

18 Information (in)efficiency in prediction markets

Erik Snowberg, Justin Wolfers and Eric Zitzewitz

18.1 Introduction

This chapter examines a new class of markets at the intersection of traditional betting and traditional financial markets. We call these 'prediction markets'. Like both financial and betting markets, prediction markets focus on uncertain outcomes and involve trading in risks. Prices from these markets establish forecasts about the probabilities, mean and median outcomes, and correlations among future events. These prices have been used to accurately predict vote shares in elections, the box office success of Hollywood movies and the probability that Saddam Hussein would be deposed by a certain date. Other names for these markets include 'virtual stock markets', 'event futures', and 'information markets'.

Financial economists have long known about the information-aggregating properties of markets. Indeed, the efficient markets hypothesis, a centre-piece of financial theory, can be stated simply as, 'market prices incorporate all available information'. While financial instruments can be very complex, prediction markets tend to be analytically simple. Their current simplicity, however, belies their powerful potential future as a way to hedge against geopolitical and other forms of risk as envisioned by Athanasoulis, Shiller and van Wincoop (1999) and Shiller (2003).

Currently, most prediction markets are quite small, with turnover ranging from a few thousand dollars on the early political markets run by the University of Iowa, to several million bet in the 2004 election cycle on TradeSports, to hundreds of millions bet on the announcement of economic indicators in Goldman Sachs and Deutsche Bank's 'Economic Derivatives' market. The most famous prediction market is the Iowa Electronic Market (IEM), which was started in 1988 to predict the vote share of the two major party presidential candidates. Since then, they have amassed a record of more accurate prediction than polls, all while limiting trading positions to a cap of \$500.

The small size and relative newness of these markets can exacerbate the types of deviations from the efficiency seen in traditional financial markets. A key focus of this chapter is understanding the conditions under which market prices are most likely to provide accurate predictions. Some of our diagnoses are well understood, and simply require increased liquidity to rectify them, while others are more speculative, and should form the basis of further research. Better understanding of the sources and types of failures in prediction markets can only enhance their eventual usefulness.

This chapter also emphasises more complex contracts that are in active use today. These contingent contracts, or 'decision markets', hold the promise not just of predicting uncertain events, but also of providing useful forecasts under alternative scenarios, which may inform decision making.

We begin by briefly describing some simple types of contracts that are currently traded. We then examine the advantages and potential pitfalls of these markets. Finally, we survey the performance of existing markets, discuss contingent contracts, and conclude.

18.2 Design of prediction markets

Prediction market contracts are simply gambles on uncertain future events. Depending on the construction of the gamble, the price yields the market's expectation of different parameters. The simplest contract is one that pays a dollar if a certain event happens. The price of that contract at any given time is simply the market's belief about the percentage chance that the event will happen.¹

Another common gamble is 'spread betting' where participants take an even money bet on a particular outcome. This sort of betting is often practised in American football and basketball: one bets that a favoured team will win by a point spread of at least y points. In a political context, this might be a bet that pays off if a candidate earns over y percent of the vote. In both cases, the market, or market-maker, must adjust y such that supply equals demand, which requires that half of the bets fall on either side. Thus, the spread reveals the market's expectation of the median of $F(y)$.²

A final type of contract, which has proved less popular in sports betting is an 'index' bet. This contract pays off at the value of a particular parameter. For instance, sports bettors can buy a contract that pays off according to the number of runs a cricket team scores. This contract would thus reveal the market's expectation of the mean number of runs. This type of contract is most commonly used to predict a political candidate's share of the vote – much as a poll might.

By using variants and/or bundles of these three types of contracts, it is possible to construct contracts that will reveal the market's expectation of higher-order moments and more complicated parameters of the distribution of outcomes. One such variant, the contingent contract, pays off only if two or more events happen simultaneously. We discuss this type in greater detail later.

18.3 Applications and evidence

Prediction markets, in their most basic form, have been around since at least the beginning of the 1900s.³ However, until recently, there were very few active markets. The proliferation of the Internet and its use for sports betting has enabled an explosion of prediction contracts. Indeed, most of the examples in this chapter are taken from contracts that have been set up in the last few years. There are still many questions that need to be rigorously examined as the data becomes available, but already we can draw a few generalisations.

First, market prices tend to respond rapidly to new information. The following anecdote provides an interesting example. On 15 October 2003 the Cubs faced the Marlins in game six of the National League Championship Series. The Cubs were favoured to win at the beginning of the game and soon built a comfortable 3–0 lead. In the top of the 8th a contract that paid \$100 if the Cubs won was trading for over \$95. Then Steve Bartman, a fan, reached over and spoiled Moises Alou's catch of a foul ball. The Marlins proceeded to score 8 runs in the remainder of the inning. By the end of the 8th, the contract on the Cubs winning was trading at around \$5. Figure 18.1 shows the rapid incorporation of information into the contract price as the game progressed.

Not only is information rapidly incorporated into prices, but additional information also contributes to the accuracy of the forecasts made by prediction markets. Figure 18.2 shows the accuracy of the predictions of the IEM vote share market as a function of the time before election day. It is clear that as election day approaches and more information is revealed and incorporated into market prices, the accuracy of the prices as predictors increases.

Second, very few arbitrage opportunities exist. They appear briefly and represent small profit opportunities. Figure 18.3 shows the bid and ask prices on a contract that paid \$100 if Arnold Schwarzenegger was elected California's Governor in 2003, sampling data on bid and ask prices from two online exchanges every four hours. Both prices show substantial variation, but they move in lockstep. Arbitrage opportunities are virtually absent.

