typically have a lot of new employees joining Google, which is to the advantage of prediction markets, so I don't know how scalable this participation scheme would be at many other companies.

The Quarterly: *What factors have contributed to making prediction markets more popular now?*

James Surowiecki: There's been a confluence of things coming together. Prediction markets now have a track record in several fields, which has helped convince people that they can be useful in real-world situations. Technology has made a big difference, since you can now use intranets to cheaply and efficiently aggregate information across a big organization. Also, although this is rarely mentioned, the prediction markets trend is really part of a broader Web 2.0 bottom-up movement. There's increasing recognition that large groups of people can solve problems together and come up with interesting answers, and that you don't necessarily need formal hierarchies to accomplish this.

The Quarterly: *Are prediction markets getting the traction they deserve?*

Todd Henderson: The notion that markets are the best aggregators of dispersed information is not new. After all, the price of a stock is just a prediction market about that company's future cash flows. The puzzle is really why not every company is using these markets. One reason may be internal resistance high in the hierarchy, even though prediction markets are not meant to replace executives and experts. They complement senior advisors, market research, and surveys—they're simply another tool in the toolkit.

James Surowiecki: There's a tough organizational challenge in how you push this stuff through, because the concept behind prediction markets runs deeply counter to some of our basic assumptions about expertise, knowledge, and also about power and the way organizations should work. Although it is just about getting access to better information, many people have a hard time believing, despite the evidence, that it can really provide good answers consistently. That's made it harder for the idea to take off.

The Quarterly: *What does it take to implement prediction markets?*

Jeff Severts: Corporations have a taboo against even considering the possibility that an important initiative may fail. To issue a contract that implies that this could happen is to betray the company in some way.

Important decisions in building prediction markets

Define the variable you are trying to forecast

- Express variables in a precise, intuitive unit (such as “2nd-quarter revenue, in euros, for new product X”) to avoid confusion among participants.
- Give the market a relatively short time duration to keep it interesting and boost participation.

Decide how comfortable you are sharing the results

- Be prepared for management embarrassment (“everyone thinks we should shelve our new product launch”).
- Consider legal issues (“nobody in the company thinks we will meet the earnings targets our stock price implies”).

Decide who should participate

- Markets involving only internal participants are easiest to organize, though adding external participants can help companies achieve the law of large numbers. Information sharing challenges exist for both internal and external participants.
- Front-line employees often are the most active and excited participants.

Decide on the nature of the market

- Markets with real-time buying and selling of contracts yield rich, continuous results but require large numbers of participants, some of whom may need training.
- Simple surveys and other single-point forecast mechanisms are easier to administer. Companies getting started may want to proceed gradually through a series of increasingly sophisticated experiments.

Decide on incentives

- Cash prizes boost participation but run the risk of looking like “internal betting pools”—where employees can bet on (or against) the company—which could cause legal problems.
- Another option is combining symbolic incentives, such as public recognition of strong forecasters, with prizes such as gift cards.

Decide on the role of experts

- Departments dedicated to forecasting will see the establishment of a prediction market as a threat.
- A key challenge for companies using prediction markets is shifting the mind-sets of experts about their roles, from “knowing all the answers” to asking the right questions to analyzing the answers in creative ways and using them to guide decision making.

Renée Dye is a consultant in McKinsey’s Atlanta office.

So we found that support from very senior executives is essential if you want to issue contracts on anything that might be controversial. “Air cover” is a must or you’ll find yourself trading on what kind of casserole we’re having in the cafeteria on Thursday.

Todd Henderson: Setting up a prediction market isn’t as easy as snapping your fingers. There are many potential issues to wrestle with. What if an employee working on a team whose effort is closely related to what the prediction market is trying to forecast has important exclusive information? You could imagine situations where people would have perverse incentives to go around the team effort in order to personally profit by betting. Such scenarios raise questions about how to structure the trading. Should this team be allowed to trade? The individuals on the team? And when?

Jeff Severts: You can control for this risk by limiting the market to “play money” and modest prizes. Employees would not sacrifice the other, larger benefits of their jobs, such as their salary and stock options, to earn a small prize in a game.

Related articles on mckinseyquarterly.com

- How to improve strategic planning
- Distortions and deceptions in strategic decisions
- Beating the odds in market entry



The Quarterly: *What role will prediction markets play in most major corporations five years from now?*

Bo Cowgill: I would be cautious and say, probably not that much of a role,

because of the perceived threats to existing structures. And that’s a shame, because I think the information coming from prediction markets is good, and there are so many ways, some not even tested, you can use them to understand your organization better.

Todd Henderson: There are legal obstacles, too. Lots of US companies I’ve talked to worry about what the Securities and Exchange Commission would say if prediction markets start pushing information way down the hierarchy. Take the employee who sees a prediction market price on her dashboard and realizes, with some degree of confidence, that a certain drug is going to be a success. Is it illegal if she trades on this information in the real stock market? Is she an insider because she now has information that only a few top people had before? What kind of disclosure obligations does that put on a US public company? Gambling laws are another issue. Should prediction markets be viewed as an unregulated form of betting? These are enormous question marks for US public companies, and until

they are straightened out I agree with Bo's assessment that the impact will be modest.

Jeff Severts: We don't agree. Smartly applied, this tool can help management listen to voices, throughout the company, that otherwise go unheard. While private companies have an advantage, large public companies with distributed, customer-facing operations, such as Best Buy, are great candidates for exploration. We can't explore the topics that are reserved for official announcements to investors, but we can improve the information flow around countless smaller issues, and that may help managers make decisions more quickly and accurately. We're definitely constrained within the current ambiguous regulatory environment, but we are optimistic that we can find a lot of value here.

James Surowiecki: I wouldn't be surprised to see prediction markets used in many more companies than today, not least as a tool to forecast sales. Consumer-facing companies should be particularly interested. Technology companies are also natural candidates because they tend to be instinctively interested in new ideas. **Q**

Renée Dye is a consultant in McKinsey's Atlanta office. Copyright © 2008 McKinsey & Company. All rights reserved.

We welcome your comments on this article.
Please send them to quarterly_comments@mckinsey.com.