5 Information markets for decision-making

Performance and feasibility

*Karen Croxson**

1 Introduction

The success of any organization hinges on its ability to make wise decisions in an uncertain environment. When, where and how should scarce resources be committed to achieve institutional objectives? A product manager must devise a strategy commensurate with the profit objective. Politicians must decide how best to allocate the public purse to achieve policy goals. A disaster relief agency must make tough decisions as to where to concentrate efforts and supplies on the ground. Making the right decision requires weighing the effectiveness of alternative actions based on relevant information, which often is widely dispersed. It can be challenging to identify genuine insights and somehow aggregate these into a meaningful intelligence. Until recently, decision-makers have tended to rely on a number of traditional approaches to informing their decisions, including consulting experts, appointing groups to deliberate and conducting polls and surveys. We can think of each as an information mechanism – with its merits, but also recognized shortcomings. For instance, 'groupthink' tends to bias the views of committees; those we approach as 'experts' may not be true experts (and may not give honest opinions); a simple poll fails to weight individual opinions according to how informative they actually are.

In recent years a band of innovative organizations, including Hewlett Packard, Google, General Electric and Microsoft, has experimented with a new mechanism – the information market. Information markets essentially are in-house betting markets. They are designed specifically to tap the knowledge of a dispersed 'crowd' and aggregate this into an accurate forecast. Participants in an information market buy and sell assets whose payoffs are tied to the realization of a future event, perhaps the effectiveness of a current sales campaign or even the success of a merger. The market price is interpreted as a collective prediction. There are sound theoretical reasons to expect this forecast to be more reliable than forecasts provided by alternative methods, by getting participants to 'put their money where their mouth is' without fear of recrimination (thanks to anonymous trading), a market addresses many of the difficult incentive problems that can undermine more traditional information mechanisms. Public betting markets

have a rich tradition in forecasting political elections and sports events, performing impressively against polls and pundits (Berg *et al.*, 2008; Vaughan Williams, 2005); many laboratory experiments have demonstrated the ability of markets to aggregate the information of traders effectively (Plott and Sunder, 1982, 1988; Forsythe *et al.*, 1982; O'Brien and Srivastava, 1991), and where information markets have been applied inside corporations they appear to have delivered some encouraging results (Ortner, 1998; Plott and Chen, 2002; Cowgill *et al.*, 2009).

This combination of theoretical appeal and promising empirical evidence has led to speculation that in-house information markets will revolutionize decision-making across private and public sector organizations.[1] Meanwhile, the business columns of international magazines and newspapers have led many high-profile discussions of the potential for the 'corporate prediction market' to become a killer Web 2.0 application.[2] Despite all the enthusiasm and expectation, real-world utilization of information markets remains far behind the hype. The vast majority of the business world is showing little interest in the internal market as a serious alternative (or even complement) to standard forecasting methods, even for more pedestrian business matters. Moreover, even among the handful of firms embracing the technology most enthusiastically, none is known to use the mechanism to guide major strategic decisions. The public sector, meanwhile, seems barely to have begun to experiment with the concept, despite the potentially huge wins from 'policy markets'.[3] This chapter reviews some of the evidence surrounding the comparative performance and feasibility of information markets, discusses some of the practical challenges inhibiting their mainstream deployment inside organizations, and touches on priorities for future work.

2 Traditional information mechanisms

2.1 Expert opinion

A time-honoured approach to informing decisions is to seek out an expert opinion. Appealing to expertise has intuitive appeal; most of us would prefer to act on the medical advice of a qualified doctor than rely on the judgement of our friends or neighbours. There are difficulties with the expert model, however. Whereas in some cases it may be trivial to identify the right person – their background or specialist training might mark them out for the role – in many other cases identifying a true expert is challenging, and those who appear to be expert in a matter may possess little actual expertise. Tetlock (2005) describes a 20-year study in which several hundred experts drawn from many fields, from professors to journalists, were asked to make thousands of predictions about the future. The predictions turned out to be only marginally more accurate than chance, and the most recognized experts – those who advise governments and business leaders, appear on television and are regularly quoted in newspapers – performed particularly poorly. Even when a real expert has been located, simply asking for an

opinion may not be the best way to ensure an honest response. Experts can suffer conflicts of interest: they may desire to tell us what they think we'd like to hear (Prendergast, 1993), or feel it is safer to herd with other experts (to avoid being wrong when everyone else is right).[4] A narrow focus on conventional circles of expertise might cause valuable knowledge to be missed; recent work has shown that official forecasts of unemployment can be improved using insights from internet search data (D'Amuri and Marcucci, 2009). Similar results have emerged for the detection of influenza (Ginsberg *et al.*, 2009; Polgreen *et al.*, 2007). Instead of relying on a single, potentially biased, opinion we might think about consulting several experts and somehow pooling their opinions. This may go some way towards mitigating biases, although it still leaves open the problem of distinguishing true experts. It also raises the non-trivial question of which pooling mechanism should be used.[5]

2.2 Group deliberation

Appointing a group to deliberate would appear to provide an attractive basis for decision-making: gather together individuals with relevant perspectives, encourage them to engage with each other, share insights, resolve differences, and so converge on an enlightened collective viewpoint. Certainly, meetings are used heavily inside most organizations. But in practice, important limitations can arise, particularly where meetings are face-to-face. Minorities with relevant information tend to be silenced too readily by social pressure, and junior members of a meeting can be reluctant to say what they really think. A troubling 'groupthink' can emerge,[6] a lack of anonymity can inhibit honesty in communication, and it is often necessary to exclude some who may have relevant information – a many-way discussion can become quite unmanageable as the number of people involved grows large.

2.3 Polls and surveys

For some issues, it might be useful to conduct a poll across the organization. Polls can provide a quick snapshot of broad opinion, which could be informative, particularly if the poll is anonymous. Instant polls are becoming an everyday part of life in some organizations, spurred by the development of handheld voting devices. These devices were first used in the 1960s by motion picture and television studios as a way to gauge audience responses to unreleased films, TV shows and commercials. Nowadays, they are used inside companies to take the pulse of opinion during meetings.[7] A theoretical drawback with polls is that, by default, they weigh individual opinion equally, and yet people may be quite differently informed. A potential response would be to assign different weights to individual opinions to reflect their relative expertise, but in this case the problem of identifying experts resurfaces. Polls also fail to reward people when their opinion is accurate and so may inhibit participation and truthful revelation of information. Surveys can be used to give a more nuanced view of opinions, but

they tend to suffer from similar weaknesses: they fail to reward individuals for reporting their insights truthfully, and they still leave the problem of how to interpret and weight the responses. Surveys also need to be designed, and their implementation can be very costly.

