Information Markets: Feasibility and Performance

David Bray, Karen Croxson and William Dutton

Abstract

Information markets are markets designed specifically for the purpose of generating accurate information. Participants buy and sell assets whose payoffs are tied to the realization of future events. In theory, as individuals trade on their private information, an asset's price will move to incorporate all relevant news; the price, as a result, will constitute a collective prognosis which at any point in time is the best forecast available. Empirical evidence supports this theoretical finding: in a host of real-world applications, information markets deliver impressive prediction accuracy: probabilities implied by market prices typically calibrate to actual event frequencies extremely well and market forecasts routinely beat prognoses from recognized experts and polls. Markets can provide a remarkable crystal ball into the future, it seems, both for public events, such as sporting and political outcomes, but also for internal corporate events, such as whether a project deadline will be met. The idea that markets efficiently aggregate information is not new to economists but the application of the market mechanism explicitly for the purpose of forecasting the future is a recent innovation, driven to considerable extent by the very public success of the Iowa Election Markets in predicting US presidential elections. Excitement currently surrounds the potential for markets to revolutionize forecasting and decision-making in a host of settings, including in the corporate domain.

The contribution of the current paper is to try to broaden the prevailing perspective on information markets. We note that the current discourse focuses disproportionately on forecast accuracy, and suggest that this encourages an inappropriately narrow consideration of the merit of information markets as a real-world forecasting institution. We set out to offer a multidimensional look at performance and feasibility. Prediction accuracy, naturally, is a key consideration, but also important, we suggest, are such factors as costs of implementation, perceived legitimacy, and possible leakage of sensitive information. Our evaluation is only meaningful if comparative: alongside information markets, we consider the performance and feasibility of other Information Aggregation Mechanisms (IAMs), such as asking an expert for a forecast; obtaining and pooling multiple expert opinions; conducting a poll or survey; and leaving a group to deliberate. The evidence we review points to a somewhat mixed picture (summarized in tabular form in Tables II and III, p.31): not only is the feasibility of IAMs context-sensitive, we find (for instance, information markets cannot be used to aggregate information in some settings due to the need to write unambiguous contracts and settle markets objectively), but even where all IAMs are feasible, the optimal choice of mechanism (as judged by the balance of costs and benefits) depends on the setting: is this a corporate event or an event of general public interest? If corporate, is the data generated of a sensitive nature? Is it important that the IAM not only predict but also offer some explanation for its forecast?

The problem of information aggregation does not have a one-size-fits-all solution, we find. Rather the relative costs and benefits of different IAMs will vary depending on the setting. The literature has delivered insights to facilitate some of the necessary comparisons but important gaps exist. For instance, a number of findings point to the superior forecast accuracy of information markets compared to predictions from individual experts but little has been said about the relative costs of implementation and operation and risk factors, such as possible leakage of sensitive information. Even less work exists to inform other important comparisons, such as whether overall a firm should opt for quick polling (as some now do) as an internal mechanism for a quick collective view on a matter or set up and operate an information market for this purpose. Overall, the study highlights the need for more "open" research on comparative IAM performance. The aim should be to shed objective light on the merits of alternative mechanisms for different categories of real-world forecasting problem. Ideally this research should involve randomized trials in realistic field settings (much experimentation is carried out in laboratories or by firms behind closed doors).

Contents

Introduction

It has long been recognized that a byproduct of speculative markets is the efficient aggregation of information. Recently, however, and particularly since the publication of James Surowiecki's (2004) "The Wisdom of Crowds," interest has grown in harnessing the power of markets specifically for the purpose of generating accurate information. `Information markets' - also called `prediction markets', `decision markets', `event derivatives', `event futures,' and `idea futures' - allow participants to buy and sell assets whose payoffs are tied to the realization of future events. To illustrate, an information market might be used to forecast box office receipts for a new movie. A contract (or asset) might, for instance, pay $1 if ticket sales are above some pre-specified level by a pre-specified date and $0 otherwise. A market price of $0.43 could be interpreted as a 43% chance of success, so defined.

Probably the best known information markets are the Iowa political markets.[1] These markets were set up in 1988 by Iowa university academics to allow the public to "bet" on US presidential elections. Since their creation many more applications have emerged and now markets predict a wide range of events. Researchers at the University of Iowa have also developed a market to forecast outbreaks of Avian Flu,[2] whilst their contemporaries at the University of Miami have established a Hurricane Futures Market.[3] Elsewhere, Hollywood play-money markets invite the public to predict opening weekend box office sales and the Oscars,[4] and, significantly, some corporations have begun to explore the potential of information markets to harness collective wisdom internally. Firms such as Hewlett Packard, Google, General Electric and Microsoft are leading the way in experimenting with business information markets.[5] Trading there typically is restricted to an internal group of "experts" and contracts written on such matters as whether a project deadline will be met or a sales target exceeded. Interest is spreading in the potential use of markets to generate conditional estimates ("Will our product ship on time *if* we take the following action?"). Decision-makers in many settings could benefit form the potential for conditional information markets to provide neutral guidance for actions.

The weight of available evidence suggests that information markets, carefully designed and sensitively implemented, can provide a remarkable crystal ball into the future. Iowa University's US election markets have consistently beaten opinion polls and political pundits (Berg et al., 2001), Hollywood play-money markets pick Oscar winners for opening weekend box office receipts more successfully than critics (Wolfers and Zitzewitz, 2004). Orange juice futures provide more accurate forecasts of Florida's weather than traditional forecasting channels, and futures traders correctly guessed the supplier of the faulty component in the 1986 Challenger disaster — within minutes of the incident. There are just a few of the available examples of the predictive powers of markets - many academic articles and press reports further attest their success.[6] In the case of corporate applications, much of

[1] http://www.biz.uiowa.edu/iem/

[2] http://flupredictiion.uiowa.edu/fluhome/Market_AvianInfluenza.html

[3] http://hurricanefutures.miami.edu/

[4] http://www.hsx.com/

[5] http://en.wikipedia.org/wiki/Prediction_market

[6] See, particularly, Wolfers and Zitzewitz (2004), Ledyard (2006), and the references therein, and the references to press coverage at http://www.aei-brookings.org/pages/index.php?id=37.

the ongoing experimentation is not publicly reported unfortunately, but the indications are encouraging: the predictions from internal markets are often remarkably accurate, update `in-running' (continually, as information arrives) and appear superior to those obtained under more traditional methods. Hewlett Packard's internal prediction markets forecast printer sales more accurately than inside experts (Chen and Plott, 1998), for instance.

Economic theory can explain the uncanny accuracy of information markets. According to the efficient market hypothesis, a market is efficient if all available information is always fully reflected in its prices (Fama, 1970). It follows that if an information market is efficient then the price of a particular contract related to a potential outcome of an uncertain event reflects all relevant information and, so, serves as the best available forecast – more accurate than the forecast obtained by appealing to other mechanism. The rationale for why an information market should be efficient (why its trading mechanism should aggregate and display individual assessments in stock prices) is based upon the so-called Hayek hypothesis. Hayek hypothesized that the price mechanism in a competitive market will be the most efficient instrument to aggregate the asymmetrically dispersed information of market participants (Hayek, 1945). Because individual assessments are tradable via contracts in information markets, participants can compete on the basis of their individual assessments of the future. Thus, their aggregated information should be reflected in prices. In theory then, the forecast generated by an information market will be the very best available – better than that which can be obtained via any other Information Aggregation Mechanisms (IAM), including: asking experts, asking multiple experts and somehow pooling their opinions, polling people, conducting surveys, and leaving groups to deliberate.

Less formally, markets outperform in terms of forecast accuracy because: they provide individuals with strong incentives to gather information and act upon this by trading; and because the price mechanism provides a way of aggregating this information succinctly and continuously ("in-running"). We discuss in this paper some of the features of other IAMs that can undermine prediction accuracy:

- voting aggregates views but does not weight these by relevance (by contrast, in markets, misinformed traders will suffer heavy losses) and does not reward people for being right;
- individual experts are prone to biases (whereas markets help overcome these – experts who trade on biased views suffer losses);[7]
- deliberative groups, such as committees, are prone to "group think" and other biases (but traders submit views independently and anonymously).[8]

In this paper we ask how, overall, the performance of an information market should be assessed and compared to alternative IAMs. The quality of the information generated (e.g. its accuracy) is an obvious consideration, but not the exclusive concern, we suggest. We identify further dimensions of performance that will be relevant for those seeking to aggregate information, including the costs associated with setting up and running an IAM and the possible loss of control over sensitive information generated. An important consideration is also feasibility, since without

[7] Many problems have appeared in practice. For example, corrections are often needed for cognitive biases. Risk-aversion, uncommon priors, and state-dependent utility can also make it hard to infer the information people have from the estimates they make. Additionally, whenever people are simply asked to state their expected outcome, incentive compatibility is a concern: does the chosen individual have the right incentives to report the truth? Vaughan Williams (2005) reviews the efficacy of expert predictions in comparison to market-based forecasts.

