Prediction markets for business and public policy

Andrew Leigh and Justin Wolfers

Recent successes in prediction markets, predicting both public events and corporate outcomes, have generated substantial interest among social scientists, policy makers and the business community. While these markets have their limitations, they may be useful as a supplement to more primitive mechanisms for predicting the future.

In 2006, Google decided that it wanted to obtain better forecasts of what the company and its competitors were planning to do. As with most firms, 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 her rank in the company, she 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.[1]

This example describes an increasingly important information-aggregation tool, known as 'prediction markets'.[2] 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 recent successes in these markets, both with regard to predicting public events and corporate outcomes, have generated substantial interest among social scientists, policy makers and the business community.

We begin by outlining the types of prediction markets currently 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 discuss the kinds of inferences that can be drawn from prediction markets and conclude with our own prognostications on what the future might hold.

Prediction markets — an overview

Perhaps the best-known type of prediction market is 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. These academics operate an electronic market in which traders can purchase 'futures contracts' that consist of a promise to pay $1 if the candidate wins the popular vote. Thus the price of this contract reflects the probability of a candidate winning the election.[3] Assuming market efficiency, these prices should yield assessments that reflect all available information — including polls, the state of the economy, and recent policy pronouncements.

> 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.

For US 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.[4] 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 Australian Capital Territory 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. Australian gamblers wagered over $1.5 million on the 2001 federal election, and over $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) and Betfair.com, and pseudo-markets (in which participants trade virtual currency) such as Newsfutures.com and Ideosphere.com. These websites often take the lead on defining a contract (e.g. in 2002–03, Tradesports ran a market on whether Saddam Hussein would be deposed as Iraqi leader), and then allow individuals to post their offers to buy and sell contracts. In this respect, the market structure largely shadows that in standard financial markets.

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

Table 1: Some Popular Prediction Markets

Market	Focus	Typical turnover on an event ($US)
Iowa Electronic Markets <www.biz.iowa.edu/iem> Run by University of Iowa	Small-scale election markets. Similar markets are run by: UBC (Canada) <www.esm.buc.ca> and TUW (Austria) <http://ebweb.tuwien.ac.at/apsm/>	Tens of thousands of dollars (Traders limited to $500 positions)
Centrebet <www.centrebet.com> For profit company	Northern Territory bookmaker, offering odds on election outcomes, current events, sports and entertainment.	Millions of dollars
TradeSports <www.tradesports.com> For profit company	Trades in a rich set of political futures, financial contracts, current events, sports and entertainment.	Hundreds of thousands of dollars, sometimes millions
Economic Derivatives <www.economicderivatives.com> Run by Goldman Sachs and Deutsche Bank	Large-scale financial market trading in the likely outcome of future economic data releases.	Hundreds of millions
Newsfutures <www.newsfutures.com> For profit company	Political, finance, current events and sports markets. Also technology and pharmaceutical futures for specific clients.	Virtual currency redeemable for monthly prizes (such as a TV)
Foresight Exchange <www.ideosphere.com> Non-profit research group	Political, finance, current events, science and technology events suggested by clients.	Virtual currency
Hollywood Stock Exchange <www.hsx.com> Owned by Cantor Fitzgerald	Success of movies, movie stars, awards, including a related set of complex derivatives and futures. Data used for market research.	Virtual currency

of economic statistics, including employment, retail sales, industrial production, and business confidence. Some markets also forecast private-sector returns. The Hollywood Stock Exchange allows people to use virtual currency to speculate on movie-related questions like opening-weekend performance, total box office returns, and likely Oscar winners. In several cases (such as the Google example above), private firms have found innovative ways to use prediction markets as a business forecasting tool.

How accurate are prediction market forecasts?

Arguably, the most important issue with these markets is their performance as predictive tools. In the political domain, the Iowa Electronic Markets have both yielded very accurate predictions for US elections and also outperformed the major pollsters.[5] In a comparison of election-eve forecasts, the Iowa markets predicted candidates' vote shares with an average absolute error of 1.5 percentage points, 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.[6]

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 that they were extremely accurate. In 2001, the Centrebet favourite won in 43 out of 47 marginal seats.[7]

Arguably, the most important issue with these markets is their performance as predictive tools. In the political domain, the Iowa Electronic Markets have both yielded very accurate predictions for US elections and also outperformed the major pollsters.

In 2004, the Centrebet favourite won in 24 out of 32 marginals.[8] However, 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.[9] Some firms have also begun to experiment with internal prediction markets. An internal market at Hewlett-Packard produced more accurate forecasts of printer sales than the firm's own processes.[10] Gerhard Ortner described an experiment at Siemens in which an internal market predicted that the firm would definitely fail to deliver on a software project on time, even when traditional planning tools suggested that the deadline could be met.[11]

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 firms 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 around 30 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.[12]

When several markets run simultaneously on the same event, arbitrage across markets becomes possible. Yet — consistent with market efficiency — most studies find little evidence of unexploited profit opportunities. For example, our analysis of data from five betting exchanges in the 2004 Australian election suggested only one potential arbitrage opportunity throughout the entire campaign (and even that would have yielded only a 2 per cent profit margin). Despite significant fluctuations, betting odds appear to co-move very closely.