3 The promise of information markets

Information markets are betting markets established specifically to aggregate dispersed information into a collective forecast. They go by many other names, including prediction markets, event futures, event derivatives and virtual stock exchanges. Participants in an information market are traders – they buy and sell contracts which have payoffs tied to the realization of a future event. Suppose a computer games executive wishes to know whether a new game will ship on time. She could set up an internal information market to help her estimate whether things are on track. After she has procured some simple software to run her market (which is available off-the-shelf from many specialist technology providers, or might be built in-house), she defines the contract 'Product X will ship on 1 November 2011 at the latest', she stipulates that this contract will pay out $1 (which could be virtual currency) if the designated success occurs and $0 otherwise. Their next step is to invite a group of participants to trade the contracts, perhaps providing each with an initial endowment. Employees from around the firm might be invited to trade, or only those closely acquainted with Product X. To kick things off, the executive might set the initial contract price at $0.50, implying a 50:50 chance that the product ships on time. Some of the participants may find this belief too pessimistic based on what they know and so will be incentivized to buy contracts in the market. Others may feel even more pessimistic about delivery than the current price implies and so will wish to sell contracts. As individuals trade in the market the contract price will move around to reflect their information. After a few weeks the contract might be priced at $0.80, from which the product manager should infer that the chances that Product X ships on time are around 80%.

There are good reasons to think that a forecast generated in this way could be highly accurate - potentially more reliable than a forecast generated in any other way. Economists have long recognized that a byproduct of markets designed for speculation and hedging is that they can aggregate private information effectively (Hayek, 1945):[8]

- Markets enable insights to be gathered from a diverse crowd of participants, and diversity has been shown to be important for problem solving (Surowiecki, 2004; Page, 2007).
- Traders can submit their views independently and anonymously in a market – unlike in deliberative group settings.
- Markets provide the right incentives to participants to reveal their information quickly and truthfully – those who have information and are first to trade on this can profit.

- Markets provide incentives for research and information discovery – individuals who actively acquire information can earn additional profits through trading.
- The market provides an algorithm for aggregating opinions into a single collective viewpoint.

Markets offer predictions that update 'in-running'. They have been described as a 'pull' forecasting mechanism, as compared to the 'push' mechanisms offered by polls and surveys.[9] As new information emerges, traders in prediction markets have the incentive and opportunity to react quickly by changing their positions, and so market price should move rapidly to incorporate any news.

Recent years have seen the creation of many public prediction markets, dealing with election results through to the winners of reality-TV contests. Forecast performance has been impressive, with market predictions routinely topping those of professional forecasters and polls in head-to-head contests. Perhaps the best known information markets are the Iowa Political Markets, created by academics at the University of Iowa in the late 1980s. They allow the public to bet on political outcomes such as US presidential elections. The predictions from Iowa markets have beaten opinion polls and political pundits remarkably consistently over the years (Forsythe *et al.*, 1992; Berg *et al.*, 2008). Following this success, Iowa researchers have developed markets to forecast outbreaks of avian flu and to predict the Federal Reserve's monetary policy. Elsewhere, Hollywood markets invite the public to predict opening weekend box office sales and pick Oscar winners – anyone can go to the Hollywood Stock Exchange website (www.hsx.com), sign up for free, and automatically collect an endowment of Hollywood dollars to buy and sell films and actors. Around 1.8 million people are now registered to trade and the predictions from these markets tend to be more accurate than those of film critics, even though only play-money is at stake (Pennock *et al.*, 2001; Spann and Skiera, 2003). Meanwhile, real-money prediction markets have existed for sporting events for some time. The online betting exchange Betfair (www.betfair.com) is the world's largest prediction exchange. It is reported to have over two million members and offers real-money markets related to thousands of sports events (and many other events of popular interest). Studies to date have reported encouraging evidence regarding the efficiency of the exchange's prices.[10]

The earliest example of corporate information markets are thought to be the markets academic economist Robin Hanson helped establish at technology provider Xanadu in 1990. One of the claims traded related to the delivery date of the firm's product: 'Xanadu will deliver its product before Premier Deng of China dies.' Deng died before the product was delivered.[11] Since then, many more businesses have begun to experiment with their use internally. The markets seem to be used primarily to forecast such matters as whether a project deadline will be met, whether a sales target will be exceeded, or what a competitor will do. Typically, a relatively small group of employees is invited to trade, and often play-money and prizes are the only incentives offered. The results reported so

far are encouraging. An early pioneer in this area was Hewlett Packard. It began experimenting with internal markets in the late 1990s, leveraging its crowd of employees to forecast printer sales. The forecasts from HP's information markets beat the forecasts provided by its sales department – the 'experts' (Plott and Chen, 2002). Google seems to have conducted the largest experiment with corporate prediction markets to date, launching its first markets in 2005 and going on to deploy thousands internally. Its markets are used to forecast product launch dates and a range of other outcomes. Cowgill *et al.* (2009) analyse data from Google's markets for the period 2005–2007. Though they find evidence of biases on the part of traders (such as optimism bias, which seems strongest in newer recruits), they report that markets are reasonably efficient and become less biased over the study period as collective trading experience increases. Microsoft has explored prediction markets internally, as have retailer Best Buy and several other leading firms, including major pharmaceutical players Pfizer, Novartis, GSK and Eli Lilly.

4 Barriers to utilization

Surprised by the unfulfilled potential of information markets, James Surowiecki wrote in his 2004 best-seller, *the wisdom of crowds*, that: '...the most mystifying thing about markets is how little interest corporate America has shown in them ... companies have remained, for the most part, indifferent to this source of potentially excellent information (Surowiecki, 2004: 21–22).' By popularizing the notion of collective wisdom, Surowiecki himself did much to acquaint real-world decision-makers with the idea of information markets, leading to a leap in awareness among executives. In the six years since then, numerous articles in leading business magazines and the popular press have addressed the promise of internal markets, ensuring that their public profile remains high. Despite all this, there remains little evidence to suggest that in-house markets are becoming mainstream. Over the past few years, consultancy McKinsey has surveyed executives from a broad range of industries, regions and functional specialties about their usage of Web 2.0 applications such as blogs, podcasts, social networks and prediction markets. In the 2009 survey, only a small minority of the approximately 1,700 respondents reported that prediction markets were deployed inside their organizations – 9 per cent of executives, largely unchanged from the previous year (McKinsey, 2009).[12] Just under half of these described their institution's use of markets as 'evaluating or running limited trials' (as opposed to 'using it in our business') and around half were only, at most, somewhat satisfied with the technology.[13] Focusing on government applications, the evidence is even less encouraging. There are few signs of experimentation with prediction markets within public sector organizations, let alone serious adoption.[14] In comparison to the media hype surrounding their potential to revolutionize decision-making, it would appear that prediction markets are experiencing a peculiar 'failure to launch'.

Will we ever see organizations leverage markets to guide their large-scale strategic bets, such as whether to merge, sack the current CEO or push into a

new geographic territory? Will their deployment for more pedestrian purposes ever become routine? In the next section, we discuss some of the design challenges surrounding the application of internal markets and other practical impediments to their proliferation inside organizations. We consider where and how some of these issues might be mitigated.

4.1 Restrictions on applicability

Information markets cannot be deployed to address every issue of interest. Compared to more traditional information mechanisms, more onerous conditions must be satisfied for a situation to be amenable to the use of markets. Specifically, it must be feasible to:

- define an uncertain event unambiguously;
- write contracts related to a (small) number of mutually exclusive and mutual exhaustive possible outcomes;
- resolve uncertainty by a specified point in time (ideally not too far into the future as long-term markets tend to be less motivating);
- settle the market against objective criteria.