[8] See Sunstein (2006) for analysis of information aggregation in deliberative groups. Deliberative groups amplify cognitive errors, and fall prey to informational cascades and group polarization.

this performance is irrelevant. We find that IAMs perform differently and are differently feasible depending on the context. The problem of aggregating information does not have a one-size-fits-all solution and vastly more research is needed to understand the context-sensitivity involved. The paper highlights and discusses some of the most important open questions.

The organization of the paper is as follows. The next section reviews information markets in terms of their current application and design. This exercise culminates naturally in the production of a tabular typology (Table I). This allows us to identify interesting differences between markets designed to yield public predictions and those intended for internal corporate use. In section 3, we turn to the key issue of the paper – a review of the feasibility and performance of information markets compared to alternative IAMs, looking across multiple dimensions. Tables II and III present an overview of the comparative feasibility and performance of IAMs. Section 4 concludes, summarizing key insights, highlighting gaps in understanding and suggesting directions for future work.

1. Information Markets: A Typology

1.1 Applications

The real-world deployment of information markets is expanding at a rapid rate. Current applications fall broadly into two categories: (i) those which concern a public event, such as the direction of the next interest rate decision, and (ii) those relating to an internal corporate event, such as whether a project will be completed by its deadline.

Public Events

Many markets have been established in an attempt to leverage collective intelligence for the purposes of forecasting public events. The most well known public prediction markets probably are the Iowa Election Markets which were set up in 1995 by Iowa university academics. [9] These markets allow the public to "bet" on US presidential elections. More recently, Iowa researchers have developed a market to forecast outbreaks of Avian Flu,[10] whilst their contemporaries at the University of Miami have established a Hurricane Futures Market.[11] Elsewhere, to give just a few of the many further examples, Yahoo! Research has a public prediction market to predict tech trends,[12] Hollywood play-money markets invite the public to forecast opening weekend box office sales and the Oscars,[13] and betting markets give rise to public predictions about sports encounters, reality TV contests, and a host of other popular events.[14]

Corporate Events

A feature common to a wide variety of organisations is that valuable, decision-relevant information often is widely dispersed among members. For instance, those "on the ground" may have some idea of how changes to a design feature are likely to

[9] http://www.biz.uiowa.edu/iem/
[10] http://fluprediction.uiowa.edu/fluhome/Market_AvianInfluenza.html
[11] http://hurricanefutures.miami.edu/
[12] http://buzz.research.yahoo.com/bk/index.html
[13] http://www.hsx.com/
[14] www.betfair.com, http://www.tradesports.com/

impact upon the delivery schedule or to affect end of quarter sales. A key challenge for internal strategists and decision-takers is somehow to aggregate this effectively so that judicious commercial decisions can be taken. Traditional approaches to internal information aggregation have centred on appealing to internal experts for an opinion, polling or surveying staff, or perhaps designating a committee or work group to deliberate the matter. Recently, a number of corporations have begun experimenting with information markets as an alternative mechanism. The earliest example of this thinking was the market academic economist Robin Hanson helped to establish at Xanadu in 1990. One of the claims traded there related to the delivery date of the firm's product:

> "Xanadu will deliver its product before Premier Deng of China dies." [15]

Hewlett Packard, was another early pioneer in this regard. It began experimenting with internal corporate information markets in the late 1990s, leveraging its crowd of employees to forecast sales. There are now many further examples of corporate information markets: Google has deployed prediction markets internally to forecast a range of outcomes of strategic importance such as product launch dates,[16] GE has experimented with their use to generate new business ideas,[17] Starwoods uses information markets to develop and select marketing campaigns,[18] and prediction markets are being deployed within the pharmaceutical sector, by major players such as Pfizer, Novartis, GSK, and Eli Lilly, for the production of new product forecasts, pipeline valuation, competitor intelligence and forecasts of existing drugs.[19]

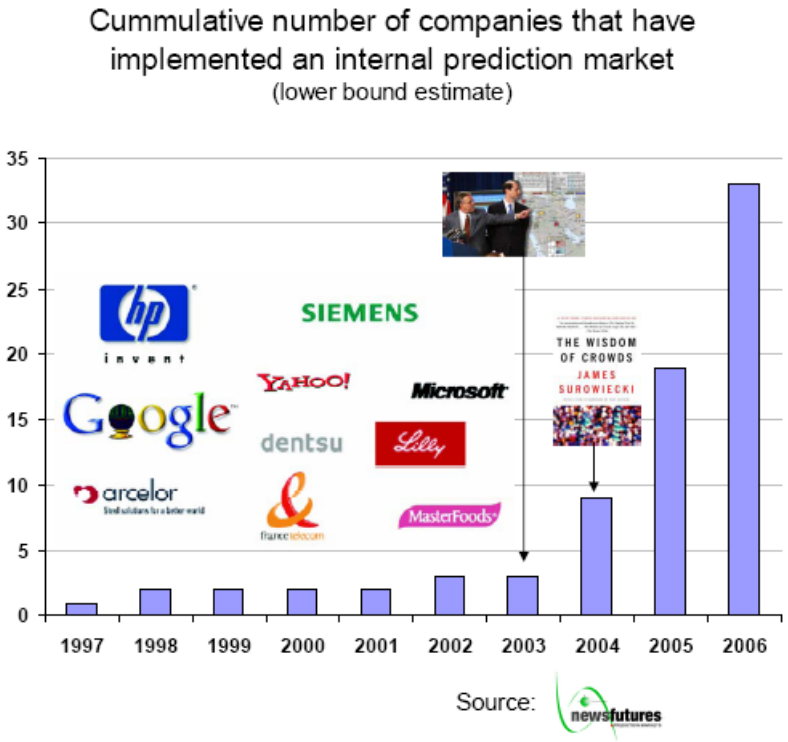

[15] http://www.overcomingbias.com/2006/11/first_known_bus.html
[16] http://googleblog.blogspot.com/2005/09/putting-crowd-wisdom-to-work.html
[17] www.consensuspoint.com
[18] Starwoods marketing department started out with some initial ideas and allowed employees to add new ideas or make changes to existing ones. Prediction markets were leveraged to select the best.
[19] http://www.pharmiweb.com/pressreleases/pressrel.asp?ROW_ID=2560

1.2 Asset Structure

The market creator needs to decide on the type of asset (interchangeably called "contract") that can be traded in the market. Two popular options are "winner-take-all" contracts and "index" contracts.

Winner-Take-All

In a "winner-take-all" model, contracts are listed which have payoffs that are tied to two or more mutually exclusive and mutually exhaustive outcomes of the event of interest. For instance, one contract might pay out if party A wins a specified election and another if instead party B is elected. Supposing that these parties are the only two contenders then one and only one of the contracts can "win" (the winning contract may pay out at $10); the losing contract becomes worthless (notionally, it is settled at $0).

Index Contract

An index contract cashes out based on the exact value of a particular variable at a specified time in the future. For example, an asset can be created that will pay its owner $1 for every percentage point of the overall vote won by party A. This is the type of contract found on standard financial futures exchanges (where the contract might deal the price of a barrel of oil to be delivered in 90 days).

Whichever form of contract is chosen, contracts must be clearly defined, easily understood and easily adjudicated. For example, 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 a certain date" should be used. 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 this particular market. Markets will only work well if there is reasonable clarity about and confidence in the contract design.

Conditional Contracts

Where the market is to be used for decision-making potentially conditional contracts will be very valuable. Conditional contracts deal with the likelihood that an outcome occurs given that another event arises. For instance, a conditional contract might relate to the question: "Will our product X ship on time if design feature A is included?" The prediction from this market could be used to guide the decision of whether or not to equip the product in question with this particular design feature. A concern raised recently in the academic literature is that working with conditional contracts can mean creating a large number of contracts and dealing with concomitant problems of market thinness (Ledyard, 2006). The practical deployment of conditional markets has yet to take off. Adam Siegel, CEO of leading platform provider Inkling Markets, confirms the lack of interest to date: "We haven't had many people do this yet but we suspect it will become more popular as market research firms for example begin to augment their methodologies to include prediction markets as a viable method to understand impact of product introductions in the marketplace, marketing on sales, brand revamps on perceptions, etc."[20]

[20]These comments were provided during an interview conducted by Karen Croxson for the current study.

1.3 Market Structure

Incentives

The market creator must ensure that the individuals they intend to participate have the right incentives to acquire information and trade in such a way as to reveal their information truthfully. Most likely, the major incentives for participation in prediction markets are: profit-seeking, entertainment and reputation-building.[21] The evidence so far suggests that all three can be important but more work is needed to unpick their individual importance in different settings.