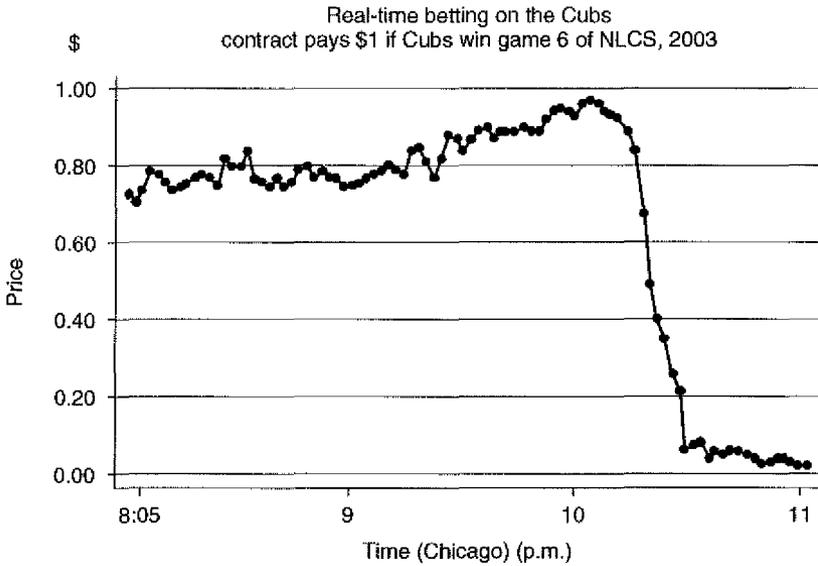


Figure 18.1 Rapid incorporation of information

Source: www.tradesports.com.

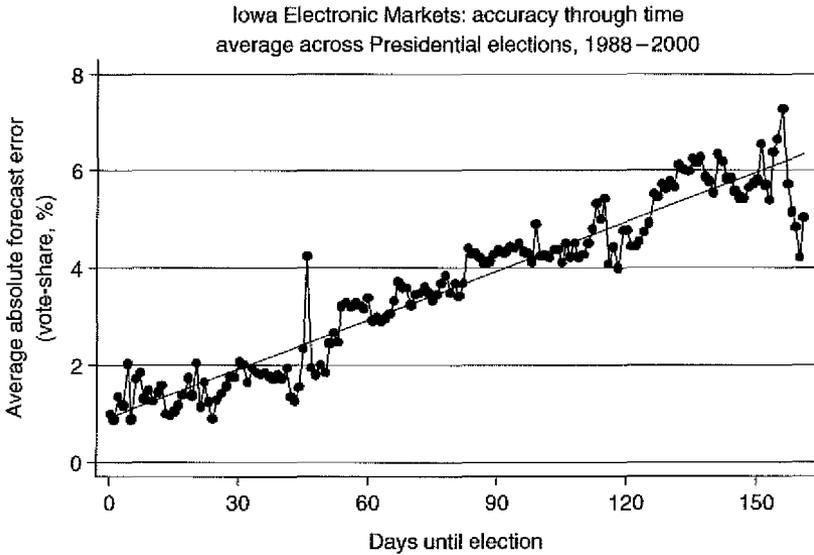


Figure 18.2 Information revelation through time

Source: Author's calculations based on data available at www.biz.uiowa.edu/iem/.

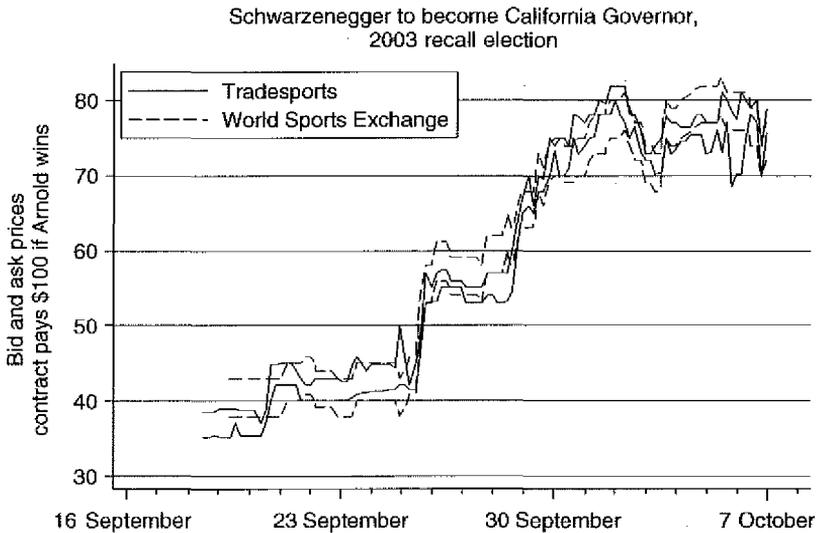


Figure 18.3 No arbitrage

Source: Prices collected electronically every four hours by David Pennock.

A third characterisation is that these markets, when well capitalised, appear to be robust to certain forms of manipulation. There are several case studies that emphasise this point. Rhode and Strumpf (2004) report that there were largely unsuccessful attempts by the big party bosses at manipulating the betting on early twentieth-century political markets. Strumpf (2004) placed random \$500 bets on the IEM and traced their effects, while Leigh and Wolfers (2002) provide examples of candidates betting on themselves in order to create a ‘buzz’. Camerer (1998) placed and cancelled large bets in parimutuel horse-racing markets. While all of these attempts at manipulation met with failure (except for brief, transitory effects) we obviously cannot draw any conclusions about the prevalence of the types of manipulation that have escaped the attention of analysts.

Finally, in most cases these markets seem to satisfy at least the weak form of the efficient markets hypothesis. There appear to be no profit opportunities from using simple strategies based on past prices. Leigh, Wolfers and Zitzewitz (2003) demonstrate this for the TradeSports ‘Saddam Security’, a contract that paid \$1 if Saddam Hussein was ousted by a particular date. Rhode and Strumpf provide evidence for early twentieth-century political markets. Tetlock (2004) reports that in general the financial and political

contracts that trade on TradeSports are efficiently priced. We provide more evidence on the accuracy of these markets in section 18.5.