This means that rather than listing a contract with the wording 'Weapons of mass destruction are not in Iraq', which specifies no end point at which the bet can be settled, contracts of the form 'WMD will have been found by date Z' should be used. Nor can the definition of the event change once the market is in progress: Ortner (1998) describes an internal prediction market used to forecast whether a software project would be delivered to the client on schedule. At some stage, the client changed the deadline, creating problems for the operation of the market. Markets will only work well only when there is reasonable clarity about and confidence in the contract design. Clearly, not all questions of interest will be compatible with the above conditions, whereas more traditional information mechanisms can be applied to less well-defined issues, including matters of opinion. Even so, a vast range of organizational issues will be amenable to markets, including potentially many 'big-ticket' questions such as: 'Will medical device Y be recalled by the end of this quarter?', 'Will our rivals GHI Inc. and JKL Inc. merge in the next six months?' or 'Will our health reform lower hospital admittances by Z% this year?'.

4.2 Limited empirical evidence

At least some of the reluctance to implement information markets can be explained by the lack of clear-cut evidence on their performance against alternatives. A related issue is confusion and uncertainty related to design choices. The number of academic articles on the topic of information markets has grown rapidly in recent years,[15] but there remains limited clarity surrounding what should be expected from markets in specific real-world settings and how exactly markets should be configured.[16]

A number of studies have confirmed the ability of information markets to aggregate information effectively in the laboratory (Plott and Sunder, 1982, 1988; Forsythe *et al.*, 1982; O'Brien and Srivastava, 1991), but other experiments have demonstrated the existence of problems such as bubbles and false equilibria (Smith *et al.*, 1988; Camerer and Weigelt, 1991; Noeth *et al.*, 1999; Lei *et al.*, 2001; Hussam *et al.*, 2008). While laboratory settings enable the experimenter to assign private information to individual participants, allowing the performance of information mechanisms to be tested cleanly, they come with an obvious downside: participants are typically inexperienced students confronted by abstract problems. As such, it is unclear whether the laboratory performance of information mechanisms extrapolates to corporate or policy contexts.[17]

Other researchers have investigated performance in the field (Forsythe *et al.*, 1992; Chen *et al.*, 2005; Christiansen, 2007; Cowgill *et al.*, 2009), and some of these have compared the accuracy of market forecasts to those generated by other information mechanisms. For instance, Chen *et al.*, (2005) analyse predictions from two public information markets for NFL games played in 2003, comparing these to 'expert' opinion pools for the same events. The opinion pools are constructed using linear and logarithmic aggregation functions to combine the subjective probability judgements of 2,000 self-identified experts. The authors find that the predictions from information markets are as accurate as the pooled expert assessments for the same time-point ahead of the game. But there are limitations to studies that compare mechanisms run side-by-side. When the mechanisms studied estimate the same outcomes (and may potentially even share some participants), this undermines the potential for clean inference regarding the efficacy of either mechanism. Information may leak from one mechanism to the other in a way unobserved by the researcher. For greater clarity, future work should pursue more systematic comparison of mechanisms in real-world settings. Graefe (2009) provides a useful recent contribution. In a field experiment the author compares the Delphi method of structured group deliberation with prediction markets, taking care to ensure that none of the market participants concurrently took part in the Delphi mechanism. Prediction markets are found to work as well as the Delphi method.

Difficult design choices have to be made in setting up an information market, but despite papers such as Spann and Skiera (2003), still little is understood about the impact of market design features on information accuracy. Two important open questions relate to participation and incentives.[18]

Participation

If active traders are few and far between, the market may become too thin to yield accurate predictions. However, recent research has suggested that, depending on the market mechanism implemented, the lower bound on numbers may be quite low (Christiansen, 2007; Chen *et al.*, 2001).

It might seem intuitive that participation should be restricted to those with some expertise about the subject matter. Assuming experts can be reliably identified, including non-experts in many conventional information mechanisms, such as meetings, would seem to add little obvious value. Prediction markets potentially are very different in this respect. The presence of uninformed traders (referred to as 'noise traders' in the finance literature) may be necessary for the market to be viable (Wolfers and Zitzewitz, 2006) and may actually increase the incentives for those with information to bring this to the market since the presence of uninformed others implies an opportunity to trade profitably on knowledge.[19] At the same time, the jury is still out on just how much non-informed trading is desirable: while some empirical studies have lent support to the view that securities mispricing is greater in illiquid markets (Kumar and Lee, 2006; Sadka and Scherbina, 2007; Chordia *et al.*, 2008), in the theoretical work of De Long *et al.* (1990), liquidity is a proxy for non-informational trading (noise trading), which may harm informational efficiency, and some recent empirical analysis of public prediction markets appears to support the idea that greater liquidity can worsen mispricing (Tetlock, 2008; Hartzmark and Solomon, 2010). Clearly, further research is needed in this area, with more investigation of the impact of non-informed trading in internal applications. Would using individual market trading performance as an indicator of expertise and then pooling 'expert' opinions yield a more informative mechanism? Work by Chen *et al.* (2001) suggests that it might.

In practice, those setting up prediction markets have taken a mix of approaches to the question of who to include. When creating the Iowa Election Markets in 1988, academics at Iowa University opted to open these markets to the general public. However, when creating their first Iowa Health Market to predict seasonal influenza (in 2004), it was decided to restrict participation to registered members of the medical community. The University of Iowa now runs two types of market: public (play money) and private (real money). Applications to play its private markets, which include many health markets, are reviewed and approved by market managers on a case-by-case basis, with those viewed as the best possible traders selected to play.[20] When retailer Best Buy experimented with internal information markets for forecasting business outcomes such as sales, the company was careful to include a wide base of participants with diverse operational knowledge, and not simply in-house forecasting experts. This approach respects the finding that sufficient cognitive diversity is important for a crowd to be wise (Page, 2007).[21]

Incentives

One of the biggest potential advantages associated with prediction markets is that participants are forced to 'put their money where their mouth is'. Compared to other information mechanisms, this should reduce biased messages. Establishing a real-money information market raises legal, financial and ethical issues, however; in many jurisdictions, including the United States, gambling is heavily restricted.[22] A group of prominent academics has called publicly for regulators to make

provisions for real-money prediction markets on the grounds of social interest (Arrow *et al.*, 2007), but for now most companies experimenting with in-house information markets have chosen to proceed cautiously using virtual money and occasional prizes. Many economists would expect a play-money market to provide weaker incentives for information acquisition and truthful revelation since traders have less at stake. A large number of laboratory studies, beginning with Siegel (1961), have supported the idea that real money is required to ensure truth-revealing incentives in experiments. However, promising results from several play-money markets inside corporations have led some observers to suggest that participants derive sufficient motivation from the pursuit of 'winner status'.[23]

A few recent investigations explore the link between incentives and prediction market performance (Servan-Schreiber *et al.*, 2004; Rosenbloom and Notz, 2006; Luckner, 2007; Gruca *et al.*, 2008), but the evidence remains inconclusive. Servan-Schreiber *et al.* (2004) compare the predictions of TradeSports, a real-money market, and NewsFutures, a play-money market, finding that the play-money markets performed as well as the real-money markets. Rosenbloom and Notz (2006) also compared the predictions from TradeSports with those from NewsFutures; they report that there was little to separate the predictions for NHL games, but that the real-money market delivered more accurate predictions for non-sports events. The recent study by Gruca *et al.* (2008) investigates the impact of incentives on prediction for movie success. The authors compare predictions from the play-money market, Hollywood Stock Exchange, with those from the real-money Iowa Electronic Market for the same movies. Contrary to Rosenbloom and Notz (2006), they find no statistically significant difference between the accuracies of the two markets for these non-sports events. As noted previously, however, a limitation of studies which compare predictions from markets for the same events is the possibility for information leakage across mechanisms. This constrains the ability to deliver clean inference. Rosenbloom and Notz (2006) have suggested that a partial explanation for their results could be the existence of publicly available betting odds for the sports events they study; these odds could have helped traders in both the real-money and play-money markets. The impact of incentives in information markets deserves further research, ideally involving randomized trials inside organizations. In the context of enterprise information markets, it will be important also to consider how incentives compare to salaries, awards or other workplace incentives.