Market manipulation represents a very serious concern about the predictive ability of these markets. Yet the profit motive has usually proven sufficient to ensure that attempts at manipulating prediction markets are unsuccessful. There have been several known attempts at manipulation of these markets, but none of them had much discernible effect on prices, except during a short transition phase. For example, one experimenter placed random $500 bets on the Iowa Electronic Markets to trace their effect.[13] Another researcher attempted to manipulate betting on horse races by cancelling large wagers at the last moment but this, ultimately,

had little effect.[14] During the 2004 US Presidential election campaign, one trader spent tens of thousands of dollars trying to shorten the odds of a Kerry Presidency on Tradesports.com but prices recovered to their previous level within an hour.[15]

Market design

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' 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-money. Legal restrictions on gambling have led some groups like 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 real-money and play-money exchanges over the 2003 NFL football season, finding that both yielded predictions that were approximately equally accurate.[16] Interestingly, both sets of prices also outperformed all but a dozen of 3,000 people in an online contest, and also easily outperformed the average assessments of these 'experts.' One practical advantage of play-money contracts is that they offer more freedom to experiment with different kinds of contracts. 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 higher entertainment value and there will be more information on whose interpretation traders can disagree. 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.

Legal constraints

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 $500 in order to receive a 'no action' letter from the Commodity Futures Trading Commission. This compromise limits the scope and depth of the 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.

Making inferences from prediction markets

How might business people or policy makers use the results from prediction markets in subsequent analysis? The most direct form of inference involves using these predictions directly. For instance, in the Hewlett Packard experiments discussed above, the markets elicited expectations of future printer sales, which were of direct interest for internal planning purposes.

Some analyses have tried to link the time series of expectations elicited in prediction markets with time series of other variables. For instance, in a paper we co-authored with Eric Zitzewitz, we tracked movements in the 'Saddam Security', a contract traded on the Tradesports exchange that paid $10 if Saddam Hussein was ousted by a given date.[17] Interpreting movements in the Saddam Security as an index for the risk of war, we estimated the extent to which oil prices and the S&P 500 index co-moved with the Saddam Security. Our conclusions (publicly released on the eve of the war) suggested that the Iraq war would lead to a $10 per barrel increase in oil prices, and a 15 per cent decline in stock prices.

> …attempts to set up markets on topics where there are insiders with substantial information advantages have typically failed.

Applying a similar methodology, Joel Slemrod and Timothy Greimel linked the price of a Steve Forbes election security in the 1996 Republican primary market with a rising interest rate premium on municipal bond prices, because Forbes's signature issue was a 'flat tax' that would have eliminated the tax exemption for municipal bond interest.[18] In a similar vein, Erik Snowberg, Justin Wolfers and Eric Zitzewitz compare movements in stock and bond prices with unexpected movements in prediction markets on election day, thereby measuring the effects of different candidates on the economy.[19]

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.

It seems quite possible to design prediction market contracts so that they bring out the connection between an event and other variables. For instance, Justin Wolfers and Eric Zitzewitz reported on experimental markets run on Tradesports in the months prior to the 2004 presidential election.[20] In one example, they ran markets linked to whether George W. Bush would be re-elected, whether Osama bin Laden would be captured prior to the election, and whether both events would occur. These markets suggested a 91 per cent chance of Bush being re-elected if Osama had been found but a 67 per cent unconditional probability.

In the Australian context, it is possible to imagine using these contingent markets to determine which party is more likely to bring down unemployment. Suppose we offer two contracts. The first pays $100 if Labor wins the 2007 federal election and unemployment falls in 2008, and if Labor loses the election, all monies are returned. The second contract pays $100 if the Coalition wins the 2007 federal election and unemployment falls in 2008, and if the Coalition loses, all monies are returned. The price difference between the two contracts may tell us something about the market's expectation as to which political party is likely to do a better job in lowering unemployment.

Robin Hanson has called these contingent markets 'decision markets', arguing (in a paper entitled 'Shall We Vote on Values but Bet on Beliefs?') that these expectations should be used to guide decision-making. Relatively simple contingent markets, as well as more complex combinatorial markets, are as yet virtually untested, and a useful focus for further research. There may be important and interesting applications in domains where these markets yield useful insights. (There remains an important caveat: these markets reveal expected correlations but correlation need not imply causation. For instance, it may be that voters are more willing to vote Liberal if they believe that unemployment is likely to fall anyway.)

Future applications

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 firms 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.

Contingent markets serve a related goal — helping to answer the question, 'If event A occurs, then how will it affect outcome X?'. Carefully designing contingent markets is trickier than simple prediction markets but the benefits from these markets are plausibly larger too.

Prediction markets doubtless have their limitations but they may be useful as a supplement to the other relatively primitive mechanisms for predicting the future like 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 Hewlett Packard, 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 firms, 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 — effectively 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. n

ENDNOTES

1 For a discussion of the Google prediction markets, see the presentation by B. Cowgill 2006, 'Google Prediction Markets: Putting Crowd Wisdom to Work', Prediction Markets Summit, New York, NY, 3 February, and a blog entry at http://googleblog.blogspot.com/2005/09/putting-crowd-wisdom-to-work.html.

