

Wisdom of the masses

Andrew Leigh and Justin Wolfers look into the future of prediction markets

In 2006, Google decided that it wanted to obtain better forecasts of what it and its competitors were planning to do. As with most companies, better predictions of product launch dates, office openings and so on are of significant strategic importance to Google. But instead of using the standard approach — forecasting the future by asking designated experts — Google opted to set up a market, in which any Google staff member could bet on the chances of an event coming true. If a Google staffer had strong information that an event would come true, then regardless of their rank in the company, they could bet on that outcome.

The results — based on the aggregated bets of thousands of Google staff members — were strong predictors of the actual outcomes. Where the trading volume was high, the forecasts were even more precise. And as closing time drew closer, the markets became steadily more accurate.

This example describes an increasingly important information-aggregation tool, known as “prediction markets”. Analytically, these are markets where participants trade in contracts whose payoff depends on unknown future events — just as in any financial or betting market. The defining feature of a prediction market is that the price of these contracts can be directly interpreted as a market-generated forecast of some unknown quantity.

Much of the enthusiasm for prediction markets derives from the efficient markets hypothesis: in a truly efficient market, the price of a financial security or prediction market contract reflects all available information. Thus, efficient prediction market prices hold the promise of yielding efficient and unbiased forecasts. Indeed, the market price may be not only the single best predictor of the event but also the most efficient aggregator of available information, such that no combination of available polls or other information can be used to improve on the market-generated forecast.

This statement does not require that all individuals in a market be rational, as long as the marginal trade in the market is motivated by rational traders. Of course, it is unlikely that prediction markets are literally efficient, but a number of successes in these markets, both with regard to predicting public events and corporate outcomes, have generated substantial interest among social scientists, policymakers and the business community.

We begin by outlining the types of prediction markets on offer, and then discuss their accuracy. We review what is known about how to design an effective market, and the constraints of the present legal regime. Finally, we conclude with our own prognostications on what the future might hold.

Perhaps the best-known type of prediction markets are election-betting markets. The Iowa Electronic Market, established by political scientists at the University of Iowa in 1988, is perhaps the world's best known election market.



Illustration: MICHAEL FITZJAMES

These academics operate an electronic market in which traders can purchase “futures contracts” that consist of a promise to pay \$US1 if the candidate wins the popular vote. Thus, the price of this contract reflects the probability of a candidate winning the election. Assuming market efficiency, these prices should yield assessments that reflect all available information — including polls, the state of the economy and recent policy pronouncements.

For United States presidential elections, the Iowa Electronic Market has tended to be more accurate than opinion polls. Research has suggested that this is because the betting market focuses on the underlying dynamics of the race, and is therefore better able to parse out events that occur several months before the election, but will not change the outcome. Conversely, the betting market responds rapidly to occurrences that affect the underlying dynamics of the race (such as the appointment of a new campaign manager), even if these events elicit relatively little response in the media or polls.

In Australia, looser regulation of sports betting in the ACT and the Northern Territory has led several bookmakers — including Betfair, Centrebet, International All Sports, SportingBet and SportsAcumen — to offer punters the ability to bet on election outcomes. True to national traditions, Australian gamblers wagered more than \$1.5 million on the 2001 federal election, and more than \$2.6 million on the 2004 federal election. While betting markets do not “look” like financial markets, the betting odds yield the same directly interpretable forecasts offered by the Iowa Electronic Markets.

Outside the political sphere, there is a growing number of web-based prediction markets, often run by companies that provide a range of trading and gambling services. Some prominent examples include

Tradesports.com (also known as InTrade.com), Betfair.com, and pseudo-markets (in which participants trade virtual currency) such as Newsfutures.com and Ideosphere.com.

Some prediction markets focus on economic statistics. For instance, Goldman Sachs, Deutsche Bank, and the Chicago Mercantile Exchange have launched markets on the likely outcome of future readings of economic statistics, including employment, retail sales, industrial production, and business confidence.

Other markets also forecast private-sector returns. In several cases (such as the Google example above), private companies have found innovative ways to use prediction markets as a business forecasting tool.

Arguably the most important issue with these markets is their performance as predictive tools. In the political domain, the Iowa Electronic Markets has both yielded very accurate predictions for US elections, and also outperformed the major pollsters. In a comparison of election-eve forecasts, the Iowa markets predicted candidates' vote shares with an average absolute error of 1 percentage point, while the average forecast error for Gallup polls was 2.1 percentage points.