It has been suggested that real money is necessary to ensure truth-revealing incentives (e.g. Gruca et al. 2003). However, the introduction of real money into a prediction market can raise regulatory complications, depending on the precise context and geography of the initiative. For instance, real money betting generally is considered illegal in the US.[22] Many market sponsors opt to side step potential complications by endowing users with virtual currency for trading. The success of many `play money' markets suggests that people may be sufficiently incentivized merely by the opportunity to pit their judgment against others and win status (player rankings have been used in many successful market settings) and by the chance to take part in a market with content which is sufficiently exciting (Servan-Schreiber et al. 2004). Bo Cowgill, Google's prediction markets architect, noted recently that "traders didn't seem to care about the cash prizes at all" in their internal prediction markets, instead, "people wanted to know about reputational prizes — their ranking in the system and *the shirts* that would identify them as winners."[23] [24] In the context of enterprise information markets, it will be important to consider how incentives compare to salaries, awards, or other incentives within the organization.

Trading Rules

The market organizer may which to place limits on the size of positions traders may take. This could be for regulatory reasons (e.g. as in the case of the Iowa Election Markets, where exposure is limited to $500), or possibly to prevent excessive shifts in price. But limits on trade size may work against forecast accuracy, preventing those truly informed from exerting sufficient influence on the collective assessment. Currently, information markets vary in this respect; sportsbooks such as Betfair place no restrictions on trade size but some others restrict this heavily.

Participation

Who should participate in the market is an important question for the market creator. If active traders are too few and far between, markets 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 - as few as sixteen traders (Christiansen, 2007) or even nine.[25] Beyond sufficient numbers, diversity of the participant pool has been identified as potentially critical to

[21] A conceivable additional motivation for participating in prediction markets is the desire to hedge specific risks. However, most prediction markets are not large enough for this purpose.

[22] IOWA academics obtained a letter of no action for the Iowa Election Markets (IEM). The IEM has received no-action-letters from the Division of Trading and Markets of the Commodity Futures Trading Commission. These extend no-action relief to the IEM's Political and Economic Markets

[23] http://bocowgill.com/2007/04/remarks-at-yahoo-confab-on-prediction.html

[24] It may well be that both real money and play money have their own advantages. Servan-Schreiber, et al. (2004), who compare predictions for the 2003-04 NFL season from TradeSports (real money) and NewsFutures (play money), find that whereas real money provides stronger incentives for information acquisition, play money delivers more efficient information aggregation (with play money the exogenous wealth distribution of players does not affect the weighting of the bets).

[25] http://groups.google.com/group/Prediction-Markets/browse_thread/thread/4c7dbfdc5756085f

the quality of predictions generated. Page (2007) explains that the wisdom of the crowd depends not only on the abilities of the people within it, but it also depends on their differences. Diversity improves the collective ability to predict and solve problems – when diverse mind sets are applied to a non-trivial challenge the outcome tends to improve.

In practice, markets for forecasting most public events tend to be open to the general public, though the Iowa Health markets in which only authenticated health professionals are eligible to trade, constitute a notable exception. Corporate market operators are more likely to hand pick groups of employees to trade, aiming to include only relative "experts."

An important design choice relates to the level of trader anonymity to be provided. Mat Fogarty at technology provider Xpree believes that anonymity of trading (especially around ship dates) is very important in the corporate context. Adam Siegel, CEO of Inkling Markets, another platform provider, concurs: "There cannot be any danger of repercussions for taking any position in the marketplace."[26]

Matching Mechanism

The way in which demand and supply is matched – the matching mechanism – is a critical further design feature. It is likely to have important implications for liquidity.

Continuous Double Auction (CDA): The preferred market structure in many professional financial markets, and also on betting exchanges, such as Betfair, is the Continuous Double Auction (CDA). A CDA allows participants to place orders to buy or sell at a specified volume and price. Whenever the orders can be matched (the sell price is lower than the buy price for compatible quantities), a trade occurs. Matching continues in this way until the best sell offer is above the offer to buy. An attractive feature of this market form is that participants can secure gains and cut losses at any time (assuming sufficiently deep markets). This encourages ongoing information acquisition. Thus, price (and hence the collective forecast) is continuously updated as traders use news to inform their decisions and move quickly to make a profit. However, if too few traders exist in a CDA-based market, traders may not be able to find liquidity to carry out the trades they desire. To counter this, a market maker may used - the market maker commits to provide liquidity on demand (by quoting prices at which participants may buy or sell). A possible problem with instituting a market marker is that they can end up making a loss.

Pari-Mutuel Market (PMM): An alternative market structure is the Pari-Mutuel Market (PMM), which pools all of the money placed on individual contracts and awards it to the eventual winner. This form of market was originally designed for sports betting. The ratio of the amount bet on an individual contract to the total pool of bets can be interpreted as the probability of a contract winning. A big advantage of the PMM is that it reduces the liquidity problem; because traders are betting into an expandable pool they can rest assured that their order will be taken (absent the need for a market maker). On the downside, a PMM limits the extent to which traders can act upon information that arrives once their initial bets have been placed (they cannot sell back their original bets). Rational traders actually will wait until the last possible moment to place bets (when all possible information is known), which may mean that contract prices up until this point are not particularly informative. A further potential drawback of the PMM relates to limitations on scope - it can be appropriate only for winner-take-all markets.

[26]These comments were provided during an interview conducted by Karen Croxson for the current study.

Some hybrid structures have emerged in an attempt to combine the best aspects of both CDA and pari-mutuel markets. Two such initiatives are particularly noteworthy:

Dynamic Pari-Mutuel Market (DPMM): David Pennock at Yahoo Labs has proposed DPMM (Pennock, 2004). This market structure combines the infinite liquidity and risk-free nature of a pari-mutuel market with the dynamic nature of a CDA.

Market Scoring Rules (MSR): Academic economist Robin Hanson has devised the MSR as a mechanism to provide infinite liquidity without the need for a risk-bearing market-maker (Hanson, 2003). Any trader can change the current prediction by replacing it with a new prediction providing they agree to pay off the person responsible for the current prediction when the market closes. So if the traders improve the prediction (i.e., move the market prices into the right direction) they can expect a positive payoff; if they do the opposite, then eventually they will lose.

Each market structure has benefits that will appeal differently to different market operators.

Disclosure of Market Information

Information generated by the market (the collective forecast) may be sensitive in its nature, particularly in the corporate setting. Mat Fogarty, CEO of Xpree, a corporate prediction market technology provider, explains:

> "the uncertainty over the Insider Trading and Reg FD (fair disclosure) rules are very worrying. The concern is that the information coming from the PM is so powerful that all recipients of that data are made insiders. Also, that 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]

Reflecting such concerns, the market creator may opt to hide some of the information generated in a prediction market from participants. Hewlett Packard has refined its proprietary platform to conceal market aggregates such as market price and now looks to be providing this platform (called BRAIN) to other corporations; Pfizer is a first client.

Further Design Considerations

Trading interface. Evidence from the field is suggestive of the importance of simplicity of the user interface and trading experience: Mat Fogarty at technology provider Xpree believes that "traders give you 1 minute or less to get it." Adam Siegel at Inkling makes adds: "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."[28]

Manipulation. The possibility of manipulation is a concern often raised in the context of discussions about prediction markets. Wolfers and Zitzewitz (2004) report that attempts to manipulate markets typically fail.29 Nevertheless, concern about possible manipulation helped derail a proposed project for trading in various forms of geopolitical risk.30 Where possible, markets should be designed so that temptation

[27] Comments made during an interview conducted by Karen Croxson for the current study.

[28] These comments were provided during interviews conducted by Karen Croxson as part of the current study.

[29] Illustrating this, Rhode and Strumpf (2005) analyze an attempt to manipulate the price of a Kerry victory on Tradesports in 2004, as well as their own attempts to manipulate prices on the Iowa Electronic Markets in 2000. Such manipulation efforts seem to have had only a very short-lived effect on prices.

[30] Defense Advanced Research Projects Agency (DARPA), a research think tank within the US Department of Defense, proposed to establish a Policy Analysis Market (PAM). The idea for PAM was

for trader manipulation is limited, and devices such as Hanson's Market Scoring Rule, which give incentives for honest trading, can be helpful here. Mat Fogarty at Xpree thinks that, at least in the corporate setting, where play-money incentives linked to modest prizes are the norm, concerns about manipulation are over-blown: "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." Adam Inkling, CEO of Inkling Markets, another platform provider, sees manipulation as a possible concern for public marketplaces but not really an issue in internal organizational markets: "We occasionally see attempts at manipulation on public sites and have implemented passive techniques to make administrators aware of such behavior. The main threat of manipulation in a public marketplace is someone creating many zombie accounts and using their starting balance in a single stock. In a private marketplace this is not an issue if access is given only to those with a certain email domain. Most corporations only allow their employees one email address, so creating an automated "swarm" is impossible. Multiple people colluding to drive a result in a market can be countered by ensuring you have a diverse participating base from multiple areas of the organization."[31]

shelved following a high-profile discussion of the risk that terrorists might profit financially from the acts of terrorism they perpetrate.

