EXPERT SERIES

Prediction Markets

A Practitioner's Guide

1/14/2013

Dr. Alpheus Bingham
Consulting Partner
YourEncore

Yiftach Nagar
Massachusetts Institute of Technology

EXPERIENCE MATTERS

INTRODUCTION

A few years ago, Microsoft used an internal prediction market to forecast whether a certain internal software tool would ship on time. Three minutes after trading began, the share ‘price’ for the ‘on time shipment’ contract was down to 3 cents per share, indicating a majority view that the chances for shipping on time were slim. As the story – which has become part of the lore of prediction market enthusiasts – goes, the project manager quickly held a meeting in which the decision to strip out some features was made. The price of the on-time contract jumped up. Then, internal customers complained and the features were put back into the product. The on-time contract price went down again. Eventually, the product shipped several months late, as the market predicted.

At Hewlett-Packard, prediction markets helped reduce errors in price predictions of a computer memory component called DRAM (a highly volatile commodity) – enabling improved purchasing timing with the subsequent financial benefits. Using the markets, HP was also able to shorten the forecasting process from multiple meetings over weeks to a single meeting and an hour-or-so per person.

Microsoft and HP are not alone. In the past decade, following growing evidence of the success of prediction markets in different settings (industry, government, academic and public markets), many other companies, like Best Buy, Eli Lilly, France Télécom, General Electric, Google, IBM, Intel, Siemens, and Yahoo! have deployed prediction markets. These companies have started experimenting and using them to augment traditional methods and processes used for forecasting sales, schedules and events.

What drives these companies to use prediction markets? What are prediction markets anyway? Are they really a forecasting panacea? How do they work? Under which conditions? And how can your company gain from them? This short paper provides an overview and offers some answers to those questions.

WHAT ARE PREDICTION MARKETS? HOW DO THEY WORK?

Prediction markets are markets where participants trade in contracts whose payoff depends on the outcome of future events (Wolfers and Zitzewitz 2004). The simplest example usually used to explain how these markets work is a *‘winner takes all’* market, e.g. a presidential-elections prediction market. Participants can buy and sell a contract (Contract A) that says ‘Candidate A will win the elections’ for a price that ranges 0-100 cents. The market opens some time before the elections and is then closed just before the winner is announced. Each unit of Contract A then pays $1 if candidate A has won, and nothing otherwise. Usually there will be another contract (Contract B) for ‘Candidate B will win the elections’ traded in the same market (and additional contracts if there are additional candidates).

What should a rational participant do? If you believe there is a 67% chance for Candidate A to win the elections, then you also believe the mean expected value of the respective contract to be 67 cents. Therefore you should be willing to buy units of that contract if and only if the price is less than or equal to 67 cents. (If it is trading much above this, then a rational person may well opt to sell the contracts they are holding.) Importantly, note that this has nothing to do with which candidate you would *prefer* to win based on your political views. Some other participant may believe that Candidate A has a higher chance (say, 80%) of winning. That participant may then be willing to buy for up to 80 cents per contract unit. The overall market price for a contract is an aggregate of the beliefs of all market participants. It represents the market expectation, or in other words, the aggregate expectation (not necessarily the “average” expectation as will be explained) of all participants, that the event of interest will indeed occur.

Other types of contracts can be constructed to estimate other things. For example, an *index* contract that pays $1 for every $1M of sales of Product A in a certain quarter. The market price of this contract represents the market expectation of the sales of Product A during the quarter. In this case, the market will only have one contract. Alternatively, rather than trying to get a point estimation of sales, a market can be constructed to get an estimation of the probability distribution of sales using intervals. In that case, there may be, for example, the following contracts traded in the market: A (sales will be $0-$1M), B (sales would be $1M<S<$2M), C (sales would be >$2M). Each contract pays $1 if and only if sales fall in the range described in the contract. In this variation of the winner-takes-all market we can get an entire probability distribution.

Additional types of contracts can be constructed to help estimate various measures such as the median or the standard deviation of a distribution, etc. In practice, the winner-takes-all and index contracts are most common and are used to estimate project schedules, sales, prices, events, etc.