18.4 When will prediction markets yield accurate predictions?

There are three main facets to prediction markets. First, the market structure is essentially an algorithm for aggregating (and sharing) opinions. Secondly, the financial and other incentives inherent in the market mechanism provide for truthful revelation. Finally, potential winnings provide robust incentives for information discovery. These features provide the power of prediction markets, and when one or more are missing the market's ability to predict will be undermined (Wolfers and Zitzewitz, 2004). We will address the problems in each category in turn.

18.4.1 Information aggregation

As the old saying goes, in the short-run markets are a voting machine, and in the long-run a weighing machine. Much of the power of markets derives from the fact that they provide an algorithm for aggregating diverse opinions: weighting the votes of market participants according to their willingness to back them with money.

However any algorithm will fail if it is deployed on an unclear task. Thus, contracts in prediction markets must be clear, easily understood and enforceable. A contract such as 'Howard Dean will win the presidential election' appears to satisfy the first two conditions, but could easily be challenged by a sore loser on the grounds that although Dean was clearly out of the running in the 2004 election, he may win in 2008. Adding a date, – i.e. 'Howard Dean will win the 2004 presidential election' – may not be enough as 'win' could refer to either the popular vote or the Electoral College. The requirement of clarity can be harder to satisfy than it appears at first glance. For instance, the day after the 1994 US Senate elections Senator Richard Shelby (Democrat Alabama) switched parties, throwing what seemed like a well-written contract on how many seats each party would take into confusion. Sometimes there is a trade-off between contractibility and capturing the event of interest. In 2003, TradeSports ran markets in 'Will there be a UN Resolution on Iraq (beyond #1441)?' and 'Will Saddam be out of office by June 30?' The former is clearly more contractible, but the latter is what traders wanted to bet on.

The key information aggregator is the market mechanism, and most prediction markets are run as a continuous double auction. Buyers and sellers submit bids and asking prices, respectively, and trade occurs when they reach a mutually agreeable price. Other markets, such as those used to

predict announcements of economic statistics, are run according to a parimutuel system. There is not enough data at this point to determine which market designs work best in which situations, particularly when markets are thin.

Since prediction markets have designs similar to many gambling markets, we can learn a lot about potential problems from studies of gambling. The longest-standing stylised fact regarding horse-race betting is the favourite-longshot bias, which is depicted in figure 18.4. Close examination of this phenomenon suggests that the behaviour it embodies is of concern in prediction markets, a point emphasised by Manski (2004).

On average, gamblers lose about 18 cents of every dollar wagered, and this ratio approximately holds for most horses – those with a 5 per cent to 50 per cent chance of winning. At the extremes, however, there are substantial deviations. Wagers on longshots produce much lower returns, offset by somewhat higher (albeit still negative) returns for betting on favourites. The overbetting of longshots ties in with a range of experimental evidence suggesting that people tend to overvalue small

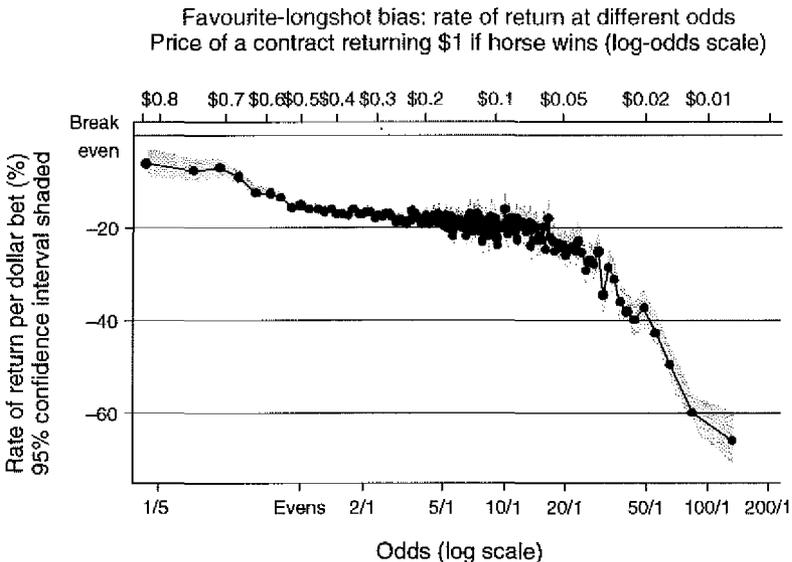


Figure 18.4 Miscalibration of small probabilities

Note: Sample is all horse races in the US, 1992–2002. $n = 5,067,832$ starts in 611,807 races.

Source: Trackmaster, Inc.

probabilities and undervalue near-certainties. We see that these errors persist beyond the psychology lab in equilibrium even in large and extremely active markets (Snowberg and Wolfers, 2004).

A related phenomenon is the 'volatility smile' found in options. This refers to overpricing of strongly out-of-the-money options, and underpricing of strongly in-the-money options relative to Black-Scholes benchmarks; thus the 'smile' refers to the shape of the relationship between implied volatility and strike price. Ait-Sahalia, Wang and Yared (2001) argue that the conclusion of miscalibration is less clear-cut in this context, because these prices may be driven by small likelihoods of extreme price changes. Additionally, when dealing with the pricing of options, one must take into account non-probabilistic factors such as wealth-dependent risk aversion, margin requirements and time to maturity. The effects of these constraints are more likely to be felt in small, poorly capitalised and long-horizon markets, so one should be especially careful when interpreting prices in such markets.