[31] Both these quotes are extracted from interviews conducted by Karen Croxson as part of the current study.

Table I. Typology of Information Markets

				PUBLIC EVENTS								CORPORATE EVENTS			
				Political		Health	Economy	Tech.	Entertainment	Sports		Bus. ideas	Sales	Project Completion	
				IOWA Election Markets	Intrade	IOWA Health Markets	IOWA Fed Funds Rate Market	yahoo tech trends	Holywood Stock Exchange	Betf air	Trade sports	GE	HP	Google	Microsoft
MARKET STRUCTURE	**Incentives**	Play Money	Prizes			√						√	√	√	
			No Prizes					√	√						
		Real Money		√	√		√			√			√		√
	Trading Rules	Position Size	Limit	√		√	√						?		?
			No Limit		√					√	√				
	Participation	Open	Anonymous	√	√			√	√			√?	√ ?	√ ?	√ ?
			Registration required	√	√			√	√	√					
		Closed	Expert			√	√								
			By Invitation			√						√	√	√	√
	Matching mechanism	CDA		√	√	√	√		√	√	√	?		√	
		MSR										?	√		√
		DPMM						√							
	Market Information Disclosure	Price	Visible	√	√	√	√	√	√	√	√	√		?	
			Not visible										√		
		Depth	Visible	√	√	√	√	√	√	√	√	?		?	
			Not visible										√		
ASSET STRUCTURE	**Winner-take-all**			√	√	√	√			√	√	√	√	?	√
	Index			√						√					
	Conditional Assets														

2. Feasibility and Performance: Markets vs. Alternatives

In this section we consider the feasibility and performance of information markets, comparing this to the feasibility and performance of other Information Aggregation Mechanisms (IAMs), such as consulting an expert, pooling expert opinions, carrying out a poll or survey, or asking a committee to deliberate.[32] Consideration of performance, we suggest, will come down to identifying and reviewing the major costs and benefits of each IAM. On the benefit side the primary consideration must be the quality of information generated, but we consider not only its accuracy but also the extent to which the prediction updates "in-running" and is explained by the IAM. The most obvious costs of any IAM are the time and money required to establish and operate the mechanism. Beyond this we consider also the possible loss of control over sensitive information which the IAM may generate (the market forecast itself may be commercially sensitive information, for instance). We assess feasibility for each IAM in terms of three main aspects: perceived legitimacy, scope, regulation, and participation.

2.1 Information Markets

Quality of Information

Accuracy: Considerable evidence suggests that information markets yield remarkably accurate predictions in comparison to polls, expert judgement, surveys, and group deliberation. Where accuracy is formally tested, three main tests seem to be preferred:

A *calibration test* is the primary tool used to evaluate the quality of probabilistic forecasts. It asks how often a forecaster's prediction of an event, E, is correct given the probability attached with that particular prediction, P(E). For instance, if the probabilities implied by prices are accurate, events priced such that the implied probability is 10 percent should occur about 10 percent of the time. The following chart depicts the impressive calibration of Storage Markets.[33]

[32] There are other forecasting methods beyond those considered in our study. See Armstrong's (ed.) (2001) book for discussion of role playing as a method for forecasting the decisions of others ("How will competitors respond if we raise price?"), and of econometric forecasting, in particular.
[33] http://storagemarkets.typepad.com/my_weblog/2007/07/updated-how-wel.html

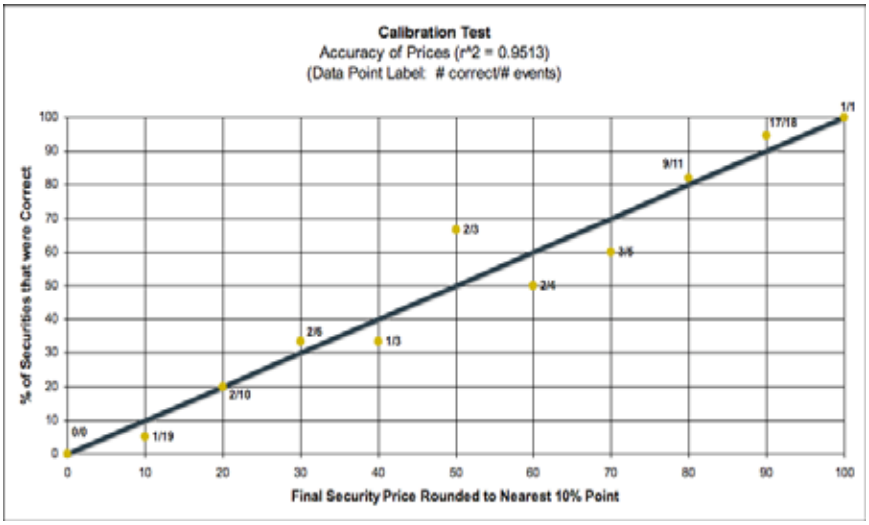

Similar performance has been found in markets for forecasting other events, both public and corporate.34 Even very small markets that are thin with few participants have performed well in calibration tests (Christiansen, 2007). The evidence to date is encouraging but more experimentation and analysis is needed.

A second popular test is to compute the *logarithmic score*. This compares the quality of the market prediction to that of a so-called "expert". Markets reviewed have also tended to perform well on this basis (Servan-Schreiber et al., 2004).

Entropy is a third measure sometimes used to gauge prediction quality. Entropy is a concept of decidedness in decision/information theory – the lower the entropy the more decisive the prediction.[35] The black line in the following chart reports the entropy of google's internal prediction market over time, benchmarking this to the entropy of illustrative probability distributions. The google market, this suggests, grows more confident in its prediction over time.[36]

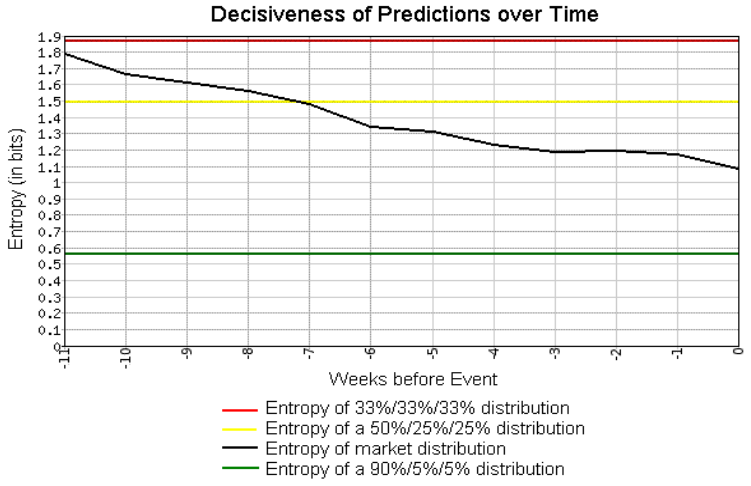

[34] For instance, Google's internal markets for forecasting product launch dates, new office openings, and many other things, have also been successful according to the calibration test: http://googleblog.blogspot.com/2005/09/putting-crowd-wisdom-to-work.html.

[35] See Cover and Thomas (2006) for an introduction to entropy in the concept of information theory.

It appears that entropy is not always studied (or at least publicly reported) in connection with analysis of information markets. This is regrettable: it offers a useful additional perspective on performance.

"In-running:" Beyond accuracy, markets offer the advantage of predictions that update "in-running." They can be thought of as a "pull" forecasting mechanism, as compared to the "push" mechanisms offered by polls and surveys. As new information emerges, traders in prediction markets react quickly by changing their positions and the market price moves rapidly to incorporate that new information.[37]

Explanatory content: By default, markets yield predictions but not necessarily any explanation for these predictions. This may or may not be a concern, depending on the context. Some corporations deploying markets internally have begun to experiment with the running blogs and wikis alongside their marketplaces, in an effort to capture trading motivations. The use of such tools could also have other advantages. For one, it might stimulate trading interest, leading to deeper markets. Meanwhile, LaComb et al. (2007) report that an anonymous blogging capability added to an internal GE "ideas" market provided a 'press' element to the market, allowed for a (possibly biased) discussion of the pros and cons of various ideas, and allowed traders to ask for clarification of ideas and for the idea's originators to respond with details.