4.3 Sensitivity of market information

Perhaps the biggest barrier to the mainstream deployment of internal corporate markets relates to the sensitivity of the information generated. Managers may have concerns about the release of information regarding:

- the nature of the organization's problem (the fact that the firm is seeking product ideas, improved sales forecasts or entertaining the idea of a merger);
- the prevailing market forecast.

In their default implementation, prediction markets reveal to all participants an uncensored 'in-running' (and potentially very accurate) forecast related to the event of interest. However, this disclosure often sits uncomfortably with an institutional desire to bias official forecasts and potentially to conceal these from employees. In many interesting cases, it may simply be unacceptable to have an unbiased metric about sensitive organizational outcomes visible to regular employees this may impact organizational morale adversely or be leaked to the stock market in violation of insider trading provisions. A third worry is that sensitive intelligence might reach competitors.[24] We expand briefly below on the first two of these concerns.

Motivation and morale

Targets may be set to stretch employees and are hence potentially unrealistic by design. Evidence from academic studies lends support to this approach: psychologists have demonstrated the existence of the Pygmalion effect, which refers to the tendency for actual performance to converge to the positive expectation held by others (Eden, 1984).[25] Equally, morale may be sensitive to expectations about external market conditions or a firm's ability to retain its key staff. Consider an internal market set up to forecast employee retention which is predicting an exodus of staff over the coming quarter. A visible prediction of this nature could become a destabilizing focus within the organization, undermining the ability of leaders to manage.[26]

Complications related to 'insider trading'

Some managers have identified insider trading rules as an important factor discouraging them from using information markets to support 'big ticket' business decisions, such as whether to merge, introduce new products or expand into a new territory (Hanson, 2008). If employees receive key corporate information that is not publicly available, they become 'insiders' in a legal sense. If they were to trade on this information in the market for the company's stock, they would engage in 'insider trading' and civil or criminal sanctions might follow. Mat Fogarty, CEO of prediction markets provider Crowdcast, explains:

> The concern is that the information coming from the PM is so powerful that all recipients of that data are made insiders. Also, if management is aware of any 'material' information, it should disclose this to its shareholders. PMs are designed to deliver plenty of 'material' information.[27]

Bell (2009) has clarified the legal issues surrounding prediction markets and current insider trading provisions and suggests practical actions firms might take to mitigate problems associated with this potentially thorny issue. One idea is to create separate markets that could be traded only by officers of

the company and other existing 'insiders', so that the company's existing framework of controls for inside information could be applied. Another suggestion is that the firm admonish participants in its private markets against trading on the information generated. Additionally, participants could be notified that claims and prices constitute the corporation's 'trade secrets', which, according to Bell, would give the firm a misappropriation claim against anyone who trades the company's shares based on its internal market information (rather than a joint liability for their insider trading). An interesting fourth suggestion is that firms might seed their internal information markets with a number of fake claims and prices, so that participants cannot tell which are real. Only traditional insiders would be told which claims were genuine and which were decoys. Looking beyond the current legal regime, Hanson (2008) has argued for changes to existing insider trading rules to allow firms to explore the gains from prediction markets more freely while still preserving the benefits traditionally associated with the existing rules that limit insider trading (the encouragement of investment in public corporations).

Reflecting the above concerns, some corporate adopters of information markets have shown interest in the possibility of preventing employees from viewing the current market price. Crowdcast, a commercial provider of prediction market solutions, has developed a system of 'blind betting', in which participants submit their predictions without observing the current market consensus.[28]

4.4 Fears about manipulation

A concern often raised in the context of discussions about information markets is the possibility of manipulation (Wolfers and Zitzewitz, 2004). The fear is that participants might be tempted to manipulate either the event outcome itself (in order to profit in the market) or market prices (in order to influence a real-world decision that is to be informed by the market price).

Manipulation of event outcomes

In some applications, the outcome of interest is exogenous to those trading the prediction market, as when members of the public bet on the outcome of a soccer match, or company employees speculate on the actions of a competitor firm. But in many interesting real-world implementations, those trading a prediction market will themselves have some influence over the outcome of interest. An obvious case arises where a small group of employees working on a particular project is invited to predict whether the project will be completed on time. The fact of their participation in the information market may create perverse incentives for employees to manipulate the outcome of the project. See Hanson (2006b) and Wolfers and Zitzewitz (2006) for informal discussions of outcome manipulation. A first formal analysis is provided by Ottaviani and

Sorensen (2006). They show how, in theory, outcomes might be manipulated by participants in corporate prediction markets.

Manipulation of market prices

An additional potential worry is that participants might feel tempted to manipulate prediction market prices, particularly where high-stakes decisions are to be based on these (Wolfers and Zitzewitz, 2006). This scenario has been analysed theoretically by Hanson and Oprea (2007) and Hahn and Tetlock (2007). Several empirical studies suggest manipulation is likely to be ineffective (Rhode and Strumpf, 2008; Hanson *et al.*, 2006). For instance, Rhode and Strumpf (2008) analyse an attempt to manipulate the price of a Kerry victory on the public market TradeSports in 2004, as well as their own attempts to manipulate prices on the Iowa Electronic Markets in 2000. These manipulation efforts seem to have had only a very short-lived effect on prices. It has even been suggested that manipulation may enhance market accuracy (Hanson and Oprea, 2007). Price manipulation remains an ongoing area of research with some recent investigations pointing to its potential effectiveness in particular settings (Veiga and Vorsatz, 2009, 2010).

In his book, *Predictocracy*, Michael Abramowicz (2007), offers the following advice on how to safeguard against market manipulation in practical settings:

> First, where there is a discrete group of potential manipulators, those individuals can be barred from participation. Of course, there is always a danger that these potential manipulators can pay off other market participants, but legal or contractual sanctions can reduce that possibility. Second, prediction markets might be limited to a group of authorized traders who are believed to have no incentive to manipulate the outcome.[29]

Others have suggested that incentives associated with the market should be kept below those associated with achieving the outcome of interest. That is, the personal return to ensuring that organizational goals are met should comfortably exceed any reward attainable from betting against these in the prediction market.