The miscalibration that causes the favourite-longshot bias and the 'volatility smile' appears in the pricing of certain securities related to financial variables on TradeSports. Table 18.1 reports the price of securities that paid off if the S&P finished 2003 in certain ranges. These securities can be approximated using December Chicago Mercantile Exchange (CME) S&P options. Comparing TradeSports prices with the state prices implied by CME option prices suggests that deep-out-of-the-money options are relatively overpriced on TradeSports. In the case of the most bearish securities, the price differences created a (small) arbitrage opportunity, one which persisted for most of the summer of 2003. Similar patterns existed for TradeSports' state securities on other financial variables (e.g. crude oil and gold prices, exchange rates, other indices). This is consistent with the favourite-longshot bias being more pronounced on smaller-scale exchanges.

While these behavioural biases may affect pricing in prediction markets, to the extent that they are systematic it remains possible to de-bias market prices so as to yield efficient forecasts.

18.4.2 Truthful revelation

Prediction markets must provide incentives for truthful revelation of information. However, these incentives do not necessarily need to be monetary. Indeed, the thrill of placing bets and the bragging rights of correct predictions may be enough to motivate traders. Some sites, such as NewsFutures.com use play money, where those who amass the largest play fortunes may be eligible for prizes. There is not enough evidence to

Table 18.1. *Price of S&P state securities^a on TradeSports versus CME market close, 23 July 2003*

S&P level at end of 2003	Price on TradeSports		
	Bid	Ask	Estimated state price ^b from CME S&P options
1200 and over	2	6	2.5
1100 to 1199	11	16	13.2
1000 to 1099	28	33	33.3
900 to 999	25	30	30.5
800 to 899	14	19	13
700 to 799	3	8	5
600 to 699	4	7	2
Under 600	5	8	1
S&P level on 23 July 2003		985	

Note:

^aPrices given are the price of a security that pays \$100 if S&P finishes 2003 in given range

^bState prices are estimated from CME option settlement prices using the method in Leigh, Wolfers and Zitzewitz (2003), adjusting for the thirteen-day difference in expiry date.

ascertain whether the use of real money makes an economically significant difference, although Servan-Schreiber, Wolfers, Pennock and Galebach (2004) provide suggestive evidence that play-money markets predicted NFL results as well as real-money markets. Since the only way to amass play-money is through a history of accurate prediction, it may even be that play-money outperform real-money exchanges. Since real- and play-money exchanges are not arbitrage linked there exist differences in the prices on the different types of exchanges. For example, in August 2003, Bush was a 2 to 1 favourite to win re-election on real-money exchanges, but was even-money on NewsFutures. By exploiting these differences in sufficiently large samples, it should eventually be possible to determine the factors driving the relative accuracy of real- and play-money exchanges.

Trading in prediction markets is much less attractive when the person you are betting against has control over the event in question, or if a relatively small group possesses most information on an event. Indeed, attempts to set up markets on topics where there are insiders with substantial information advantages have typically failed. For instance, market-makers withdrew

liquidity from markets on the winner of the pre-recorded reality show *Survivor* after CBS employees were accused of insider trading. Perhaps for the same reason, the TradeSports contracts on the next Supreme Court retirement have generated very little trade, despite the inherent interest in the question.

Finally, there is some evidence that the smaller-scale prediction markets are slower to incorporate information than deeper related financial markets. For example, Leigh, Wolfers and Zitzewitz (2003) found that changes in the 'Saddam Security' lagged war-related changes in the S&P or oil prices by 1–2 days. This is to be expected given that deeper financial markets have more traders investing larger sums of money, so there is more attention to buying and selling quickly when news breaks.

18.4.3 Information discovery and sharing

The incentives provided by a prediction market must be large enough to motivate the collection and sharing of information through the market mechanism. It is important to note here that although the vast amount of money in prediction markets may be uninformed, it is the marginal, not average dollar that sets prices. Thus, the presence of a few informed traders can still lead to very accurate predictions. It is because of this distinction between the average and marginal dollar driving prices that one cannot simply earn a profit betting against the New York Yankees (although one may derive some pleasure from doing so).

Figure 18.5 shows the price of a contract on whether or not weapons of mass destruction (WMD) will be found in Iraq. Note that at some points the value of the contract exceeded 80 per cent, yet weapons were never found. It is likely that this market performed poorly since the cost of gaining new information was quite high. Since WMD can be non-existent almost everywhere, but still exist somewhere, it was difficult to bet against the strong case made by the White House, at least initially.

Even if the market designer can avoid the above pitfalls a market will fail unless there is a motivation for trade. Trade in these markets can be motivated by a desire to hedge against risk, the thrill of pitting one's judgement against others, or a perceived profit opportunity on both sides because of divergent opinions over outcomes.

None of the prediction markets run on the websites we surveyed are large enough to truly hedge against significant risk. George W. Bush could not take a large enough stake against himself in order to ensure a win–win in the upcoming election. By providing contracts that are better linked to the underlying source of risk in individual portfolios, it seems likely that prediction markets will become more liquid, yielding more accurate pricing.

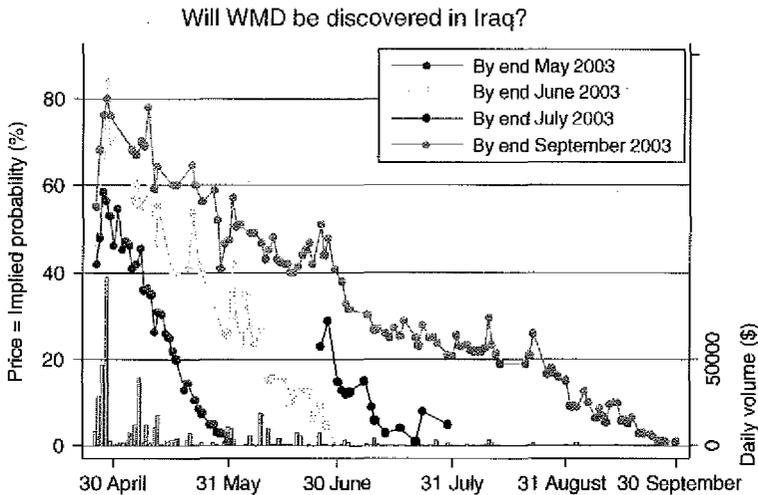


Figure 18.5 Inefficiencies in prediction markets

Source: Average daily price data provided by www.tradesports.com.

