



MARKET FORECASTS

Prediction markets can be uncannily accurate — sometimes. Scientists have begun to understand why they work, and how they can fail.

BY ADAM MANN

It was a great way to mix science with gambling, says Anna Dreber. The year was 2012, and an international group of psychologists had just launched the ‘Reproducibility Project’ — an effort to repeat dozens of psychology experiments to see which held up¹. “So we thought it would be fantastic to bet on the outcome,” says Dreber, who leads a team of behavioural economists at the Stockholm School of Economics.

In particular, her team wanted to see whether scientists could make good use of prediction markets: mini Wall Streets in which participants buy and sell ‘shares’ in a future event at a price that reflects their collective wisdom about the chance of the event happening. As a control, Dreber and her colleagues first asked

a group of psychologists to estimate the odds of replication for each study on the project’s list. Then the researchers set up a prediction market for each study, and gave the same psychologists US\$100 apiece to invest.

When the Reproducibility Project revealed last year that it had been able to replicate fewer than half of the studies examined², Dreber found that her experts hadn’t done much better than chance with their individual predictions. But working collectively through the markets, they had correctly guessed the outcome 71% of the time³.

Experiments such as this are a testament to the power of prediction markets to turn individuals’ guesses into forecasts of sometimes startling accuracy. That uncanny ability ensures

ILLUSTRATION BY SÉBASTIEN THIBAUT

that during every US presidential election, voters avidly follow the standings for their favoured candidates on exchanges such as Betfair and the Iowa Electronic Markets (IEM). But prediction markets are increasingly being used to make forecasts of all kinds, on everything from the outcomes of sporting events to the results of business decisions. Advocates maintain that they allow people to aggregate information without the biases that plague traditional forecasting methods, such as polls or expert analysis (see page 304).

In science, applications might include giving agencies impartial guidance on the proposals that are most worth funding, helping panels to find a consensus in climate science and other fields or, as Dreber showed, giving researchers a fast and low-cost way to identify the studies that might face problems with replication.

But sceptics point out that prediction markets are far from infallible. “There is a viewpoint among some people that once you set up a market this magic will happen and you’ll get a great prediction no matter what,” says economist Eric Zitzewitz at Dartmouth College in Hanover, New Hampshire. That is not the case: determining the best designs for prediction markets, as well as their limitations, is an area of active research.

Nevertheless, prediction-market supporters argue that even imperfect forecasts can be helpful. “Hearing there’s an 80 or 90% chance of rain will make me take an umbrella,” says Anthony Aguirre, a physicist at the University of California, Santa Cruz. “I think there’s a big space between being able to time travel and physically see what will happen, and then throwing up your hands and saying it’s totally unpredictable.”

THE MAGIC OF GAMBLING

People have been betting on future events for as long as they have played sports and raced horses. But in the latter half of the nineteenth century, US efforts to set betting odds through marketplace supply and demand became centralized on Wall Street, where wealthy New York City businessmen and entertainers were using informal markets to bet on US elections as far back as 1868. These political betting pools lasted into the 1930s, when they fell victim to factors such as stricter gambling laws and the rise of professional polling. But while they lasted they had an impressive success rate, correctly picking the winners of 11 out of 15 presidential races, and correctly identifying that the remaining 4 contests would have extremely tight margins.

The prediction-market idea was revived by the spread of the Internet, which dramatically lowered the entry barriers for creating and participating in prediction markets. In 1988, the University of Iowa’s Tippie College of Business

launched the not-for-profit IEM as a network-based teaching and research tool; ahead of the 8 November presidential election that year, they set up a market to predict the fraction of votes that would go to each of the presidential candidates (see ‘How a market predicts’). The fractions changed daily as traders interpreted fresh

participated in prediction markets that beat the firm’s official projections of quarterly printer sales 75% of the time. And in September 2002, six months before the US-led invasion of Iraq, the Dublin-based betting site TradeSports.com gained international notoriety when it ran a prediction market on when Iraqi dictator Saddam Hussein would be ousted. By the time the war began in March 2003, bettors were 90% certain Hussein would be out by April and 95% sure he’d be gone by May or June. He was deposed in April.

MARKET RESEARCH

Prediction markets have also had some high-profile misfires, however — such as giving the odds of a Brexit ‘stay’ vote as 85% on the day of the referendum, 23 June. (UK citizens in fact narrowly voted to leave the European Union.) And prediction markets lagged well behind conventional polls in predicting that Donald Trump would become the 2016 Republican nominee for US president.