Some of those engaged with the practical deployment of information markets inside organizations appear to take the view that the fears surrounding manipulation are overblown, at least for play-money markets.[30]

4.5 Difficulties sustaining participation

Users often report difficulties with the trading interface. This needs to be kept as simple as possible. Some participants are simply uncomfortable with the trading metaphor, and the idea of translating their knowledge into a price (Green *et al.*, 2007). Adam Siegel at Inkling, another platform provider, has observed a shift away from trading screens that look like stock exchanges:

> When we first launched Inkling, we were greeted with great skepticism because our application looked nothing like a stock trading platform. Now if you look at the newcomers in the space, they all try to highlight 'ease of use' as their differentiator.[31]

Once participants have mastered how to play an information, the challenge becomes how to sustain their interest. One strategy is to offer random prizes for participation, as well as winner prizes. The utilization of information markets has so far been almost exclusively confined to forecasting near-term events, and it has been suggested that markets, at least in their current incarnation, are not well-suited to forecasting long-term outcomes because of the difficulties sustaining participation:

> Incentives lose power if the payoff is too remote, and feedback is important for driving participation and performance. Forecasting a result within a few quarters seems to work, but over a year begins to feel like a stretch. We are experimenting with alternative market structures that might help forecast the distant future while paying incentives more quickly.
>
> (Hopman, 2007)

4.6 Perceived lack of legitimacy

Finally, information markets may fail because of a lack of perceived legitimacy. Sunstein (2006b) notes that information markets remain relatively unfamiliar and suggests that their use might breed confusion and distrust. By contrast, deliberation has been found to increase confidence and decrease variance in the group's prediction, which gives deliberative groups greater perceived legitimacy. Recent work by Graefe (2009) has sought to compare the acceptability of different information mechanisms in laboratory experiments. Participants were observed to discount market results more than those from other forecasting mechanisms, despite the fact that this harmed accuracy. Some may have a simple lack of faith in market predictions, whereas others may feel positively threatened by their arrival: information markets throw into question the role of in-house experts, and the forecasts generated may temper executive decisions, thereby challenging existing hierarchical structures. For these reasons, management and employees may be hostile to the introduction of markets and these may fail to become integrated into the normal workflow. When this happens, time spent trading may be seen as a distraction from proper work duties, rather than a valuable exercise. In GE's experimental idea futures markets, limits were imposed on trading hours; to ensure that the market did not interfere with regular work time, participants were asked to trade before or after work, during lunch or for only a few minutes at a time during work hours (LaComb *et al.*, 2007).

Jeff Severts, who introduced prediction markets at US consumer electronics retailer Best-Buy, considers it vital to secure executive buy-in for such initiatives:

> support from senior executives is essential if you want to issue contracts on anything that might be controversial. 'Air cover' is a must or you'll find yourself trading on what kind of casserole we're having in the cafeteria on Thursday.
>
> (Dye, 2008)

One way to improve openness of management to information markets might be through the education of current and future executives. Business schools could play a greater role in raising awareness of the problem of aggregating dispersed information inside organizations and the merits and limitations of mechanisms to achieve this. Andrew McAfee of Harvard Business School has incorporated examples of prediction markets into some of his teaching to show executives how Web 2.0 can be used to tap collective intelligence.[32] Daphne Raban and Dorit Geifman have experimented with prediction markets in MBA teaching at the University of Haifa's School of Management, exposing students to the topic through web-based demonstrations. Their recent article (Raban and Geifman, 2009) discusses the pedagogical value of this approach.

5 Concluding remarks

Recent years have seen active experimentation with internal information markets, leading some to foresee that they will one day revolutionize decision-making. On current evidence, markets are some way off delivering on this expectation. Despite many years of experimentation on the part of innovative firms, with many encouraging results, information markets have so far failed to become established as a mainstream forecasting solution. Nor are there any signs that those who utilize markets currently are deploying these to guide 'big ticket' business decisions. If markets provide superior forecasts then the biggest gains will come from using them to inform key strategic choices, such as whether to merge with a rival or replace the CEO.

This chapter has highlighted some of the advantages of information markets over other mechanisms, but also the many practical barriers to their wider deployment in real-world organizations. Potentially, markets provide an ingenious solution to the incentive problems that can beset alternative approaches to informing organizational decision-making. At the same time, markets raise novel practical challenges; they are not suitable for all settings, and need to be designed and implemented carefully and sensitively to be effective. Markets may leak sensitive information in a way that other mechanisms do not (or do less), and this in turn can impact morale and motivation negatively, create legal complications by turning participants into 'insiders' and hand competitors vital commercial intelligence. Markets may be subject to manipulation (though the fear of this may be greater than the reality), and it can be challenging to sustain participation. However, executives and policy-makers must keep in mind the relevant counterfactual. The practical performance of markets should not be compared to some absolute ideal, rather to the merits and limitations of real-world

alternatives. Holding meetings, relying on 'expert' opinion, conducting surveys, or polling employees are all information mechanisms with costs and potential weaknesses. Often, their limitations are not sufficiently appreciated and challenged.

To date, much of the experimentation with information markets has been carried out either inside laboratories or by corporations behind closed doors. Producing accurate forecasts is essential to the success of all organizations and more of these should consider collaborating with academics to road-test competing information mechanisms scientifically, ideally allowing randomized trials within their own realistic settings. Organizations should look to do this as part of a broader information-based strategy (Davenport, 2009).[33] Prejudices against information markets, related to negative feelings about 'gambling' in the workplace, the threat to existing authority, or simply the counterintuitiveness of crowd wisdom, could be challenged more aggressively through education. In particular, business schools could play a big part in stimulating practitioners to engage with the concept of information markets through classroom experience and open, scientific debate.

Notes

* This work was completed at New College, Oxford University, and the Oxford-Man Institute of Quantitative Finance. I am grateful to both institutions for their generous support. Parts of the chapter build on earlier background research carried out for a study on the performance of distributed problem-solving networks, which was supported by the Oxford Internet Institute and McKinsey & Company. Several colleagues provided valuable comments and input, particularly Bill Dutton and David Bray at the Oxford Internet Institute, Michael Chui and James Manyika at McKinsey, and Clare Leaver and Ian Jewitt at the Department of Economics at Oxford.

1 Sunstein (2006a) makes the case for a wide range of information markets to address a range of important issues, including predicting damage from natural disasters. Meanwhile, Robin Hanson has suggested that 'Decision markets will one day revolutionize governance, both public and private' (excerpt from Robin Hanson's review of Abramowicz (2007), full text available online at www.overcomingbias.com/2008/01/predictocracy.html).

2 See, for instance, coverage in *The Economist* (2005), the *New York Times* (Lohr, 2008), and the *Wall Street Journal* (Dvorak, 2008). The *Special Interest Group on Prediction Markets* at www.forecastingprinciples.com provides an updated overview of media coverage related to information markets.

3 The likely gains from markets for public policy have been articulated by several academic economists (Hanson, 2006a; Ledyard *et al.*, 2006; Hahn and Tetlock, 2005).

4 The tendency of experts to herd has been observed empirically in many settings, including economic forecasting: 'economic forecasters all tend to be wrong in the same way. Their incentives to flock together are obvious enough.' Tim Harford, *Financial Times*, 9 August, 2008. Models of 'career concerns' can explain this phenomenon theoretically. See, for example, Scharfstein and Stein (1990) for an agency theoretic model in which concern for professional reputation causes experts to mimic the decisions of other experts.