That said, as risk aversion becomes an increasingly important driver of trade, it may become necessary for researchers to adjust market prices for the risk premium, rather than interpreting them directly as probabilities.

These factors suggest that prediction markets are most likely to succeed when events are widely discussed with diverse interpretations of the available public information. The general interest creates both a larger pool of potential traders, as well as a greater thrill of being right. The public nature of the information makes it unlikely that there will be a perception of manipulation or corruption.

18.5 Performance of prediction markets

As troubling as some of the theoretical and practical problems with prediction markets may be, they generally – but not always – perform well. The evidence on this comes from a range of fields as diverse as the imaginations of the experimenters who use them. In the political domain, Berg, Forsythe, Nelson and Rietz (2001) summarise the evidence from the IEM, documenting that the market has both yielded very accurate predictions and also outperformed large-scale polling organisations. Figure 18.6 shows the aggregate forecast performance of all these experimental markets (or at least those for which data is publicly available). Each point represents the proportion of contracts trading at a given price that won. If

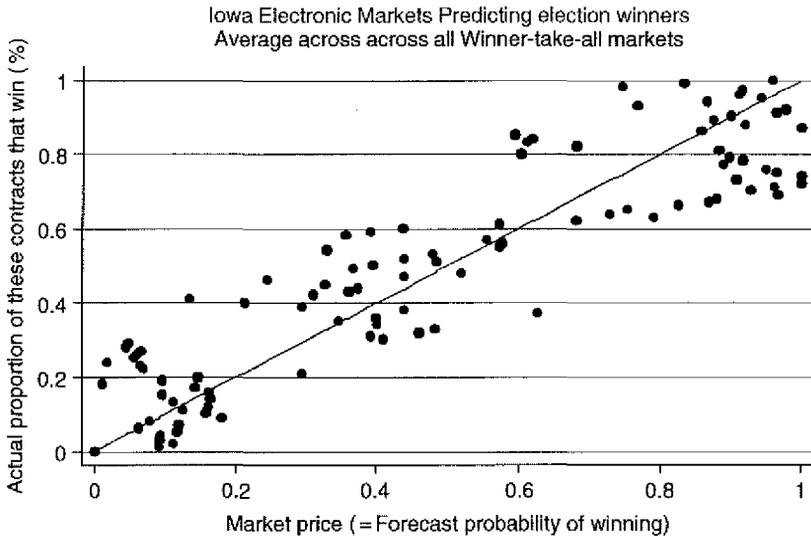


Figure 18.6 Accuracy of predictions

Note:

$n = 23,941$ daily price observations in 100 split-adjusted winner-take-all contracts over twenty-five elections.

Source: Author's calculations based on data available at: www.biz.uiowa.edu/icm/.

markets were perfectly accurate, then we would expect the data to lie along the 45 degree line. Not only are these markets typically quite accurate, but previous research has documented that they are better predictors than the Gallup poll.

In spite of the concerns we raise above about the amount of interest and liquidity necessary for a functioning prediction market, there are examples of smaller markets that work well. At the level of individual political districts there is often little interest in, or money for, local polling. Yet when Australian bookmakers started opening contracts on district-level races, Wolfers and Leigh (2002) show that they were extremely accurate.

Politicians and pundits use more than just polls when evaluating election chances and policy choices. They also rely on expert opinion. Figure 18.7 shows that the 'Saddam Security' co-moved tightly with both expert opinion (Will Saletan's 'Saddameter' – his estimate of the probability of the US going to war with Iraq) and oil prices (which respond to turmoil in the Middle East).

In a business context, Chen and Plott (2002) report that a well-designed internal market produced more accurate forecasts of printer sales than the

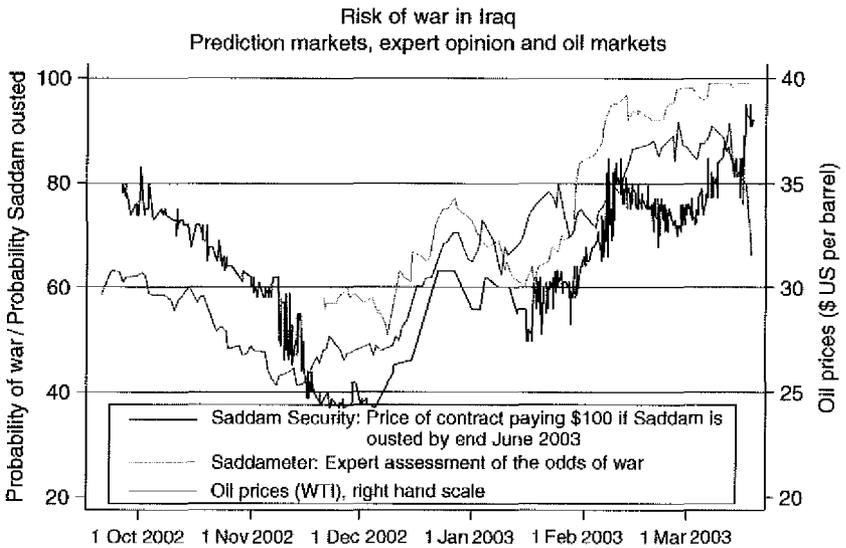


Figure 18.7 Correlation with expert opinion and other markets

Source: Trade-by-trade Saddam security data provided by Tradesports.com; Saddameter from Will Saletan's daily column in State.com.

firm's standard processes. Project planning has also been subjected to the judgement of prediction markets. Ortner (1998) launched an experimental market that predicted that a firm would definitely not meet its delivery target even when traditional planning tools suggested that it may have been on track. The market prediction proved correct. In the world of entertainment, Pennock, Lawrence, Giles and Nielsen (2001) show that the Hollywood Stock Exchange can usefully predict box office takes of films on their opening weekend and is about as accurate in picking Oscar winners as a panel of experts.