In Australia, election-eve betting markets correctly predicted that the coalition would win the 2001 and 2004 federal elections. By contrast, a coalition win was predicted by the election-eve polls of only two out of three major pollsters in 2001, and only two out of four major pollsters in 2004.

Perhaps more interesting in terms of how well prediction markets can aggregate information is the performance of markets at the level of the individual district. Typically districts are sufficiently small that there is little interest (or funding) for local polling, yet when Australian bookmakers started betting on

district-level races, we found they were extremely accurate. In 2001, the Centrebet favourite won in 43 out of 47 marginal seats. In 2004, the Centrebet favourite won in 24 out of 32 marginals.

That said, comparing the performance of markets with a mechanistic application of poll-based forecasting may not provide a particularly compelling comparison. A more relevant starting point might be to compare the predictions of markets with those of independent analysts.

For example, a survey of 10 experts published on the Sunday before the 2004 federal election found that three thought Mark Latham would win, while seven thought John Howard would win, but with a smaller majority than in 2001. None forecast the true result — a Howard victory with an increased majority. By contrast, the betting market was more confident of a Howard victory in 2004 than in 2001.

In a corporate context, the Hollywood Stock Exchange predicts opening weekend box-office success, and these predictions have been quite accurate. Further, this market has been about as accurate at forecasting Oscar winners as an expert panel. Some companies have also begun to experiment with internal prediction markets. An internal market at Hewlett-Packard produced more accurate forecasts of printer sales than the company's internal processes.

Gerhard Ortner described an experiment at Siemens in which an internal market predicted that the engineering giant would definitely fail to deliver on a software project on time, even when traditional planning tools suggested that the deadline could be met.

While the Hollywood markets have drawn many participants simply on the basis of their entertainment value, the HP and Siemens experiences suggested that motivating employees to trade was a major challenge. In each case, the

companies ran real money exchanges, with only a relatively small trading population, and subsidised participation in the market, by either endowing traders with a portfolio, or matching initial deposits. The predictive performance of even these very thin markets was quite striking.

Yet as the Google example discussed in the opening paragraph of this article suggests, internal markets with hundreds of traders are likely to have even greater forecasting power.

In another recent prediction market, traders in *Economic Derivatives* predict the likelihood that economic data released later in the week will take on specific values. The traditional approach to aggregating forecasts is to simply take an average or a “consensus estimate” from a survey of 30 or so professional forecasters. Comparing average market and consensus forecasts for four variables (non-farm payrolls, retail trade, unemployment claims and business confidence), shows that the market-based predictions of these economic indicators are not only extremely close to the corresponding “consensus” forecast, but also typically closer to the actual outcome, outperforming the experts on average.

The success of prediction markets, like any market, can depend on their design and implementation. For a prediction market to work well, contracts must be clear, easily understood and easily adjudicated. This requirement for clarity can sometimes turn out to be complex. For example, in the course of Siemens's internal prediction market on whether a software project would be delivered to the client on schedule, the client changed the deadline.

One intriguing question is how much difference it makes whether prediction markets are run with real money or with some form of play

Wisdom of the masses

From page 3

money. Legal restrictions on gambling have led some groups such as NewsFutures.com to adopt play money exchanges, with those who amass the largest play-fortunes eligible for prizes.

One suggestive experiment compared the predictive power of the prices from a real-money and play-money exchanges over the 2003 NFL football season, finding that both yielded predictions that were approximately equally accurate. Interestingly, both sets of prices also outperformed all but a dozen of 3000 people in an online contest, and also easily outperformed the average assessments of these "experts".

On "play money" exchanges, such as Foresight Exchange, one often sees quite loosely worded "contracts" such as that a "scientific study will conclude that astrology is a statistically significant predictive method to describe an individual's personality traits".

Some prediction markets will work better when they concern events that are widely discussed, since trading on such events will have a higher entertainment value and there will be more information on whose interpretation traders can disagree with. Ambiguous public information may be better in motivating trade than private information, especially if the private information is concentrated, since a cadre of highly informed traders can easily drive out the partly informed, repressing trade to the point that the market barely exists.

Indeed, attempts to set up markets on topics where there are insiders with substantial information advantages have typically failed. For instance, the Centrebet market on who will lead the Liberal Party to the 2007 federal election has generated very little trade despite the inherent interest in these questions. The same has proved true for prediction markets on when a US Supreme Court justice will choose to retire.