Cost-efficiency

The set-up costs for information markets can vary considerably. Initial outlay can be very low if an off-the-shelf platform is `rented' from a third party provider, such as Inkling or Xpree. If a highly-customized in-house system is sought then the initial outlay will be much higher. Participants may need to be educated to play the market, again this will depend on the setting. Once the market is up and running, the ongoing operational costs typically will be very modest. However, where an organization is considering using its own employees as traders it must consider the cost of distracting potentially many of these potentially much of the time with an internal market. 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). Offsetting benefits that may arise from an internal information market include greater worker morale (employees may enjoying bragging about successful trading and feeling that their voice is heard), which in turn may lower overall agency costs.

Control over Sensitive Information

Where markets are used internally within organizations, managers may be sensitive to the release of information to market participants regarding:

- the nature of the organization's problem (the fact that the firm is seeking product ideas, improved sales forecasts etc. could itself be sensitive information)
- the prevailing market forecast (if visible to participants it may be leaked to competitors, financial markets).

Adam Siegel, CEO of prediction markets provider Inkling sees this loss of sensitive information as one of the key barriers to adoption: "some people are simply scared of

[37] Illustrating this, Croxson and Reade (2007) find that prices on Betfair's major soccer markets swiftly and completely impound the news of a goal. The incorporation of relevant information may be slower than this in less liquid market situations.

exposing sensitive information – it's too politically toxic in their organizational climate."[38] HP's new BRAIN market format addresses some of these potential concerns. Traders are rewarded if they win but don't see the aggregate forecast.[39]

Feasibility

Perceived legitimacy: At present info markets remain unfamiliar and Sunstein (2006) notes that their use might breed confusion and distrust. Furthermore, information markets imply that the expertise of individual in-house experts is less valuable. Wisdom gleaned from employees through internal markets can temper executive decisions and this will not always be welcome. Robin Hanson and others have emphasized these and related points in connection with the feasibility of information markets. An important consideration that can arise for corporations is the possibility that employees technically might become "insiders" as a result of their exposure to an insight from an internal marketplace, with potentially complex legal implications.

Scope: Not all questions are amenable to being answered by information markets. A number of conditions must be satisfied:

- it must be possible to make future market situations tradable via well-defined contracts, the determination of whose payoffs is objective and clearly understood by all participants.
- participants must have some knowledge about the future market situation (otherwise, stock prices would be randomly set and the Hayek hypothesis would not apply).
- individuals must have sufficient incentive to participate and reveal their true valuations.

An example of a question not compatible with these conditions is "Patient A presents symptions A, B, C. Would you diagnose this as condition Z – yes or no?" The problem here is that the answer is a matter of personal opinion (or professional judgment) rather than some objective truth.

Regulation: Anti-gambling laws arguably are now the main barrier to wider use of real-money information markets (Hanson, 2003). Gambling is illegal in many jurisdictions, including the US. Play-money markets present no regulatory problems, however, and although some academics feel that real-money is necessary to create sufficient incentives for information gathering and revelation (e.g. Gruca et al. 2003), recent studies have suggested that play money can be as effective in practice (Servan-Schreiber et al., 2004).

Participation: Participants must be available who are sufficient in number and in diversity in order for information markets to produce accurate information. This may not always be the case. However, this potential limitation is not unique to information markets – it can restrict the ability of other IAMs to perform effectively.

2.2 Expert Opinion

A well-established and still widely-used alternative to information markets is to ask an expert for a forecast. One might, for instance, approach a weather forecaster, a betting

[38] Comments made during an interview conducted by Karen Croxson for this study.
[39] http://www.hpl.hp.com/research/ssrc/competitive/brain/?jumpid=reg_R1002_USEN

tipster, a political pundit, an economic or financial analyst, a sales manager, or the head of new product ideas, for an expert prediction depending on the event of interest.

Quality of Information

Accuracy: A number of factors can compromise the accuracy of expert judgement. An important one concerns incentive compatibility: whenever people are simply asked to state their expected outcome the concern arises: does the chosen individual have the right incentives to report the truth? Scoring rules can be used to give agents incentives to acquire information they would not otherwise possess and reveal this - at the expense of a patron who agrees to pay a reward. They have long been used in weather forecasting, economic forecasting, and many other domains. But while scoring rules solve incentive problems and do not suffer from problems with market thinness (as potentially information markets do), some other problems have appeared in practice. For example, corrections are often needed for cognitive biases (Kadane and Winkler, 1988).

Overall the weight of evidence suggests that information markets are able to provide superior forecasts to experts across many settings. Horse races are better predicted by betting markets than by professional handicappers (Figlewski, 1979) and orange juice commodity futures markets improve on government weather forecasts (Roll, 1984). Play money markets predicted the 2000 Oscar winners better than 4 out of 5 columnists (Pennock, Giles, and Nielsen, 2001), and NFL results better than almost 2,000 experts in an online football forecasting contest (Servan-Schreiber et al., 2004). When compared to official Hewlett-Packard forecasts of printer sales, internal corporate markets were more accurate six out of eight times, even though the official forecasts where made after the markets closed and with knowledge of the market prices (Chen and Plott, 1998).

"In-running:" Appealing to an expert generates a spot prediction, rather than a forecast that is updated "in-running." That is, like most other traditional approaches to information aggregation, the information provided by an expert has an episodic character: results have to be analyzed, evaluated and summarized manually, at a particular point in time.

Explanatory content: Individual experts could be asked to justify their prediction. Since only a single participant is involved the process of eliciting this explanation is straightforward compared to the task of seeking to understand why the participants in an information market may have traded as they did.

Cost efficiency

The costs involved in obtaining an expert prediction are likely to vary depending on the setting. In all cases, a suitable expert will need to be identified and this has the potential to be a costly and time-consuming exercise. The person will require payment for his or her services and may need to be provided with access to privileged information in order to build a forecast. The whole process may take some time to complete and in the corporate setting may involve disruption to other employees.

Control over sensitive information

Commissioning a single expert to provide an opinion is a way to reduce the potential for leakage of sensitive information. This may or may not be a key concern, depending on the setting.

Feasibility

Perceived legitimacy: Providing that the expert is viewed as credible, this is likely to be considered a legitimate approach to obtaining a forecast.

Scope: It is necessary to find a suitable expert for the topic in question (someone who possesses all information relevant to a particular issue) and to give this person sufficient time (and possibly access to relevant materials) to produce a forecast.

Regulation: There are no regulatory barriers to the use of individual expert forecasts.

Participation: Typically, experts will be identified from within an organization and establish themselves internally as an authority on a particular set of issues (project development managers will be regarded as experts on their particular projects), or they will emerge in the market at large as experts on public events (e.g. weather forecasters, financial analysts, sporting and political commentators). They do this for the prospect of financial reward and possibly have additional motivations related to status.

2.3 Expert Opinion Pools

Often experts are not in agreement and many methods have been proposed to combine their opinions into a single view. Chen and Pennock (2005) review opinion pools, noting their classification by Clemen and Winkler (1999) into two broad categories: mathematical approaches and behavioral approaches. 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.[40] The important assumption of behavioral approaches is that, through exchanging opinions or information, experts can eventually reach an equilibrium in which further interaction won't change their opinions. One of the best known behavioural approaches is the Delphi technique (Linstone and Turoff, 1975). Both mathematical and behavioural approaches have their pros and cons. Mathematical approaches are comparatively easy to use and less time consuming but several impossibility results (e.g., Genest, 1984) show that no aggregation function can satisfy all desired properties of an opinion pool, unless the pooled opinion degenerates to a single individual opinion, which implies a dictator. Meanwhile, behavioral approaches like the Delphi are attractive in that they allow experts to revise their information and opinions dynamically, but these revision processes are not guaranteed to converge.

Quality of Information

Accuracy: We know of two studies which have sought to compare the prediction accuracy of information markets to pooled expert opinion. Servan-Schreiber et al. (2004) compare information markets to the arithmetic average of expert opinions in the context of forecasting NFL outcomes. The predictions from two markets are analyzed, Tradesports and NewsFutures, and compared to the average predictions of nearly two thousand experts: the market-based forecasts comfortably outperform.

[40]See Genest and Zidek (1986) and French (1985) for comprehensive reviews of mathematical approaches.