New markets in 'economic derivatives' also provide a useful contrast with expert opinion. Typically, in the run-up to the release of economic numbers such as inflation surveys, non-farm payrolls, retail trade and the Institute for Supply Management (ISM) purchasing managers' index, experts offer their opinions about what the numbers will be. These numbers are aggregated into a 'consensus forecast' and the market's reaction to the release of the economic numbers is often tied to whether, and by how much, the actual number differs from the consensus forecast. In table 18.2, we compare the performance of the consensus estimate with the results of the economic derivatives auctions, from their first year of operation.

Table 18.2. *Predicting economic outcomes comparing market-aggregated forecasts with consensus surveys, mean absolute error of forecasts*

	Non-farm payrolls (monthly change, 000)	Retail trade (ex-autos) (monthly change, %)	ISM manufacturing purchasing managers' index
Consensus estimate	71.1	0.45	1.10
Economic derivatives	72.2	0.46	1.07
Prediction market sample size	16	12	11

The consensus and market-based estimates of these economic indicators are extremely close – so close that there is no statistically (or economically) meaningful difference in forecast performance. This is true if one examines either correlations with actual outcomes or average forecast errors. That is, in this case the consensus estimate appears to aggregate expert opinion about as well as the prediction market. Even so, this early sample is sufficiently small that precise conclusions are difficult to draw.

18.6 Using prediction markets in decision making

We know we can use prediction markets to make accurate assessments about uncertain future events. We now turn to how to use these predictions to better inform decision making.

The simplest approach is to just use the predictions directly. For instance, in their experiments at Hewlett Packard, Chen and Plott (2002) elicited expectations of future printer sales through a market in which employees bet against each other. These expectations are likely of direct interest for internal planning purposes.

Researchers have also tried to link the time series progression of prediction markets with other variables in order to find the correlation between the two. For instance, prior to the 2004 election, several analysts tried to find a link between the probability of George W. Bush's re-election and the price of the S&P 500. The result is a strong positive correlation between an increase in Bush's chance of re-election and the health of the stock market. While this has been trumpeted as evidence that Bush would be better for the economy than Kerry, this provides a very clear case where correlation does not imply causation. It is just as likely that a strong economy would increase the chances of Bush's re-election as the other way around.

Two further elements are required for a regression analysis to be feasible: (1) time series variation in event probabilities, and (2) a sufficiently strong correlation to allow one to distinguish the relationship from other events affecting the probabilities. By using a slightly different set of contracts, however, it is possible to estimate correlations even when these conditions are not satisfied. For example, we could sell two securities, one which pays \$ P a year from now if Bush is re-elected (where \$ P is the price of the S&P 500 a year from now) and the purchase price is refunded if Kerry is elected, and a second that pays \$ P if Kerry is elected, with the purchase price is refunded if he is not. The difference would be the market's expectation of the relationship between the election of Bush or Kerry and the S&P 500. Of course, while these securities form a contingent market – one that allows us to gauge the market's expectation of one event contingent on another event occurring – they do not resolve the issue of whether this correlation reflects a causal relationship.

Very few contingent markets have been constructed, but they are growing in popularity. In 2004 the IEM offered securities linked both to the two-party vote share of each Democratic candidate and the vote share of Bush if he were to face that particular candidate. These contracts pay nothing if that particular match-up does not occur. These securities can be used to infer the probability that a given Democrat wins the primary, as well as the expected two-party vote share if s/he were to win the nomination. The prices and calculations from two days before the Iowa Caucus appear in table 18.3.

Column (A) shows the price of a contract that pays the Democratic vote share in the general election; the bettor must also pick the Democratic nominee, or the security pays nothing. Column (B) shows the price of a contract that pays Bush's vote share, if the bettor also correctly picks the Democratic nominee and nothing otherwise. The prices thus reflect the market's assessment of both the chance of the candidate winning the Democratic nomination and the share of the vote he would take against Bush.

No matter who the candidate is, the expected Democratic and Republican shares of the two-party vote must sum to one. Thus, adding the prices of the securities shown in columns (A) and (B) yields the probability that each candidate wins the Democratic nomination (shown in column (C)).

A more interesting statistic would be the market's expectation of how each candidate would fare versus Bush if they win the nomination. As suggested by Robin Hanson (1999), this number could then be used to inform the nomination decision of the Democrats, as they presumably would like to nominate someone with a good chance of winning the general election. The calculation is done in column (D). This logic suggests that

Table 18.3. *Contingent markets*

Candidate	Candidate vote-share (A) (\$)	Bush vote-share given this candidate (B) (\$)	Prob. this candidate wins nomination (C) = (A) + (B) (%)	Expected vote-share if nominated (D) = (A)/(C) (%)
Howard Dean	0.289	0.245	53.4	45.9
Wesley Clark	0.101	0.102	20.3	50.2
Richard Gephardt	0.017	0.019	3.6	52.8
John Kerry	0.062	0.067	12.9	51.9
Other Democrats ^a	0.042	0.049	9.1	53.8

Note:

^a By this date, 'Other Democrats' was more or less the same as John Edwards.