In the United States, the legal environment has forced onshore prediction markets to make compromises that have limited their attractiveness. The Iowa Electronic Markets agreed to limit positions to \$US500 in order to receive a "no action" letter from the Commodity Futures Trading Commission. This compromise limits the scope and depth of their markets, and possibly their efficiency. Platforms such as NewsFutures, the Hollywood Stock Exchange and the Foresight Exchange operate using play money, albeit play money that can sometimes be converted into prizes. Thus these platforms evade anti-gambling legislation by operating as games of skill.

The consulting firms running corporate prediction markets have taken the same approach as experimental economists on campuses: they have subsidised participation, allowing everyone to leave a winner, albeit to varying degrees.

Finally, two markets that we are aware of, Economic Derivatives and Hedgestreet, have obtained regulatory approval to offer trading of innovative futures contracts. Hedgestreet started in 2004, and Economic Derivatives, a joint venture by Deutsche Bank and Goldman Sachs that runs markets in economic numbers, has been operating since 2002. Its traders are

large institutions, and its markets have attracted volumes of hundreds of millions of dollars. While the US government has taken a strong anti-gambling stance, it seems that innovation will continue apace with new contracts being launched instead on financial markets.

In Australia, the legal environment has prevented prediction markets from establishing themselves in most states. In addition, there is a concern that for major bookmakers such as Centrebet, trading on economic derivatives would bring them into conflict with the Australian Stock Exchange and the Australian Futures Exchange. Since Australian betting agencies already handle significant sums of money for elections and major sporting events, relaxation of the regulation governing such markets would bring little risk, but a significant public benefit.

Prediction markets are extremely useful for estimating the market's expectation of the probability of a particular event occurring. Markets on election outcomes serve a useful public purpose in providing an up to the minute barometer of the probability that a given party or candidate will be elected. For companies whose business is exposed to a large degree of political risk, our advice is to follow the betting markets, not the opinion polls or the televised talking heads.

Prediction markets doubtless have their limitations, but they may be useful as a supplement to the other relatively primitive mechanisms for predicting the future such as opinion surveys, politically-appointed panels of experts, hiring consultants or holding committee meetings.

We are already seeing increasing interest in these markets in the private sector, with the experiments at HP, Microsoft and Google now being supplemented with new markets on pharmaceuticals and the likely success of future technologies on NewsFutures.

With the exception of a short-lived attempt by a US defence agency to create information markets, governments have shied away from directly operating them. This may be a sensible policy stance — yet politicians should also recognise that if the valuable information generated by trade in these markets is not fully internalised into the profits earned by these private companies, prediction markets will be underprovided. In the area of economic derivatives, it might well be in the interests of government departments to fund the creation of markets that aggregate traders' expectations of future economic releases. Indeed, the Reserve Bank of Australia already closely tracks the difference in yields on inflation-indexed bonds and non-indexed bonds — in effect a prediction market on future inflation.

More generally, the highly restrictive stance that governments have taken towards regulating prediction markets does not appear to take any account of the informational benefits of such markets. Would there be a public benefit from liberalising gambling regulations to give prediction markets a freer hand? You bet.

Andrew Leigh is an economist in the Research School of Social Sciences at the Australian National University. Justin Wolfers is an economist in the Wharton School at the University of Pennsylvania. This is an edited version of an article that first appeared in *The Melbourne Review*.

The secrets of Casanova

Michael Dirda finds more fact than fiction

In the popular imagination, the Venetian adventurer Giacomo Casanova (1725-1798) is, above all, a smooth operator, the archetypal "man who loved women" even if no one today could possibly believe such a sweet talker. Or could they?

Whether Casanova really did all the things he is known for is a question that haunts critical opinion surrounding his enthralling and immensely long book, *History of My Life*. Over the years, this upstart son of Venetian actors seems to have insinuated himself into the palaces of cardinals and the arms of courtesans, hobnobbed with famous philosophers and notorious charlatans. By the time he was 50, Casanova — aka the Chevalier de Seingalt, aka the occult master Paralisée Galtinarde, aka the spy Antonio Pratolini — had lived in Venice, Rome, Istanbul, Corfu, Paris, Geneva, Vienna, Marseille, London, Berlin, Madrid, Moscow, Warsaw and Trieste.

Ten years later — broke, ailing and apparently cast aside by Fortune — he accepted a sinecure as the librarian at a nobleman's estate in Dux, a Bohemian backwater of the Austro-Hungarian Empire. There, from 1790 to 1798, as blood flowed in the streets of Paris, Casanova scribbled away, recalling happier times, largely for his own amusement and consolation. As he said in a letter:

"I write 13 hours a day which passes like 13 minutes. What pleasure it is to recall pleasures! But what pain it is to recollect pains. I amuse myself because I do not invent. What bothers me is the necessity I am under to disguise the names for I have no right to publicise the affairs of others."