5 Opinion pools can be classified into two broad categories: mathematical approaches and behavioural approaches (Clemen and Winkler, 1999). In mathematical approaches, the opinions of individual experts are expressed as subjective probability

distributions over outcomes of an uncertain event. They are combined through various mathematical methods to form an aggregated probability distribution. The important assumption of behavioural approaches is that, through exchanging opinions or information, experts can eventually reach an equilibrium in which further interaction won't change their opinions. The Delphi technique (Linstone, 2007) is a well-known behavioural approach. Both mathematical and behavioural approaches have advantages but also drawbacks. See Chen *et al.* (2005) for further discussion.

6 Armstrong (2006) discusses many examples of behaviour in face-to-face meetings that can lead to biased results, including the tendency for personal influence over the group outcome to be related to how loudly a person speaks, tone of her voice and physical appearance. See Sunstein (2006) for further insights regarding deliberative groups. Deliberative groups amplify cognitive errors, and fall prey to informational cascades and group polarization. Structured group deliberation processes, such as the Delphi method, have been designed to try to mitigate some of the shortcomings of unstructured face-to-face meetings.

7 US real estate company, Forest City Enterprises Inc., uses instant polling to improve the decision-making in its communications department: 'When evaluating materials for the company newsletter, website or other communications tools, the company uses the anonymity of the technology to gain input from the entire team.' http://marketingpr.suite101.com/article.cfm/audience_response_systems_for_employee_opinion.

8 The following discussion draws on points made by Wolfers and Zitzewitz (2004).

9 http://future.iftf.org/2006/12/prediction_mark.html.

10 Croxson and Reade (2010) study Betfair's major football markets, finding prices update remarkably quickly to the arrival of goals during live matches. The incorporation of relevant information may be slower than this in less liquid markets. Investigating horse-racing markets, Smith *et al.* (2006) have found Betfair prices to be less biased than those from bookmakers.

11 See Hanson's blog, 'Overcoming Bias' for a discussion of this initiative: www.overcomingbias.com/2006/11/first_known_bus.html.

12 The current recession may have influenced utilization of business prediction markets, but not obviously negatively: 665 of the total 1,695 respondents in the McKinsey survey felt that the economic downturn had increased interest in Web 2.0 technologies within their organizations, 704 considered it to have had no impact, and 187 reported that it had decreased interest in these tools.

13 I am grateful to McKinsey & Company for providing access to their proprietary detailed survey data beyond the summary results published at their website (www.mckinseyquarterly.com), and to Michael Chui at the McKinsey Global Institute for several helpful discussions on this topic.

14 To date, the most high-profile initiative to create a public sector information market has been a proposed information market for the US intelligence community – the 'Policy Analysis Market', which was the brainchild of academic economist Robin Hanson. With the blessing and seed funding of the US Defense Advanced Research Projects Agency (DARPA), 'PAM' was designed to aggregate information regarding geopolitical risks and terror attacks. It failed in a storm of controversy, largely for political reasons (Hanson, 2006a). The author is aware of selected more modest attempts to introduce information markets into public sector environments, including a recent experiment in the United Kingdom to deploy information markets to forecast demand for bed space at a large public hospital (Rajakovich and Vladimirov, 2009).

15 Tziralis and Tatsiopoulos (2007) provide a comprehensive survey of the prediction markets literature.

16 Lewis Shephard, Director of Microsoft's Institute for Advanced Technology in Governments, has commented on his blog that the evidence regarding prediction markets is not established enough to justify their deployment by governments:

> Microsoft Research has explored prediction markets, running an internal one as the 'Information Forecasting Exchange' from 2003–2006. Internal efforts at Yahoo and Google have also been noted. But, frankly, I'm not actively promoting PM's to government friends, as I don't believe we understand the results and supporting science well enough yet.
>
> (http://lewisshepherd.wordpress.com/2008/07/12/test-for-prediction-markets/)

17 See Levitt and List (2007) for a recent consideration of factors affecting the generalizability of laboratory findings.

18 Many other design choices must be made when implementing a prediction market, including which trading mechanism to use. See Tziralis and Tatsiopoulos (2007) for a recent overview of the sub-literature dealing with market modelling and design.

19 Wolfers and Zitzewitz (2006) have highlighted the problem of attracting non-informed traders as one of five open questions about prediction markets:

> Counterintuitively, the problem for most prediction markets is attracting sufficient uninformed order flow. Markets need uninformed order flow to function; when trading is conducted by rational traders, whose sole motivation is expected returns, the no-trade theorem binds, and the market unravels. Uninformed order flow can have a variety of motivations (entertainment, overconfidence, and hedging, for example), but with the exception of hedging, these are usually noneconomic, putting economists at a comparative disadvantage in predicting which markets will succeed.

20 'Potential traders are individuals who have information related to our private markets and include those in the healthcare and public health fields, such as physicians, nurses, microbiologists, epidemiologists, and public health professionals, among others.' http://iehm.uiowa.edu/iehm/content/faq.html.

21 Scott Page has demonstrated formally that the wisdom of the crowd depends not only on the abilities of the people within it, but also on their cognitive differences (Page, 2007).

22 In order to operate the Iowa Election Markets legally, Iowa academics obtained no-action letters from US regulator the Commodity Futures Trading Commission. To secure this relief it was agreed to limit positions to $500 and to operate on a non-for-profit basis.

23 Bo Cowgill, until recently the manager of Google's internal prediction markets, observed that Google employees seem more concerned with status than cash remuneration: 'on a number of occasions, I've forgotten to pay out the small cash prizes we have at Google, and nobody noticed. But everyone notices when the T-shirts that show who won don't come' (Dye, 2008).

24 Adam Siegel, CEO of prediction markets provider Inkling, sees this loss of control over sensitive information as a prohibitive concern for some organizations: 'Some people are simply scared of exposing sensitive information – it's too politically toxic in their organizational climate' (Comment made during interviews conducted by the author in 2007–2008.)

25 The Golem effect designates the opposite phenomena, whereby low expectations encourage low performance. Professor Dov Eden at Tel Aviv University has confirmed both self-fulfilling phenomena in banks, schools and the military, among other settings (Eden, 1984). His advice to leaders: 'Have high expectations and reinforce them with positive messages to the employee, even if it requires being a good actor' www.aftau.org/site/News2?page=NewsArticle&id= 6927.

26 Managing expectations about staff retention may be a particular concern around the time of a major organizational change, such as following a merger. According to press reports, 120 of Cadbury's 170 senior managers left the confectioner in the six months following its takeover by Kraft in February 2010. www.thegrocer.co.uk/articles.aspx?page=articles&ID=211285.

27 Comments were provided during interviews conducted by the author in the period 2007–2008.
28 Similarly, it appears that Hewlett Packard previously refined its proprietary information mechanism, BRAIN, to conceal aggregates such as the current forecast (Acheson *et al.*, 1997)
29 *Predictocracy* can be read online as a blog: http://predictocracy.org/blog/?p=107.
30 Mat Fogarty at prediction markets technology provider Crowdcast has expressed the view that, at least in the corporate setting, where play-money incentives linked to modest prizes are the norm, fears about manipulation seem exaggerated: 'Sometimes people ask about manipulation. With the current low level of prizing it is not a concern. I have not come across manipulation in practice.' These comments were provided during interviews conducted by the author in the period 2007–2008.
31 Comments were provided during interviews conducted by the author in the period 2007–2008.
32 http://blogs.hbr.org/hbr/mcafee/2009/12/prediction-markets-a-teaching-moment.html.
33 Thomas Davenport has urged organizations to shift to a 'test-and-learn mind-set', basing their decisions on the results of randomized internal experiments. See his recent article in the *Harvard Business Review* (February 2009): 'How to Design Smart Business Experiments', available online at: http://hbr.org/2009/02/how-todesign-smart-business-experiments/ar/1.