WHY DO THEY WORK?

Like financial markets, prediction markets are mechanisms that aggregate information and knowledge. Many economists subscribe to the efficient market hypothesis (Fama 1970), according to which share prices incorporate all relevant information (and that, therefore, no solo trader can ‘beat the market’), and some would use it to explain the success of prediction markets. However, this hypothesis has also been disputed by behavioral economists who have shown that investors are not rational, but rather, suffer many biases; and by others who pointed to the many inefficiencies observed in real markets. It is also not very likely that prediction markets are completely efficient, and in actual practice they are, at times, subject to bubbles and biases.

By the less restrictive marginal trader hypothesis (Forsythe et al. 1992), the presence of not-so-rational participants in the market is not a problem: It is enough that there are *some* traders who are rational and well informed. In case the price does not reflect all available knowledge, those traders can make profits, which they are motivated to do by the payout structure of the markets. And as they trade, they will drive the market price in the “right” direction.

A vast body of theoretical and empirical research has shown that combining forecasts from multiple independent, uncorrelated sources (whether they are statistical models, human experts, or even non-experts) that have relevant knowledge and information, leads to increased forecast accuracy (Armstrong 2001; Clemen 1989). Why does it work? In a sentence: Each individual prediction may contain bits of truth mixed with various misconceptions. The bits of truth are correlated with each other so they add up to a larger truth, whereas biases, errors and misconceptions are not correlated with each other (because they are independent) and therefore cancel each other. This is captured nicely in Page’s Diversity Prediction Theorem (Page 2007), although others have noticed and explained this before. One derivative of that is that a large number of diverse non-experts can make predictions that are better than those of a small number of experts. Indeed, this is the basis of the ‘wisdom of crowds’ described by Surowiecki (2004) and Sunstein (2006).

There are many ways to combine predictions and expertise, from face-to-face groups to various ‘mechanical’ ways, like simple and weighted averaging, majority voting, and various other statistical ways like bootstrapping and Bayesian learning. Research has shown that ‘mechanical’ aggregation has many advantages over various face-to-face modes of prediction and decision making (Sunstein 2005, 2006). Theoretical and empirical comparisons have further shown that no single method of combination is best under all circumstances. Why markets then? One key factor in the success of prediction markets is that they can *incentivize* participation from a *diverse* crowd. The lure of markets is due not only to potential financial gains, but also (especially in play-money markets) to the fun involved in betting.

Perhaps the most popular explanation for why prediction markets do better than other mechanisms (mainly: polls and pundits), is that markets drive people to put their money where their mouth is. Servan-Schreiber (2012) proposes that when predictions are posed as wagers, rather than as requests for 'unaccountable opinions', and so long as real incentives are involved (prizes, knowledge, recognition) people employ the same risk-considerations as they would in real-money markets. Indeed, most empirical comparisons found no difference in accuracy of play-money and real-money prediction markets. It is also important to note that markets not only provide an incentive for *sharing* information: They also incentivize information gathering.

"The lure of markets is due not only to potential financial gains, but also to the fun involved in betting."

Finally, an important note should be made: There are several different designs for implementing the mechanism of prediction markets, such as continuous double auctions, pari-mutuel markets, and using different kinds of market makers. The reasoning outlined above for what drives markets' accuracy does not depend on any specific market mechanism. Empirical evidence showing prediction market performance was collected in various settings, and conclusions seem to hold robust across many configurations. We discuss later on some conditions that are important for successful implementation. The market mechanism itself does not seem to be one of them.

PREDICTION MARKETS VS. OTHER METHODS

Experts, Models, Meetings, Polls and Markets

There are many ways to predict the future, and of course none is perfect. Perhaps the first method, since the days of the oracle at Delphi, is to ask experts. Unfortunately, like all of us, experts are humans, with bounded rationality, many well-documented cognitive biases and limitations, and too often, unjustified overconfidence. Hundreds of empirical studies, in many domains, have shown that the ability of experts to provide reliable predictions is astonishingly poor. Many studies have shown that almost always, simple statistical models can outperform experts.