Early chapters of his story, which he wrote in French, were read by the debonair Prince de Ligne (the so-called first gentleman of Europe), who enthusiastically compared them to Montaigne's essays in their frankness.

But the book itself was never finished. Casanova had just reached the year 1774 — and his return to Venice after nearly two decades in exile (because he was wanted for escaping from the Leads, the infamous prison in the Doge's Palace) — when he succumbed to an infection of the genito-urinary tract. (Not surprising in one who had, he said, experienced at least 11 episodes of the "pox".) According to several witnesses, his last words were: "I have lived as a philosopher, and die as a Christian."

The manuscript of the *History* then drifted into the hands of a grandnephew, who sold it in 1821 to the Leipzig publishing firm of Brockhaus, which published it in a loose German translation (1822-1828). This was followed a few years later by a French text (1826-1838), edited and reworked by a Jean Laforgue, who had been commissioned by Brockhaus to polish the occasionally Italianate diction and tone down some of its sex scenes (even though the decorous Casanova is never graphic). All subsequent translations, including the one into

English by Arthur Machen (1894), necessarily followed this 19th-century travesty. Not until 1960-1962 did Brockhaus, in conjunction with Plon in Paris, finally print the original, to the joy of Casanovists around the world. Willard Trask expertly rendered this text into English in 1966, and his is the English version to read, whether entire (12 volumes bound as six) or in the new abridged Everyman edition, which trims the narrative in half, for those with "other calls on their attention".

The memoirs themselves offer far more than some 4000 pages of genteel 18th-century sexual adventures. They provide a top-to-bottom survey of the glittering, shabby world of Europe before the French Revolution, from the *douceur de vivre* of the aristocrats to the desperation of the common people. The historians J H Plumb and John Julian Norwich have written admiringly of this portrait of the age, and writers as diverse as Stefan Zweig, Edmund Wilson, and V S Pritchett have found the memoirs compulsively readable.

In a textual note to the Everyman edition, Peter Washington compares Casanova's *History of My Life* to Proust's *In Search of Lost Time*: "Characters recur; experiences are repeated in new circumstances, places revisited, philosophies assessed, adventures pondered, all with cumulative force. . . . As the author ages, shadows fall across the brilliant surface of his story. A Proustian sense of the passage of time appears, compounded by an equally Proustian need to understand the relationship between outward events and personal evolution."

To liken Casanova's *History* to a great novel may suggest that the autobiography verges on fiction. After all, even the most empirical memoir is constantly prey to cosmetic enhancement. Lives are chaotic, messy, and crowded with inessentials, while art requires order, themes and a pleasing variety so as to transform the detritus of experience into an object of beauty. The temptation to sharpen or improve an anecdote can be irresistible. At what point does the merely artful pass into the fictional?

In Casanova's case, questions have periodically arisen about his overall reliability and truthfulness. (One 19th-century bookseller even argued that the whole *History* was made up by Stendhal.) Consider the book's most famous episode: Casanova's thrilling escape from the Leads, where he had been imprisoned for corrupting Venetian youth and being a general nuisance to certain men in power. Did he and a fellow prisoner actually break through the ceiling of his cell with a long metal bolt, and then scramble across the roof of the Doge's Palace? Or did his patron Matteo Bragadin simply bribe the jailer to let him out?

As it happens, nearly all the evidence confirms that the account in the memoirs is essentially true. The escape was famous during Casanova's lifetime and, because he was obliged to retell the story so



Hercules and Omphale by François

often, he actually published a short account of his "flight from the Leads" (1788). No one, then, questioned its authenticity. That he was the first person ever to break out of the famous prison was widely known throughout Venice. Previous inmates, with far better connections than Casanova, had never been able to buy their way to freedom. Today, virtually every scholar who has studied the Venetian state documents believes that Casanova escaped pretty much just as he says he did.

What about the lengthy conversations — such as the barbed exchanges with Voltaire — that add so much drama to the *History's* pages? Can anyone's memory be that good? In the case of the French philosopher, Casanova tells us that each night he scribbled down everything he could remember about his talks with the great man. The distinguished authority on Voltaire, Theodore M Besterman, was convinced of the essential accuracy of their encounters. At several points, Casanova casually alludes to his trunks full of papers.