2 For a comprehensive review which inspires much of this essay, see J. Wolfers and E. Zitzewitz 2004, 'Prediction Markets', *Journal of Economic Perspectives*, vol.18 no. 2, pp.107–126.

3 J. Wolfers and E. Zitzewitz 2006, 'Interpreting Prediction Market Prices as Probabilities,' *National Bureau of Economic Research* Working Paper no. 12200.

4 R. Forsythe, F. Nelson, G.R. Neumann, and J. Wright 1992, 'Anatomy of an Experimental Political Stock Market,' *American Economic Review*, vol. 82, no. 5, pp. 1142–61; D. Shaw and B. Roberts 2000, 'Campaign Events, the Media and the Prospect of Victory: The 1992 and 1996 US Presidential Elections', *British Journal of Political Science*, vol. 30, pp. 259–289.

5 J. Berg, R. Forsythe, F. Nelson and T. Rietz 2001, 'Results from a Dozen Years of Election Futures Markets Research', forthcoming in C. Plott and V. Smith (Eds) *Handbook of Experimental Economic Results*.

6 In 2001, ACNielsen and Newspoll correctly forecast a Coalition win, while Morgan forecast a Labor win. In 2004, ACNielsen and Galaxy correctly forecast a Coalition win, while Morgan forecast a Labor win, and

Newspoll's final poll had the parties 50–50.

7 J. Wolfers and A. Leigh 2002, 'Three Tools for Forecasting Federal Elections: Lessons from 2001', *Australian Journal of Political Science*, vol. 37, no. 2, pp. 223–240.

8 A. Leigh and J. Wolfers 2006, 'Competing Approaches to Forecasting Elections: Economic Models, Opinion Polling and Prediction Markets', *Economic Record*, vol. 82, no. 258, pp. 325–340.

9 D. Pennock, S. Lawrence, C.L. Giles, and F.A. Nielsen 2001, 'The Real Power of Artificial Markets', *Science*, vol. 291, pp. 987–988.

10 K-Y. Chen and C. Plott 2002, 'Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem', CalTech Social Science Working Paper no. 1131.

11 G. Ortner 1998, 'Forecasting Markets — An Industrial Application', *mimeo*., Technical University of Vienna.

12 R. Gürkaynak and J. Wolfers, 2005, 'Macroeconomic Derivatives: An Initial Analysis of Market-Based Macro Forecasts, Uncertainty and Risk', *NBER International Macroeconomics Annual.*

13 P. Rhode and K. Strumpf 2006, 'Manipulating political stock markets: a field experiment and a century of observational data', *mimeo*., University of North Carolina.

14 C. Camerer 1998, 'Can Asset Markets be Manipulated? A Field Experiment with Racetrack Betting', *Journal of Political Economy*, vol. 106, no. 3, pp. 457–482.

15 P. Rhode and K. Strumpf 2006, 'Manipulating political stock markets: a field experiment and a century of observational data', *mimeo*., University of North Carolina.

16 E. Servan-Schreiber, J. Wolfers, D. Pennock and B. Galebach 2004, 'Prediction Markets: Does Money Matter?', *Electronic Markets*, vol. 14, no. 3.

17 A. Leigh, J. Wolfers and E. Zitzewitz 2003, 'What do Financial Markets Think of War in Iraq?', NBER Working Paper no. 9587. National Bureau of Economic Research, Cambridge, MA.

18 J. Slemrod and T. Greimel 1999, 'Did Steve Forbes Scare the Municipal Bond Market?', *Journal of Public Economics*, vol. 74, no. 1, pp. 81–96.

19 E. Snowberg, J. Wolfers and E. Zitzewitz 2007, 'Partisan Impacts on the Economy: Evidence from Prediction Markets and Close Elections', *Quarterly Journal of Economics*, vol. 122, no. 2.

20 J. Wolfers and E. Zitzewitz 2004, 'Experimental Political Betting Markets and the 2004 Election', *The Economists' Voice*, vol. 1, no. 2, article 1.

21 R. Hanson 2004, 'Shall We Vote on Values, But Bet on Beliefs', *mimeo*., George Mason University (available at: http://hanson.gmu.edu)

Andrew Leigh

Andrew Leigh is an economist in the Research School of Social Sciences at the Australian National University. Email: andrew.leigh@anu.edu.au. Website: http://econrsss.anu.edu.au/~aleigh/

Justin Wolfers

Justin Wolfers is Assistant Professor of Business and Public Policy at the Wharton School, University of Pennsylvania, Faculty Research Fellow with the National Bureau of Economic Research, Research Affiliate with the Centre for Economic Policy Research, a Visiting Scholar with San Francisco Federal Reserve and a Research Fellow with Institute for the Study of Labor (IZA). Email: jwolfers@wharton.upenn.edu. Web: www.nber.org/~jwolfers