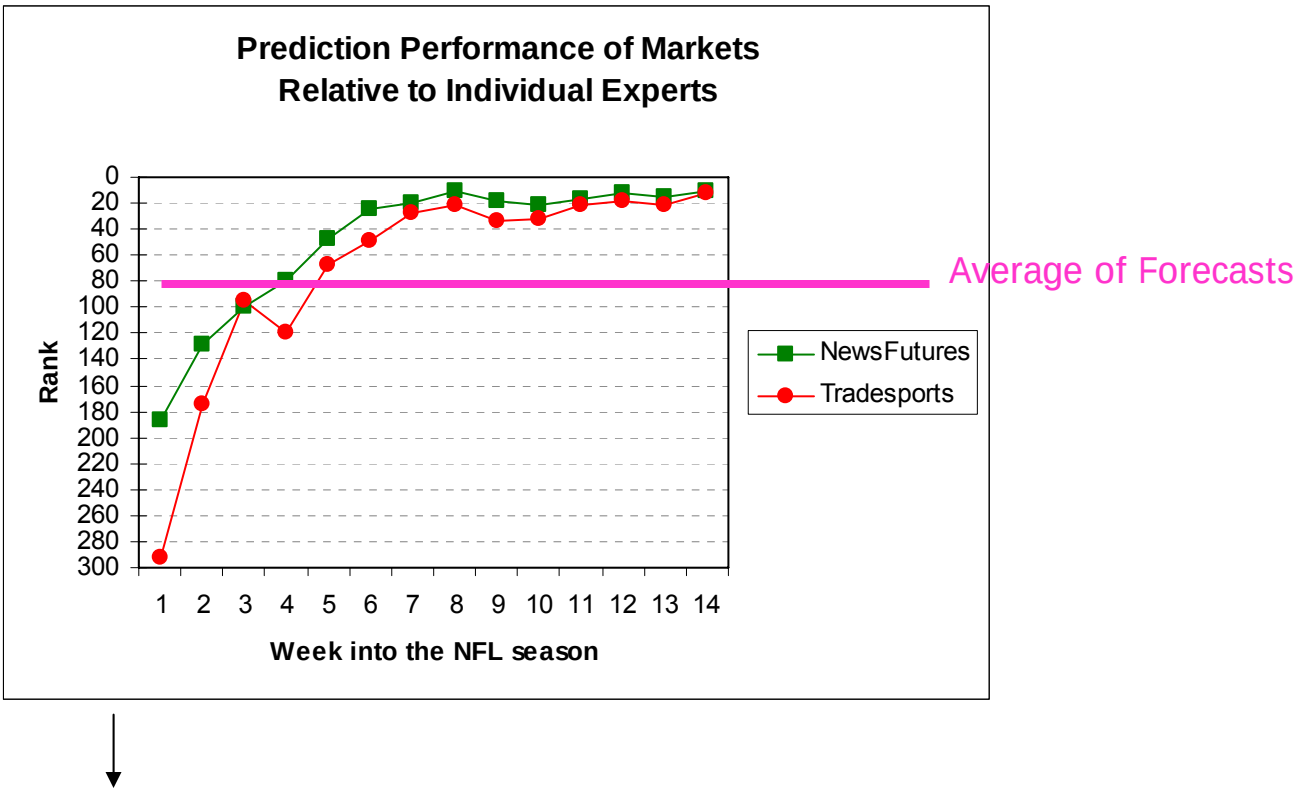

Source: Servan-Schreiber, Wolfers, Pennock & Galebach (2004)

Chen and Pennock (2005) meanwhile also study NFL outcomes and find that, at the same time point ahead of the game, information markets provide as accurate predictions as opinion pools which the authors have carefully pre-selected on grounds of promising performance. More work is needed in this area.

"In-running:" Like predictions obtained from single experts, pooled expert opinion delivers a static spot forecast (unlike prediction markets update themselves continually).

Explanatory content: Just as experts could be asked to justify their predictions, a pool of experts could be asked to submit explanations for their individual forecasts.

Cost efficiency

Costs arise in identifying experts and deciding upon belief aggregation methods. Pooling opinions effectively can be difficult.

Control over Sensitive Information

Using a mathematical approach to pooling individual expert opinions avoids the need to share sensitive information regarding the final forecast with even those experts concerned. A behavioural approach, like the Delphi method, implies that those experts called upon to participate will bear witness to the views of other experts on the matter at hand and will be aware of the final pooled opinion.

Feasibility

Perceived legitimacy: Compared to relying on the wisdom of a large crowd of individuals, consulting a select group of pre-identified experts and somehow pooling their insights perhaps appears intuitively a more legitimate approach to forecasting.

Scope: The task of constructing an expert opinion pool is constrained by space and time, compared to information markets, and the task of eliciting multiple opinions and aggregating these manually could also be constraining in terms of the scope of topics that can be addressed.

Regulation: There are no regulatory impediments to the use of opinion pools.

Participation: Experts must be available. Often a single expert is chosen to specify parameters for each variable. For example, a single person might be asked to estimate the weather in a given area.

2.4 Polls

Polls are surveys of opinion drawn from a particular sample of the population of interest. Answers provided by those in the sample are extrapolated to provide an estimate of popular opinion. Polling appears to have its roots in an early straw vote conducted by *The Harrisburg Pennsylvanian* in an attempt to forecast the 1824 US Presidential election.[41] Early polls were relatively naïve and unscientific (with no systematic effort to take account taken of sample representativeness). In the twentieth century, George Gallup approached the problem of polling in a more scientific fashion, and scored a number of key predictive successes over rivals (including in the 1945 UK national election, where the Gallup poll stood out in correctly forecasting Winston Churchill's defeat). Nowadays, polls are widespread throughout democracies and new technologies such as the Internet make lifted many practical restrictions related to the physical task of polling.

Quality of Information

Accuracy: By comparison with information markets, polls suffer at least two important disadvantages which are likely to undermine the accuracy of their predictions: (i) they weight the views of all individuals equally and not in accordance with how informed these are and (ii) they do not offer incentives for truthful revelation of information. In practice, information markets have outperformed polls in terms of prediction accuracy in many head-to-head tests, particularly election markets.[42] For example, when compared to concurrent major opinion polls on U.S. presidential elections, the Iowa Electronic Market forecasts were more accurate 451 out of 596 times (Berg, Nelson, and Rietz, 2003).

[41] http://en.wikipedia.org/wiki/Opinion_poll

[42] Consider the headlines, "Punters v. pollsters: Are betting markets a better guide to election results than the polls?" *The Economist*, April 14 2005 or "The 'Election Futures' Market: More Accurate than the Polls? As The U. of Iowa Goes, So Goes the Nation?" *Business Week*, November 11, 1996.

Item	1988	1992	1996	2000	All
# big polls	59	151	157	229	596
Poll "wins"	25	43	21	56	145
Market "wins"	34	108	136	173	451
% market	58%	72%	87%	76%	76%
P-value	0.148	0.000	0.000	0.000	0.000

Source: Berg, Nelson, and Rietz, 2003

"In-running:" A further quality consideration is that polls are not dynamically updated, rather present a static snapshot.

Explanatory content: They also have offer nothing by way of explanation for the prediction generated. The information is aggregated into a single statistic (43% said Yes to a particular question) but the reasoning behind the overall response is not captured. This may or may not matter greatly, depending on the context.

Cost

Polls can be cost effective ways to aggregate information, particularly Internet-enabled quick-polling, which is now a standard feature of many news and social-networking websites.

Control over Sensitive Information

The information generated by a poll need not be shared with the participants, if data sensitivity is a concern.

Feasibility

Perceived legitimacy:

Scope: Polls can be conducted in relation to almost any issue, including many that are not amenable to the use of markets (where contracts cannot be written and settled unambiguously e.g. "Should Steve McClaren be replaced as England Football Manager? Yes or No?" or "Would you diagnose a patient with disease A if she presents symptoms X, Y, and Z.? Yes or No?"

Regulation: There are no regulatory impediments to the use of opinion polls.

Participation: A sufficiently large and representative group of participants must be available to be polled.

2.5 Surveys

A survey constitutes a set of questions and other prompts designed to elicit relevant information from respondents. Surveys come in a wide variety of forms, from brief questionnaires with highly standardized answers through to elaborate in-depth interview formats.

Quality of Information

Accuracy: A number of issues can undermine the accuracy of the information provided by a survey. For example,

- Subjects may not be motivated to give accurate answers. Instead they may give answers that present themselves in a favourable light.
- The survey could suffer so-called response bias: those who choose to respond may be different from those who do not respond.
- An aggregation mechanism is needed to weight results and an inappropriate mechanism could bias inferences.

"In-running:" The results of a survey comprise a static piece of information, and not a continuously updated picture.

Explanatory content: A potentially positive feature of surveys is that the information that emerges from a survey may be richer in its explanatory nature than that from markets or polls; the opportunity arises naturally to ask respondents to explain their forecasts.

Cost

Surveys can be very costly to design, implement, and then to analyse once completed – the aggregation of information is not automatic.

Control over Sensitive Information

Whilst a survey generates potentially sensitive information, this need not be disseminated among participants or to third party.

Feasibility

Perceived legitimacy: Surveys may be perceived as quite legitimate if those surveyed are viewed as sufficiently expert and or representative – this will depend on the context.

Scope: Like polls, surveys can be conducted in relation to almost any issue, including many that are not amenable to the use of markets. For instance, they can be used to study attitudes, values, beliefs, and past behaviours.