Edwards did not have a security tied to him until four days after the Iowa caucuses.

Source: Closing prices, 17 January 2004, IEM.

they should choose Edwards or Kerry as the nominees (Gephardt was already largely out of the running). This implication has led these contingent contracts to sometimes be called 'decision markets'.

We are optimistic that contingent contracts can be used to inform decision making; however, some care must be taken when doing so. There are many plausible stories one could come up with for the reason why the Kerry security is trading higher than Bush|Kerry. For instance, the markets may believe that Kerry won't win the nomination unless the country makes a dramatic shift to the left, but that if this does happen it is likely that Kerry will win both the nomination and the election. Simply nominating Kerry based on these contingent contracts would then be a mistake, since it will not make the country swing to the left, and Kerry would thus be more likely to lose the general election than, say, Edwards.

Irrespective of such issues, the predictions based on these contingent contracts seem to be consistent with subsequent events. On 19 January 2004, Howard Dean lost the Iowa Caucus in spectacular fashion, and that evening self-destructed as he uttered the now infamous 'Dean Scream'. His likelihood of winning the Democratic nomination tumbled from 53.4 per cent to 24.5 per cent by the end of that night. John Kerry, who won that day, saw his probability rise from 12.9 per cent to 25.8 per cent while John

Edwards, who came second, saw his rise from 9.1 per cent to 22 per cent. These candidates were predicted to fare much better against Bush, and accordingly, Bush's expected share of the two-party vote fell – from 52.1 per cent the night before the Iowa caucuses to 48.5 per cent the night after. (In an analogous example, Berg and Rietz, 2003, found that as it became clear in 1996 that Bob Dole would win the Republican primary, Bill Clinton's re-election chances soared.)

We also have preliminary results from an experimental contingent contract we ran on TradeSports. This experimental security paid \$1 if both Bush were re-elected and Osama bin Laden were captured by the election. It seems likely that bin Laden's capture would have a positive effect on Bush's re-election chances, and the markets agree. In mid-June a contract on Bush's re-election was trading at \$57 and an implicit contract on bin Laden being captured by 2 November was trading at around \$27. The joint contract requiring both events to occur was trading for approximately \$21. Using the prices and method above, this tells us that the market assessed the probability of Bush winning if bin Laden were captured at 77 per cent. It also tells us that the market thought that the chance of Bush being re-elected if bin Laden were not captured was 50 per cent.⁴

A cleaner example of the difficulty of untangling correlation and causality comes from a second contract we ran on TradeSports. This contract paid \$1 if Bush won the 2004 election and the terror alert on election day was at its peak level of 'red'. The market put the probability of this occurring at 8.0 per cent, and the probability of red alert on 1 November (the day before the election) at 8.2 per cent. Using these two numbers, we infer that the market believes if the terror alert level is at 'red' then Bush has a 97 per cent chance of winning the election. This estimate seems rather high. There is probably some imprecision due to the problem of miscalibration of small-probability events and the small amount of trading in this market.

If we take this estimate at face value, however, we are confronted with another problem. One explanation might be that the increased threat of terrorism would cause Americans to rally around Bush and re-elect him. However, recall that in Spain in early 2004 a terrorist attack caused the incumbent party to lose the election. If terrorists think a similar thing might happen in the US, we might be tempted to infer that the market believes that if Bush looks strong in the election, this may increase the threat of a terror attack, raising the alert level.

If we were to pass an econometrician data on the likelihood of Bush winning the election and the terror alert level in many states of the world, the econometrician would note a strong correlation between Bush winning

and an elevated terror alert level. However, she would not declare a causal relationship between the two. Instead, she would note that there are 'selection effects' – that is, the states of the world in which the country is on red alert are not random.

Just as an econometrician uses a selection model to correct for selection bias (Heckman, 1979) one can simply add another security or contingency tied to a variable that is driving the terror alert level (such as reports of terrorist activity overseas). If the probability of a certain contingency is high, then only stories that include it are plausible explanations of what will cause a 'red' alert. However, this eliminates scenarios only in a piecemeal fashion, and to the extent that there are an infinite number of possible scenarios involving an infinite number of variables, not all of which are observable, it will never be possible to absolutely pin down causation.

The preceding paragraphs may make it sound as if there are extreme difficulties with prediction markets that make their use in this domain hopeless. However, the difficulties here are no different than those in any other econometric situation. These issues should be the topic of further research and application. In the meantime, simple prediction markets continue to be extremely useful for estimating the market's expectation of moments or distributions – even multivariate ones.

18.7 Looking forward

This chapter has focused jointly on the promise and the limitations of prediction markets. While these markets manifest the pathologies of all financial markets more deeply, it is important to keep in mind that they also outperform many other prediction tools, often at lower cost. One's optimism about the further use of prediction markets in business, government and finance depends a lot on what sorts of mechanisms for prediction one is comparing the market-generated prices with.

Furthermore, there is a broad pool of research into more common financial markets that has not yet been applied to these markets. Currently the level of sophistication of prediction markets in practice is such that they can be understood using very basic financial tools and rules of thumb. As these markets prove themselves and become better capitalised there will be an incentive to apply more advanced methodologies to their execution. This in turn will lead to more effective and efficient markets that will embody fewer of the problems we have outlined and allow for true hedging against geopolitical and other risks.

We have also focused on an emerging, more complex form of markets that try to predict the probability of multiple events happening

simultaneously. These contingent contracts, or 'decision markets', can be used in conjunction with simpler securities to tease out the market's perception of factors important to public decisions. As we note, there are difficulties in separating correlation from causality but, carefully applied, we believe that there are domains in which these markets will be useful public policy inputs.