The trouble with models is that they take time and effort to develop, and that they are limited in context. They can work very well under well-prescribed assumptions and circumstances, but they are less applicable in more dynamic or "fuzzy" environments, where patterns are not easy to discern – or may not exist at all. While a model can do a good job of predicting sales of a familiar commodity, factoring in some economic parameters, seasonality, etc., it seems much less likely that a model can help predict sales of a completely new type of product, and even less likely that a model can predict strategic moves of competitors.

Another common way of prediction and decision making is using face-to-face meetings. The organizational literature is filled with descriptions of what every manager knows: meetings are a poor way for creating good prediction. Phenomena such as information cascades, group polarization, groupthink etc. interfere with the process of eliciting and weighing information. Sunstein (2006) and Thompson (2012) offer thorough discussions of the limitations of face-to-face groups.

One common conclusion is that mechanically aggregating predictions yields much more accurate predictions, and for reasons discussed above, prediction markets seem better than other ways of doing so. But do they really succeed in practice?

Prediction Markets Accuracy: Just how well do they work?

In the public sphere, open prediction markets have been shown to empirically produce predictions that are better than those of polls, experts, and even statistical models – in many different domains, such as presidential elections (e.g., Berg, Forsythe, et al. 2008), sports (e.g., Spann and Skiera 2009), movie sales, and other domains (e.g., Pennock et al. 2001). Prediction markets also fared well against other methods of combining predictions such as simple averaging, weighted averaging and logarithmic regression (Berg, Nelson, and Rietz 2008; Y. Chen et al. 2005).

What happens in companies? While less data is publically available, there are encouraging signs that prediction markets often yield accurate and reliable results. In early experiments done at HP, for example, sales forecasts created in markets were more accurate and more stable than official company predictions in a vast majority of cases (K.-Y. Chen and Plott 2002; Malone 2004). Intel has tested markets for predicting product demand, and similarly found them to be at least as accurate (often better; and as much as 20% more accurate) than official forecasts, “an impressive result given that the official forecasts have set a rather high standard during this time period with errors of only a few percent” (Hopman 2007, 131). In one case, the market indicated no confidence in the official forecast, which later indeed turned out to be off significantly (Hopman 2007).

The most comprehensive publically available data on corporate prediction markets comes from Google, who has been running prediction markets since 2005. During the first 2.5 years, Google ran 270 markets for making predictions about demand, performance, industry news, and other subjects, in which over 1,400 people participated. Overall, market prices closely approximated actual event probabilities (Cowgill, Wolfers, and Zitzewitz 2009).

Additional Advantages

As Nagar and Malone (2011) note, prediction markets may be appealing in some settings for reasons beyond accuracy improvement.

By tying compensation to performance while giving participants a sense of both fun and challenge, they serve to increase both extrinsic and intrinsic motivation, and lead to the increase of attentive participation. Participants have an economic incentive for both gathering more information that would improve their performance in the markets, and for sharing it through trading.

The markets can help expose uncomfortable truths. As one manager said: “If you let people bet on things anonymously, they will tell you what they really believe because they have money at stake. This is a conversation that’s happening without politics. Nobody knows who each other is, and nobody has any incentive to kiss up” (Brodkin 2008). Further, the actions of traders within them can identify and reveal new experts. We discuss that point at more length below (see ‘The Village Well’).

IMPLEMENTING ORGANIZATIONAL PREDICTION MARKETS

Selecting Participants: Who should participate, and how many?

Theoretically, the more people who have relevant information who participate, the better (as explained above). In practice, there are of course trade-offs involved in that (e.g. the need to train people). Some experiments in companies yielded calibrated[1] predictions with even as little as twelve experts participating in any given contract in the market. We would still argue that while 12-20 people are a good number for initial pilots, it is probably a better idea to include more people. The total is dependent on the degree of trader overlap for each contract (or dependent collection of contracts). Think dozens per contract and 100+ for an overall market with reasonable diversity of predictions such that each trader trades in about 20% of available contracts.