Regulation: There are no regulatory impediments to the use of surveys.

Participation: A sufficiently large and ideally diverse group of participants must be available to be surveyed.

2.6 Deliberating Groups

Deliberation groups are commonly used in public and private organisations alike. Examples are committees and juries.

Quality of Information

Accuracy: In a process of group deliberation, greater emphasis tends to be given to shared information, and social pressures can lead members not to say what they know (Sunstein, 2006). Four failures tend to arise as a consequence:

- *Amplification of cognitive errors*. In attempting to process information, individuals use heuristics that lead them to predictable errors; they are also subject to identifiable biases. There is evidence to suggest that group deliberation amplifies these errors and biases. For instance, MacCoun (2002) finds that, if individual jurors are biased (perhaps

because of exposure to extra-evidentiary information, such as misleading pretrial publicity or the defendant's unsightly physical characteristics), then not only can juries not be relied upon to correct these biases but they are likely actually to amplify them.

- *Hidden profiles and common knowledge* A "hidden profile" is defined as an accurate understanding which a group (based on the information at its disposal) could but does not obtain. The reason for this failure is known as the common knowledge effect – information which is held by all group members is given excessive weight relative to that which only a few individuals possess (Stasser and Titus, 2003; Wallace, 1999).

- *Cascades.* Sunstein defines a cascade as "a process by which people influence one another, so much so that participants ignore their private knowledge and rely instead on the publicly stated judgments of others." He goes on to distinguish two types of cascade: "In informational cascades, people silence themselves out of deference to the information conveyed by others. In reputational cascades, they silence themselves so as to avoid the opprobrium of others." Both can undermine the accuracy of information generated through the process of group deliberation (see, for instance, Anderson and Holt, 1997).

- Group polarization describes the process by which "members of a deliberating group end up adopting a more extreme version of the position toward which they tended before deliberation began" (Brown, 2006). Group polarization is a well-established tendency supported by hundreds of case studies. All manner of deliberative groups, including groups of judges and juries, will tend to polarize.

A high-profile example of the failure of deliberation to produce accurate information is contained in the 2004 report of the Senate Select Committee on Intelligence, which suggests that the Central Intelligence Agency (CIA) fell prey to groupthink in the run-up to the invasion of Iraq, failing to explore alternatives to the thesis that Iraq posed a serious threat and failing to use the information held within its organization effectively (Sunstein, 2006). As Sunstein records, the report found that despite having its own formalized methods "to challenge assumptions and 'group think,' such as 'red teams' and 'devil's advocacy,' the CIA showed a "tendency to reject information that contradicted the presumption" that Iraq had weapons of mass destruction.

Overall, the evidence suggests that information markets yield more accurate predictions than deliberating groups. This is because information markets offer economic rewards for individuals to disclose the information they have and strong incentives for the marginal trader prepared to exploit biases in prevailing "group think."

"In-running:" Like other traditional IAMs, group deliberation delivers a static spot forecast and does not dynamically update to new information.

Explanatory content: If the deliberative process is tracked by those interested in its outcome then considerable explanation for forecasts can be gleaned.

Cost

Deliberating groups can be costly. There is evidence from inside corporations that information markets may be time- and cost-efficient, compared to using rounds of deliberative meetings to produce forecasts; HP replaced a forecasting process that took many meetings over a period of weeks with a single meeting and an hour or so per person inputting predictions into an internal market.[43]

[43] http://future.iftf.org/2006/12/prediction_mark.html

Control over Sensitive Information

All group members will become party to any information generated through the process of deliberation.

Feasibility

Perceived legitimacy: Deliberation has been found to increase confidence and decrease variance in the group's prediction, and it has been suggested that this gives deliberative groups greater perceived legitimacy than information markets (Sunstein, 2006).

Scope: Deliberative groups can in principle be put in place to consider any issue of interest. Some of the requirements that constrain the use of information markets, such as the ability to define an event unambiguously and settle it (fairly soon) objectively, do not apply. As a result, the scope for their application is potentially very broad.

Regulation: Regulatory concerns do not apply to the use of group deliberation.

Participation: Sufficient individuals in possession of some relevant information and enough time to engage in deliberations must be available, identified, and recruited.

Table II. Feasibility: Information Markets vs. Alternatives

	Expert Opinion	Pooled Expert Opinion	Polls	Surveys	Deliberative Groups	Information Markets
Perceived legitimacy	HIGH – it is intuitive that experts should be consulted.	HIGH – it is intuitive that experts should be consulted.	MED?	MED?	HIGH: Deliberation increases confidence and decreases variance (=> perceived legitimacy)	LOW: at the present time where info markets remain unfamiliar in many settings and can "breed confusion and distrust." Managers may feel undermined and so question legitimacy.
Scope	HIGH?	MED-HIGH?	MED-HIGH	HIGH	HIGH: handle any question.	MED- Not all questions are suitable for mkts.
Regulation	HIGH? – no problems?	HIGH? – no problems?	HIGH? – no problems?	HIGH? – no problems?	HIGH: no regulatory concerns.	LOW-MED - US restricts use of real-money markets. However, money may not matter.
Participation	HIGH	HIGH	HIGH	HIGH	HIGH.	HIGH

Table III. Performance: Information Markets vs. Alternatives

		Expert Opinion	Pooled Expert Opinion	Polls	Surveys	Deliberative Groups	Information Markets
Quality of Information	**Accuracy**	LOW-MED – lack of incentive for honest forecasting?	LOW-MED – difficult to aggregate.	LOW-MED – polls assign arbitrary weight to info.	LOW-MED – how to aggregate results?	LOW (problems with bias, "group-think" polarization, etc.)	HIGH - provided they are well-designed, markets aggregate information more efficiently than other IAMs..
	"In-running"	LOW-MED	LOW-MED	LOW-MED	LOW – snapshot.	LOW?	HIGH: markets yield continuously updated - "in-running" - predictions.
	Explanatory content	HIGH – can ask expert to justify.	HIGH – can ask for explanations.	LOW-MED no explanation for voting by default but could use wikis	HIGH – surveys can elicit explanations for forecasts/views.	HIGH – principal can understand the process that lead to collective forecast.	LOW-MED-Design dependent; if wikis/blogs used to elicit trading motivations then okay.
Cost-Efficiency	**Setup**	LOW-MED (need to identify experts)	LOW-MED (identify experts and design mech.)	HIGH – easy and cheap to setup online.	LOW – costly to set up.	MED – identify participants, schedule meetings etc.	LOW-MED – markets need to be designed and implemented technologically. Participants may need training.

	Operation	HIGH	HIGH	HIGH – easy and cheap to run.	LOW – costly to administer/analyze	LOW – deliberation can take long time.	MED-HIGH- once operational markets can be very cheap to run.
	Control of sensitive info,	HIGH (only expert and principal need know the forecast etc.)	HIGH (experts and others need not see collective intelligence)	HIGH (participants need not see collective intelligence)	HIGH (participants need not see collective intelligence)	LOW-MED. All group members see sensitive information.	LOW-HIGH. By default all participants see current forecast. However, some firms (e.g. HP) are experimenting with designs which restrict the visibility of market aggregates.
	Positive side effects e.g. fun	LOW?	LOW	MED - people like to have their voices heard.	LOW – people like to have their say but may resent long surveys.	LOW? People tend to find meetings low fun.	MED-HIGH – trading can be fun and people like to have their say (an lower agency costs inside an organization). However, pessimistic forecasts can upset morale and motivation.

3. Conclusion and Directions for Future Research

Information markets are markets designed specifically for the purpose of aggregating information to produce a collective forecast of an uncertain event. This paper compares information markets to other Information Aggregation Mechanisms (IAMs) in terms of their feasibility and performance. The other IAMs considered are: expert judgment; pooled expert opinion, group deliberation; and conducting a poll or survey.

A prerequisite for this comparative exercise is a view on how the performance of an IAM should be measured. Performance, we suggest, must relate to the major costs and benefits in each case. Considering the benefits of IAMs, the obvious consideration is the quality of information it generates. Based on our review of experiences in the field, we identify several dimensions of quality that matter, chiefly: prediction accuracy; whether information is updated "in-running;" and (depending on the setting) the explanatory content of the information. Evidence from the field suggests that information markets, carefully designed and implemented, tend to provide superior forecast accuracy. They also deliver predictions that are continuously updated, whereas many IAMs, including surveys and polls, produce a a static forecast, which is not dynamically revised to incorporate any new information. Turning to consider cost-efficiency, a natural concern is the time and money required to establish and operate the IAM. For instance, information markets can be very cheap to set up, if an off-the shelf platform is rented from one of the many technology specialists, or very expensive (certainly compared to such alternatives as simple polling) if a more bespoke in-house platform is required. In both cases, once the information market is established operating costs tend in practice to be very low, particularly low compared to the costs associated with ongoing group deliberation. Evidence from the field suggests that less explicit costs that arise can also be important. Depending on the IAM in question, these extra costs include the potential distraction to employees (e.g. excessive market-watching in the information markets case) and possible loss of control over any sensitive information the IAM generates (e.g. it may be commercial undesirable for employees to see a pessimistic collective forecast regarding the likelihood that a product will ship on time). Somewhat offsetting these costs, there could be additional positive side-effects of IAMs; participation may or may not be fun. Alongside performance, it is natural to consider the overall feasibility of each IAM, since without feasibility performance is irrelevant. Dimensions of feasibility we put forward for consideration in this report include: the scope of the questions that can be asked; the perceived legitimacy of the mechanism; possible regulatory concerns in each case; and the participation prospects.