Prediction markets are, at their core, a tool for deriving consensus estimates and assessments from a diverse body of people and opinions. To the extent that there exist questions that are important enough to generate interest, and thus liquidity, prediction markets may be used to replace or augment more primitive technologies such as frequent meetings or arbitrary algorithms.

Notes

The authors would like to thank David Pennock, Emile Servan-Schreiber of NewsFutures, David Dempsey and John Delaney of TradeSports, David Siegel and Scott Hereld of Trackmaster and George Neumann of IEM for help with data. Thanks to Kay-Yut Chen, Andrew Leigh, Rohan Wolfers, Betsey Stevenson, Tim Taylor, Hal Varian and Craig Yee for stimulating discussions. Brett Danaher, Doug Geysler, Chris Lion, Paul Reist and Ravi Pillai provided outstanding research assistance. Justin Wolfers would like to acknowledge the financial support of the Hirtle, Callaghan & Co. – Arthur D. Miltenberger Research Fellowship.

- 1 The price of a winner-take-all security is essentially a state price, which will equal an estimate of the event's probability under the assumption of risk-neutrality. The sums wagered in prediction markets are typically small enough that assuming that investors are not averse to the idiosyncratic risk involved seems reasonable. But if the event in question is correlated with investors' marginal utility of wealth, then probabilities and state prices can differ. In what follows, we leave this issue aside and use the term 'probability' to refer to risk-neutral probability. For more on this topic, see Wolfers and Zitzewitz (2005).
- 2 There is a subtle, an almost metaphysical question here: What is the 'market's' expectation anyway? Throughout, we will speak as though the market is itself a representative person, and that 'person' has a set of expectations. Consequently there are important but subtle differences between parameters such as the market's median expectation and the median expectation of market participants.
- 3 Rhode and Strumpf (2004) investigate turn-of-the-twentieth-century markets that were used to predict the outcomes of presidential elections. If you see sports gambling as a rudimentary form of prediction markets, then obviously prediction markets are quite a bit older.
- 4 This last figure can be calculated by Bayes' Rule: $(57-21)/(100-27) \approx 50$ per cent.

References

- Aït-Sahalia, Yacine, Wang, Yubo and Yared, Francis (2001) 'Do Options Markets Correctly Price the Probabilities of Movement of the Underlying Asset?', *Journal of Econometrics*, 102, pp. 67–110
- Athanasoulis, Stefano, Shiller, Robert and van Wincoop, Eric (1999) 'Macro Markets and Financial Security', *Economic Policy Review*, 5, pp. 21–39
- Berg, Joyce, Forsythe, Robert, Nelson, Forrest and Rietz, Thomas (2001) 'Results from a Dozen Years of Election Futures Markets Research', in Charles Plott and Vernon Smith (eds.), *Handbook of Experimental Economic Results*, New York: Elsevier Science
- Berg, Joyce and Rietz, Thomas (2003) 'Prediction Markets as Decision Support Systems', *Information Systems Frontiers*, 5(1), pp. 79–93
- Camerer, Colin (1998) 'Can Asset Markets be Manipulated? A Field Experiment with Racetrack Betting', *Journal of Political Economy*, 106(3), pp. 457–82
- Chen, Kay-Yut and Plott, Charles (2002) 'Information Aggregation Mechanisms: Concept, Design and Field Implementation for a Sales Forecasting Problem', Social Science Working paper, 1131, Pasadena: California Institute of Technology
- Hanson, Robin (1999) 'Decision Markets', *IEEE Intelligent Systems*, 14(3), pp. 16–19
- Heckman, James J. (1979) 'Sample Selection Bias as a Specification Error', *Econometrica*, 47(1), pp. 153–61
- Leigh, Andrew, Wolfers, Justin and Zitzewitz, Eric (2003) 'What do Financial Markets Think of War in Iraq?', NBER Working Paper, 9587
- Manski, Charles (2004) 'Interpreting the Predictions of Prediction Markets', NBER Working Paper, 10359, March
- Ortner, Gerhard (1998) 'Forecasting Markets – An Industrial Application', Technical University of Vienna, mimeo
- Pennock, David, Lawrence, Steve, Giles, C. Lee and Nielsen, Finn Arup (2001) 'The Real Power of Artificial Markets', *Science*, 291, pp. 987–8
- Rhode, Paul and Strumpf, Koleman (2004) 'Historical Prediction Markets: Wagering on Presidential Elections', *Journal of Economic Perspectives*, 18(2), pp. 127–42
- Servan-Schreiber, Emile, Wolfers, Justin, Pennock, David and Galebach, Brian (2004) 'Prediction Markets: Does Money Matter?', *Electronic Markets*, 14(3), pp. 243–51
- Shiller, Robert (2003) *The New Financial Order: Risk in the Twenty-First Century*, Princeton: Princeton University Press
- Snowberg, Erik and Wolfers, Justin (2004) 'Understanding the Favorite-Longshot Bias: Risk Preferences versus Misperceptions', University of Pennsylvania, mimeo
- Strumpf, Koleman (2004) 'Manipulating the Iowa Political Stock Market', University of North Carolina, mimeo
- Tetlock, Paul (2004) 'How Efficient are Information Markets? Evidence from an Online Exchange', University of Texas at Austin, mimeo

- Wolfers, Justin and Andrew, Leigh (2002) 'Three Tools for Forecasting Federal Elections: Lessons from 2001', *Australian Journal of Political Science*, 37(2), pp. 223–40
- Wolfers, Justin and Zitzewitz, Eric (2004) 'Prediction Markets', *Journal of Economic Perspectives*, 18(2), pp. 107–26
- (2005) 'Interpreting Prediction Market Prices as Probabilities', University of Pennsylvania, mimeo