References

Abramowicz, M. (2007), *Predictocracy: Market Mechanisms for Public and Private Decision Making*, Yale University Press.

Acheson, D., R. Mazzilli and L. Prata (1997), 'Prediction Markets', mimeo, MIT Sloan Management.

Armstrong, J.S. (2006), 'How to Make Better Forecasts and Decisions: Avoid Face-to-Face Meetings', *Foresight: The International Journal of Applied Forecasting*, 5, 3–15.

Arrow, K.J., S. Sunder, R. Forsythe, R.E. Litan, M. Gorham, E.W. Zitzewitz, R.W. Hahn, R. Hanson, D. Kahneman, J.O. Ledyard, S. Levmore, P.R. Milgrom, F.D. Nelson, G.R. Neumann, M. Ottaviani, C.R. Plott, T.C. Schelling, R.J. Shiller, V. L. Smith, E.C. Snowberg, C.R. Sunstein, P.C. Tetlock, P.E. Tetlock, H.R. Varian and J. Wolfers (2007), 'Statement on Prediction Markets', *AEI-Brookings Joint Center Related Publication No. 07–11*.

Bell, T. (2009), 'Private Prediction Markets and the Law', *The Journal of Prediction Markets* 3(1), 89–110.

Berg, J., R. Forsythe, F. Nelson and T. Rietz (2008), 'Results from a Dozen Years of Election Futures Market Research', *in* C.R. Plott and V.L. Smith (eds), *Handbook of Experimental Economics*, Elsevier.

Camerer, C. and K. Weigelt (1991), 'Information Mirages in Experimental Asset Markets', *Journal of Business* 64, 463–493.

Chen, K.-Y., B.A. Huberman and L.R. Fine (2001), 'Forecasting Uncertain Events with Small Groups', *Proceedings of the 3rd ACM Conference on Electronic Commerce (EC–2001)*, Tampa, Florida, USA, 14–17 October.

Chen, Y., C. Chu, T. Mullen and D.M. Pennock (2005), 'Information Markets vs. Opinion Pools: an Empirical Comparison', *Proceedings of the 6th ACM Conference on Electronic Commerce*, Vancouver, BC, Canada, June 2005.

Chordia, T., R. Roll and A. Subrahmanyam (2008), 'Liquidity and Market Efficiency', *Journal of Financial Economics* 87, 249–268.

Christiansen, J. (2007), 'Practical Experiments in Small Markets and Behaviours Observed', *The Journal of Prediction Markets* 1, 17–41.

Clemen, R.T. and R.L. Winkler (1999), 'Combining Probability Distributions from Experts in Risk Analysis', *Risk Analysis* 19(2), 187–203.

Cowgill, B., J. Wolfers and E. Zitzewitz (2009), 'Using Prediction Markets to Track Information Flows: Evidence from Google', *Working Paper.*

Croxson, K. and J. Reade (2010), 'Information and Efficiency: Goal Arrival in Soccer Betting', mimeo, Oxford University.

D'Amuri, F. and J. Marcucci (2009), *'Google it!' Forecasting the US Unemployment Rate with a Google Job Search Index*, Institute for Social and Economic Reform.

Davenport, T.H. (2009), 'How to Design Smart Business Experiments', *Harvard Business Review*, 87, 68.

De Long, J.B., A. Shleifer, L.H. Summers and R.J. Waldmann (1990), 'Positive Feedback Investment Strategies and Destabilizing Rational Speculation', *Journal of Finance* 45, 379–395.

Dvorak, P. (2008), 'Best Buy Taps "Prediction Market"', *Wall Street Journal*, 16 September.

Dye, R. (2008), 'The Promise of Prediction Markets: A Roundtable', *McKinsey Quarterly.*

Economist, The (2005), 'Market, Market, on the Wall. Technology Trends: If Prediction Markets are so Good at Making Forecasts, Why Not Use them to Identify Emerging Technologies?', 10 December.

Eden, D. (1984), 'Self-Fulfilling Prophecy as a Management Tool: Harnessing Pygmalion', *The Academy of Management Review* 9(1), 64–73.

Forsythe, R., T. Palfrey and C. Plott (1982), 'Asset Valuation in an Experimental Market', *Econometrica* 50, 537–567.

Forsythe, R., F. Nelson, G.R. Neumann and J. Wright (1992), 'Anatomy of an Experimental Political Stock Market', *The American Economic Review* 82, 1142–1161.

Ginsberg, J., M.H. Mohebbi, R.S. Patel, L. Brammer, M.S. Smolinski and L. Brilliant (2009), 'Detecting Influenza Epidemics Using Search Engine Query Data', *Nature* 457, 1012–1014.

Graefe, A. (2009), 'Prediction Markets versus Alternative Methods: Empirical Tests of Accuracy and Acceptability', PhD Dissertation, University of Karlsruhe.

Green, K.C., J.S. Armstrong and A. Graefe (2007), 'Methods to Elicit Forecasts from Groups: Delphi and Prediction Markets Compared', *Foresight – The International Journal of Applied Forecasting.* Online: http//marketing.wharton.openn.edu/documents/research/delphi-wpv33.pdf.

Gruca, T.S., J. Berg and M. Cipriano (2008), 'Incentive and Accuracy Issues in Movie Production Markets', *The Journal of Prediction Markets* 2(1), 29–43.

Hahn, R.W. and P.C. Tetlock (2005), 'Using Information Markets to Improve Public Decision Making', *Harvard Journal of Law and Public Policy* 29, 214–289.

Hahn, R.W. and P.C. Tetlock (2007), 'Optimal Liquidity Provision for Decision Makers', *AEI-Brookings Joint Center Working Paper No. 06–18.*

Hanson, R. (2006a), 'Decision Markets for Policy Advice', *in* A.S. Gerber and E.M. Patashnik (eds), *Promoting the General Welfare: New Perspectives on Government Performance*, Brookings Institution Press.

Hanson, R. (2006b), 'Foul Play in Information Markets', *in* R.W. Hahn and P.C. Tetlock (eds), *Information Markets: A New Way of Making Decisions*, AEI Press.

Hanson, R. (2008), 'Insider Trading and Prediction Markets', *Journal of Law, Economics, and Policy* 4(2), 449–463.

Hanson, R. and R. Oprea (2007), 'A Manipulator can Aid Prediction Market Accuracy', *Economica* 76, 304–314.

Hanson, R., R. Oprea and D. Porter (2006), 'Information Aggregation and Manipulation in an Experimental Market', *Journal of Economic Behavior and Organization* 60, 449–459.

Hartzmark, S. and D. Solomon (2010), 'Efficiency and the Disposition Effect in NFL Prediction Markets', *Working Paper.*

Hayek, F.A. (1945), 'The Use of Knowledge in Society', *American Economic Review* 35(4), 519–530.