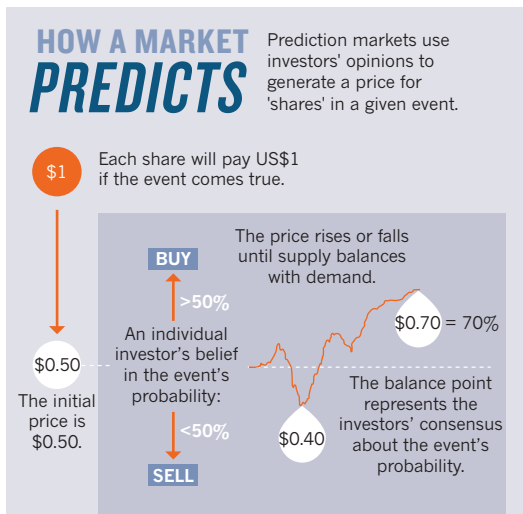
Such examples have inspired academics to probe prediction markets. Why do they work as well as they do? What are their limits, and why do their predictions sometimes fail?

Perhaps the most fundamental answer to the first question was provided in 1945 by Austrian economist Friedrich Hayek. He argued that markets in general could be viewed as mechanisms for collecting vast amounts of information held by individuals and synthesizing it into a useful data point — namely the price that people are willing to pay for goods or services.

Economists theorize that prediction markets do this information gathering in two ways. The first is through ‘the wisdom of crowds’ — a phrase popularized by business journalist James Surowiecki in his book of that name (Doubleday, 2004). The idea is that a group of people with a sufficiently broad range of opinions can collectively be cleverer than any individual. An often-cited case is a game

in which participants are asked to estimate the number of jelly beans in a jar. Although individual guesses are unlikely to be right, the accumulated estimates tend to form a bell curve that peaks close to the actual answer. When investor Jack Treynor ran this experiment on 56 students in 1987, their mean estimate for the number of beans — 871 — was closer to the correct answer of 850 than all but one of their guesses⁵.

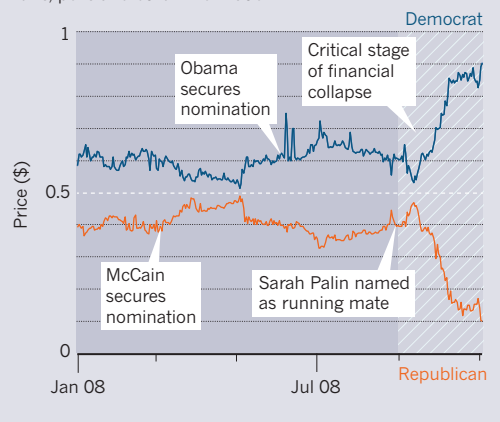
As Surowiecki and others have emphasized, however, crowds are wise only if they harbour a sufficient diversity of opinion. When they don’t — when people’s independent judgements are skewed by peer pressure, panic or even a charismatic speaker — the wisdom of crowds can easily fall prey to collective breakdowns. The housing bubble of the mid-2000s, which was a major contributor to the 2007–08 financial



EXAMPLE: 2008 US PRESIDENTIAL ELECTION

Barack Obama (Democrat) versus John McCain (Republican)

The market's prediction constantly shifted to accommodate news, polls and other information.



information about polls, the economy and other issues. On the eve of the election, the market predicted that the Republican nominee, George H. W. Bush, would be victorious with 53.2% of the vote — which is exactly what he got. And in 2008, a study found that the IEM’s predictions across five presidential elections were more accurate than the polls 74% of the time⁴.

The success of the IEM helped to inspire the creation of dozens of other prediction markets. In 1996, for example, the Hollywood Stock Exchange was launched to forecast opening-weekend box-office take and other movie-related outcomes; its markets correctly predicted that *Hamlet* would be a flop that year and that *Jerry Maguire* would be a hit. In the early 2000s, employees of information-technology company Hewlett-Packard

crash, was one such breakdown of judgement. But this is where the second market mechanism comes in. Sometimes called the marginal-trader hypothesis, it describes how — in theory — there will always be individuals seeking out places where the crowd is wrong. In the process, these traders will identify undervalued contracts to buy and overvalued contracts to sell, which tends to push prices back towards a sensible value. An example can be seen in the 2015 film *The Big Short*, which dramatizes the true story of a hedge fund that bet against the irrational exuberance of the US housing market and gained substantially from the crash.

“WHEN SOMEONE STARTS TO SUGGEST A BET, PEOPLE IMMEDIATELY START TO CLARIFY WHAT THEY MEAN.”

Laboratory experiments have been used to test many aspects of this theoretical framework, including how well prediction markets aggregate information under different conditions. In a 2009 experiment⁶ that was designed to mimic scientific research and publishing, researchers set up three prediction markets in which participants tried to predict which hypothesis about a fictitious biochemical pathway would end up being true.

FIELD-TESTING THE FUTURE

In one market, key pieces of information about the pathway were available to all participants; the traders quickly converged on the correct answer. In another, analogous to proprietary corporate research, information was privately held by individuals; the traders often failed to reach a consensus. And in the third, analogous to results being discovered in different labs and then published in journals, information was initially kept private and then made public. The market was able to find the right answer — but the individuals who discovered useful information first could use their private knowledge to anticipate the markets and extract a small profit.