Taking this fairly broad view of comparative performance and feasibility, a mixed picture emerges (summarized in tabular form in Tables II and III, p.30). Not only is the feasibility of IAMs context-sensitive (for instance, information markets cannot be used to aggregate information in some settings due to the need to write unambiguous contracts and settle markets objectively) but even where all IAMs are feasible, the optimal choice of mechanism (as judged by the balance of costs and benefits) depends on the setting: is this a corporate setting or a public interest matter? If corporate, is the data generated of a sensitive nature? Is it important that the IAM not only predict but also offer some explanation for the forecast? The problem of information aggregation does not have a one-size-fits-all solution.

The study highlights many significant gaps in the prevailing understanding of IAMs, pointing to a compelling agenda for further research. There have so far been very few "real-world" experiments to test prediction markets in the academic literature; most articles cite predictions from the established public event markets such as the Hollywood Stock Exchange, the Iowa Electronic Markets, and TradeSports' sports and events markets. Christiansen's (2007) work, which considers a small, relatively illiquid market with a dearth of informed traders and finds this to be well-calibrated is a useful exception, but further real-world experimentation is greatly needed. This research should be undertaken as a priority and conducted in a scientific and "open" way, ideally involve randomized trials, conducted in realistic field settings. This is particularly true where information markets are concerned, since market adoption is still at an early-stage and much of the interesting experimentation has been carried out in academic laboratories or by corporations behind closed doors.

Predicting the performance of information markets remains to a large extent guesswork. For instance, early insights suggest that information markets can deliver more accurate predictions, but despite papers such as Spann and Skiera (2003), still very little is understood about the impact of market design features on information accuracy, and practitioners and the academic literature alike have had very little to say on the cost side. Information markets almost certainly have operational limits (they appear to work best in `thick' situations, where many traders possess small bits of information, behave competitively, are good at Bayesian updating, and also do not have a significant stake in the outcome) but we have only a partial understanding of these. Comparative analysis, which puts the market-mechanism head-to-head with alternatives such as quick voting, should be a further priority.

The possibility of using conditional information markets as decision-making guides deserves further research. Information markets have the potential to revolutionize the way decision-making is guided in many setting, including ultimately the policy realm. However, the point has been made that detailed decision analysis could require potentially many conditional markets, which would tend to thin markets out (Ledyard, 2006). Research should explore this issue (and possible workarounds) systematically. Alternative market designs, such as call markets and market scoring rules, have been proposed to address some of the problems such as thinness, but more work is needed to gauge the effectiveness of these mechanisms across different real-world settings. Further complications that warrant the time and attention of researchers include the implications of running markets in which traders have a large stake in the outcomes of interest and often also exclusive and important pieces of information that they may withhold strategically (as often may arise in business settings). A further exciting avenue for further research lies in exploring whether some combination of predictions information markets and other channels (e.g. polls) could offer superior forecasts (Chen and Pennock, 2005). The relaxation of existing US regulatory provisions would make more viable (Arrow et. al, 2007) these and other related lines of research.

References

Anderson, Lisa, R. and Charles A. Holt (1997): "Information Cascades in the Laboratory," *The American Economic Review*, 85(5).

Armstrong, J. Scott (ed.) (2001): *Principles of Forecasting: A Handbook for Researchers and Practictioners*, Springer.

Arrow, Kenneth, et al. (2007): "Statement on Prediction Markets," *AEI-Brookings Joint Center for Regulatory Studies*. This paper can be downloaded free of charge from the Social Science Research Network at: http://ssrn.com/abstract=984584.

Berg, J., F. Nelson and T. Rietz (2003): "Accuracy and Forecast Standard Error of Prediction Markets," Departments of Accounting, Economics and Finance, Henry B. Tippie College of Business Administration, University of Iowa.

Brown, Roger (1986): *Social Psychology: The Second Edition,* New York, N.Y.: Free Press.

Chen, K-Y and D. Pennock (2005): "Information Markets vs. Opinion Pools: An Empirical Comparison," Tech. rep., California Institute of Technology.

Chen, K-Y and C. R. Plott (1998): "Information Aggregation Mechanism: Concept, Design, and Implementation for a Sales Forecasting Problem," Tech. rep., California Institute of Technology.

Christiansen, Jed D. (2007): "Prediction Markets: Practical Experiments in Small Markets and Observed Behaviour," *The Journal of Prediction Markets*, 1, pp. 17–41

Clemen, R. T., and R. L. Winkler (1999): "Combining Probability Distributions from Experts in Risk Analysis," *Risk Analysis*, 19(2):187–203, 1999.

Cover, Thomas M., and Joy A. Thomas (2006): *Elements of Information Theory*, Wiley.

Croxson, Karen and J. James Reade (2007): "Information and Efficiency: Goal Arrival in Soccer Betting," *mimeo*.

Fama, E.F. (1970), "Efficient Capital Markets: A Review of Theory and Empirical Work," *The Journal of Finance*, 383.

Figlewski S. (1979): "Subjective Information and Market Efficiency in a Betting Market," *Journal of Political Economy,* 87(1):75–88.

French, S. (1985): "Group Consensus Probability Distributions: A Critical Survey," *Bayesian Statistics*, 2:183–202, 1985.

Genest, C. (1984): "A Conflict Between Two Axioms for Combining Subjective Distributions," *Journal of the Royal Statistical Society*, 46(3):403–405.

Genest, C. and J. V. Zidek (1986): "Combining Probability Distributions: A Critique and an Annotated Bibliography," *Statistical Science*, 1(1):114–148, 1986.

Gruca, Thomas, S., Joyce Berg, and Michael Cipriano (2003): "The Effect of Electronic Markets on Forecasts of New Product Success," *Information Systems Frontiers* 5 (1): 95-105.

Hanson, R. (2003): "Combinatorial information market design," *Information Systems Frontiers*, 5, 107–119.

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

Kadane, Joseph B., and Robert L. Winkler (1988): “Separating Probability Elicitation from Utilities,” *Journal of the American Statistical Association*, 83(402):357-363.

LaComb, Christina Ann, Janet Arlie Barnett, and Qimei Pan (2007): “The Imagination Market,” *Inf Syst Front*, 9:245–256

Ledyard, John O. (2006): “Designing Information Markets for Policy Analysis,” in *Information Markets: A New Way of Making Decisions*, AEI Press.

Linstone, H. A. and M. Turoff (1975): *The Delphi Method: Techniques and Applications*. Addison-Wesley, Reading, MA.

MacCoun, R. (2002): “Comparing micro and macro rationality,” In M. V. Rajeev Gowda and Jeffrey Fox (Eds.), *Judgments, decisions, and public policy.* New York: Cambridge University Press.

Ortner, G. (1998): “Forecasting Markets - an industrial application, Part II,” *Working Paper*, Technical University of Vienna, Department of Managerial Economics and Industrial Organization..

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

Pennock, D. M., Giles, C., Nielsen, F. (2001): “The Real Power of Artificial Markets,” *Science,* 291:987–988.

Rhode, P. W. and K. S. Strumpf (2006): “Manipulating political stock markets: A field experiment and a century of observational data,” *University of North Carolina Working Paper*.

Roll R. (1984): “Orange Juice and Weather,” *The American Economic Review,* 74(5):861–880.

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

Spann, M., and E. Skiera (2003): “Internet-Based Virtual Stock Markets for Business Forecasting,” *mimeo*.

Stasser, Garold and William Titus (2003): “Hidden Profiles: A Brief History,” *Psychological Inquiry*, 14 pp. 304.

Sunstein, Cass R. (2006): “Deliberation and Information Markets” in *Information Markets: A New Way of Making Decisions*, AEI Press.

Surowiecki, James (2004): *The Wisdom of Crowds*, Little, Brown.

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

Wallace, Patricia (1999): *The Psychology of the Internet*, Cambridge University Press, 82.

Wolfers, Justin and Eric Zitzewitz (2004): "Prediction Markets," *Journal of Economic Perspectives*, *18*(2), pp. 107-26.