Hopman, J.W. (2007), 'Using Forecasting Markets to Manage Demand Risk', *Intel Technology Journal* 11(2), n.p.

Hussam, R.N., D. Porter and V.L. Smith (2008), 'Thar She Blows: Can Bubbles Be Rekindled with Experienced Subjects?', *American Economic Review* 98(3), 924–937.

Kumar, A. and C. Lee (2006), 'Retail Investor Sentiment and Return Comovements', *Journal of Finance* 61, 2451–2486.

LaComb, C.A., J.A. Barnett and Q. Pan (2007), 'The Imagination Market', *Information System Frontiers* 9, 245–256.

Ledyard, J., R. Hanson and T. Ishikida (2006), 'Designing Information Markets for Policy Analysis', *in* R.W. Hahn and P.C. Tetlock (eds), *Information Markets: A New Way of Making Decisions*, AEI Press.

Lei, V., C.N. Noussair and C. Plott (2001), 'Nonspeculative Bubbles in Experimental Asset Markets: Lack of Common Knowledge of Rationality vs. Actual Irrationality', *Econometrica* 69(4), 831–859.

Levitt, S. and J. List (2007), 'What do Laboratory Experiments Measuring Social Preferences Tell Us about the Real World?', *Journal of Economic Perspectives* 21(2), 153–174.

Linstone, H.A. (2007), *The Delphi Method: Techniques and Applications*, Addison-Wesley.

Lohr, S. (2008), 'Betting to Improve the Odds', *New York Times*, 9 April.

Luckner, S. (2007), 'Prediction Markets: How do Incentive Schemes Affect Prediction Accuracy?', *in* N. Jennings, G. Kersten, A. Ockenfels and C. Weinhardt (eds), *Negotiation and Market Engineering*, IBFI.

McKinsey (2009), 'How Companies are benefiting from Web 2.0: McKinsey Global Survey Results', *McKinsey Quarterly*, September.

Noeth, M., C. Camerer, C. Plott and M. Weber (1999), 'Information Aggregation in Experimental Asset Markets: Traps and Misaligned Beliefs', mimeo.

O'Brien, J. and S. Srivastava (1991), 'Dynamic Stock Markets with Multiple Assets', *Journal of Finance* 46, 1811–1838.

Ortner, G. (1998), 'Forecasting Markets – An Industrial Application, Part II', *Working Paper*, Technical University of Vienna, Department of Managerial Economics and Industrial Organization.

Ottaviani, M. and P. Sorensen (2006), 'Outcome Manipulation in Corporate Prediction Markets', *Journal of the European Economic Association* 5, 554–563.

Page, S. (2007), *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*, Princeton University Press.

Pennock, D.M., S. Lawrence, C. Giles and F. Nielsen (2001), 'The Real Power of Artificial Markets', *Science* 291(5506), 987–988.

Plott, C. and K.-Y. Chen (2002), 'Information Aggregation Mechanism: Concept, Design, and Implementation for a Sales Forecasting Problem', *Working Paper, No. 1131*, California Institute of Technology, Division of the Humanities and Social Sciences.

Plott, C. and S. Sunder (1982), 'Efficiency of Experimental Security Markets with Insider Information: An Application of Rational Expectations Models', *Journal of Political Economy* 90, 663–698.

Plott, C. and S. Sunder (1988), 'Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets', *Econometrica* 56(5), 1085–1118.

Polgreen, P.M., Y. Chen, D.M. Pennock and F.D. Nelson (2007), 'Using Internet Searches for Influenza Surveillance', *Clinical Infectious Diseases* 47, 1443–1448.

Prendergast, C. (1993), 'A Theory of "Yes Men"', *American Economic Review* 83(4), 757–770.

Raban, D.R. and D. Geifman (2009), 'Designing Online Information Aggregation and Prediction Markets for MBA Courses', *Interdisciplinary Journal of E-Learning and Learning Objects: IJELLO Special Series of Chais Conference 2009 Best Papers* 5, 247–262.

Rajakovich, D. and V. Vladimirov (2009), 'Prediction Markets as a Medical Forecasting Tool', *The Journal of Prediction Markets* 3(2), 78–106.

Rhode, P.W. and K.S. Strumpf (2008), 'Manipulating Political Stock Markets: A Field Experiment and a Century of Observational Data', *Working Paper*, June.

Rosenbloom, E.S. and W. Notz (2006), 'Statistical Tests of Real-Money versus Play-Money Prediction Markets', *Electronic Mar*kets 16(1), 63–69.

Sadka, R. and A. Scherbina (2007), 'Analyst Disagreement, Mispricing, and Liquidity', *Journal of Finance* 62, 2367–2403.

Scharfstein, D.S. and J.C. Stein (1990), 'Herd Behavior and Investment', *American Economic Review* 80, 465–479.

Servan-Schreiber, E., J. Wolfers, D.M. Pennock and B. Galebach (2004), 'Prediction Markets: Does Money Matter?', *Electronic Markets* 14(3), 243–251.

Siegel, S. (1961), 'Decision Making and Learning under Varying Conditions of Reinforcement', *Annals of the New York Academy of Science*, 89, 766–783.

Smith, M.A., D. Paton and L. Vaughan-Williams (2006), 'Market Efficiency in Person-to-Person Betting', *Economica* 73(292), 673–689.

Smith, V.L., G.L. Suchanek and A.W. Williams (1988), 'Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets', *Econometrica* 56(5), 1119–1151.

Spann, M. and E. Skiera (2003), 'Internet-Based Virtual Stock Markets for Business Forecasting', *Management Science* 49(10), 1310–1326.

Sunstein, C. (2006a), *Infotopia: How Many Minds Produce Knowledge?*, Oxford University Press.

Sunstein, C.R. (2006b), 'Deliberation and Information Markets', *in* R.W. Hahn and P.C. Tetlock (eds), *Information Markets: A New Way of Making Decisions*, AEI Press.

Surowiecki, J. (2004), *The Wisdom of Crowds: Why the Many are smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*, Little, Brown.

Tetlock, P.C. (2008), 'Liquidity and Prediction Market Efficiency', *Working Paper*.

Tetlock, P.E. (2005), *Expert Political Judgment: How Good is it? How Can we Know?*, Princeton University Press.

Tziralis, G. and I. Tatsiopoulos (2007), 'Prediction Markets: An Extended Literature Review', *The Journal of Prediction Markets* 1, 75–91.

Vaughan Williams, L. (2005), *Information Efficiency in Financial and Betting Markets*, Cambridge University Press.

Veiga, H. and M. Vorsatz (2009), 'Price Manipulation in an Experimental Asset Market', *European Economic Review* 53, 327–342.

Veiga, H. and M. Vorsatz (2010), 'Information Aggregation in Experimental Asset Markets in the Presence of a Manipulator', *Experimental Economics*.

Wolfers, J. and E. Zitzewitz (2004), 'Prediction Markets', *Journal of Economic Perspectives* 18(2), 107–126.

Wolfers, J. and E. Zitzewitz (2006), 'Five Open Questions About Prediction Markets' *in* R.W. Hahn and P.C. Tetlock (eds), *Information Markets: A New Way of Making Decisions*, AEI Press.