One of the first prediction markets devoted exclusively to scientific questions grew out of a project started in 2011 by economist Robin Hanson of George Mason University in Fairfax, Virginia. Eventually known as SciCast, the project included a website where participants could wager on questions such as, “Will there be a lab-confirmed case of the coronavirus Middle East Respiratory Syndrome (MERS or MERS-COV) identified in the United States by 1 June 2014?” (There was.) SciCast’s assessments were more accurate than an uninformed prediction model 85% of the time (see go.nature.com/2dm6llp).

SciCast was discontinued in 2015, when

its funding ended. But it helped to inspire Metaculus, a market launched in November 2015 by Aguirre and his colleague Greg Laughlin, an astrophysicist now at Yale University in New Haven, Connecticut. The site grew out of Aguirre’s interest in finding ‘superpredictors’ — people whose forecasting skills are far above average. Metaculus asks participants to estimate the probabilities of such things as, “Will a clinical trial begin by the end of 2017 using CRISPR to genetically modify a living human?” or “Will the National Ignition Facility announce a shot at break-even fusion by the start of 2017?”

As in SciCast, Metaculus participants do

not use actual money: players instead move a slider representing their belief in the likelihood of an answer and accrue a track record for being correct. The lack of cash bets is partly a matter of practicality, says Aguirre. “When it’s ‘Will Hillary win?’, zillions of people will buy on that. But if it’s ‘Will this new paper on arXiv get more than ten citations?’, you’re not going to find enough people with real money to make an accurate prediction.” But it’s also the case that real money isn’t strictly necessary for a successful prediction market: several studies^{7,8} have shown that traders can be equally well motivated by the prestige of being right.

Metaculus currently has around 2,000 active users, although its creators hope to accrue 10,000 or more. Already, the site has produced evidence that successful prediction is a skill that can be learned. The best players work out the optimal time to adjust their guesses up or down, and their performance gradually improves.

Laughlin and Aguirre suggest that Metaculus could be useful to journalists and other members of the public who want to know which questions most interest scientists. Funding agencies might similarly be attracted to its results. “Having prediction markets that are getting an even-handed assessment is potentially a way of aiding the decision for what projects are most worth funding,” says Laughlin.

But scientific prediction markets have yet to gain much traction with researchers or the public. One important reason is that most political and business questions get clear-cut answers in relatively short time periods, and this is where prediction markets excel. But few would-be traders have the patience to endure the decades of effort, ambiguity and experimentation that are often required to answer questions in science.

This problem is hardly unique to prediction

markets, however: “It is in general easier to make short-term than long-term predictions,” says Aguirre. As long as prediction markets offer a way to update guesses in light of new information, proponents argue, they will do as well or better than other forecasting methods.

Scientific prediction markets also suffer more from ambiguity issues than do political or economic ones. In an election, one person is eventually declared the winner, whereas in science, resolutions are rarely so neat. But prediction-market advocates don’t think that this is necessarily a cause for concern. “When someone starts to suggest a bet, people immediately start to clarify what they mean,” says Hanson. Aguirre says that he and Laughlin take great pains on Metaculus to ensure that predictions are well-defined and easy to understand.

Whether prediction markets can work for science remains an open question. When Dreber’s team repeated 18 economics experiments as part of a follow-up to her psychology investigation, both the prediction markets and surveys of individuals overestimated the odds of each study’s reproducibility⁹. Dreber isn’t sure why this happened. She points out that the psychologists in the first study were all already interested in replication — whereas the economists in the second were not involved in the reproducibility project — so they might have been better at collectively estimating reproducibility.

Prediction markets in general still need to deal with challenges such as how to limit manipulation and overcome biases. Yet conventional representative polling, which once relied on answers from phonecalls to randomly sampled landlines, is being jeopardized by the movement to mobile phones and online messaging. Because the accuracy of prediction markets is at least on par with, if not better than, polls, economist David Rothschild of Microsoft Research in New York City thinks that prediction markets are well placed to take over if polling goes into decline. “I can create a poll that can mimic everything about a prediction market,” he says. “Except markets have a way of incentivizing you to come back at 2 a.m. and update your answer.” ■

Adam Mann is a freelance writer in Oakland, California.

1. Open Science Collaboration *Perspect. Psychol. Sci.* **7**, 657–660 (2012).
2. Open Science Collaboration *Science* **349**, aac4716 (2015).
3. Dreber, A. et al. *Proc. Natl Acad. Sci. USA* **112**, 15343–15347 (2015).
4. Berg, J. E., Nelson, F. D. & Rietz, T. A. *Int. J. Forecast.* **24**, 285–300 (2008).
5. Treynor, J. L. *Financ. Anal. J.* **43**, 50–53 (1987).
6. Almenberg, J., Kittlitz, K. & Pfeiffer, T. *PLoS ONE* **4**, e8500 (2009).
7. Rosenbloom, E. S. & Notz, W. *Electron. Mark.* **16**, 63–69 (2006).
8. Servan-Schreiber, E., Wolfers, J., Pennock, D. M. & Galebach, B. *Electron. Mark.* **14**, 243–251 (2004).
9. Camerer, C. F. et al. *Science* **351**, 1433–1436 (2016).