Importantly, the people who participate in the market should be diverse. We would propose to push for diversity on several dimensions: discipline (engineering, marketing, sales, etc.); rank (include both senior and junior people); geography, etc. An important pattern revealed at Google was that people who sit near each other (e.g. in nearby cubicles, or in offices on the same floor), showed similar trading patterns. There was much less correlation across trades from people on different floors.

"Importantly, the people who participate in the market should be diverse."

One way to draw larger participation and to potentially raise diversity is to extend the trader population. This idea may seem radical to many people, and obviously might not be appropriate for any market. For example, there may be questions that management cares about which cannot be exposed too broadly for legal reasons. But the principle of opening the market to many people, rather than small groups has advantages. Self-selection could draw the more knowledgeable people, and you may not even know who these are sometimes.

Another option to consider is to include experts from outside the company. While this approach may not be applicable in every situation and for every type of question (e.g. you might not want to expose some proprietary data), it has several advantages. First, it can help avoid information cascades – especially if those experts have no ties inside the company, and among themselves. Second, it provides a way to outsource some (or even all) of the activity, and thus simplify the operation of markets.

How to Motivate Participation

A key challenge in implementing organizational collective-intelligence systems such as wikis, blogs, etc. is motivating participation. In a business environment, the main obstacle is the duo of opportunity cost, and seeing the value. People are often skeptical about new systems, and it can take time to convince people of the value they can get. In "play money" markets, there are no real financial gains to be made, so why would busy people 'play'?

[1] When a weather forecaster predicts "60% chance for rain", how do you estimate the quality of her predictions? One measure is calibration. Look at her record over a long period. If it rained in 10-20% of the days of which her predictions were in that range (10-20%), on 20%-30% of the days her predictions were in *that* range, and so on, then she is "perfectly calibrated". To do this assessment, we usually bin the predictions in 10% ranges as explained, and draw a graph of actual vs. expected. A straight 45 degrees line is the ideal. Of course the same measure applies to any kind of probabilistic prediction.

The good news is that companies do manage to attract employees to participate in the markets by combining some or all of the following measures:

Incentives

Small prizes like T-shirts, lunch vouchers, etc. are often given as a prize to the best predictors – often tracked within the market platform as a “leader board.” Another approach is to hold lotteries/raffles for winning small gadgets, etc., where the amount of ‘tickets’ one gets is proportional to one’s “cash balance” (so better performers have a higher chance of winning the lottery). At Google, simply making public the identities and announcing top traders proved to be effective in itself. It helps create a ‘buzz’ and make things fun by adding a competitive element even in the absence of money.

‘The Village Well’

A market place can be more than just a mechanism of matching buyer bids and seller offers. Much like financial websites (e.g. Yahoo! Finance), that incorporate conversation platforms, there are prediction markets that incorporate conversation forums in the platform. If these are not a part of the software itself it is still worthwhile to add them even as a separate webpage. Such discussion forums (regardless of the exact implementation) can foster productive controversy, which contributes to the ‘buzz’ around the markets, and which may expose additional information. Adding a component of “why” to the pricing information of “what”, these conversations can be linked to trading data and strengthen the creation of business intelligence. Even beyond that, they sometimes help to expose specific subject matter experts that may not otherwise be known. This is a distinct possibility even in closed (internal to organization) markets and certainly a frequent occurrence when markets are opened to external traders.

Including ‘Fun’ Markets

About 30% of the markets reported in Google’s paper were ’entertainment’ markets, on topics not related to the company’s business (e.g. markets about sport events, movies, etc.). These helped in creating the buzz, generating active participation and frankly making things fun for employees. Such markets help keep engagement high because once people logged in to trade in those markets, they also were more active on the professional markets (King 2006).

Lowering Barriers to Entry By Simplifying the Interface

Many people may find trading in markets intimidating, and learning how to so, a daunting task. But as discussed above, prediction markets work well as an information gathering and processing tool regardless of the exact mechanism implemented. Several vendors have realized that the learning curve was too steep and now offer simple and intuitive graphical user interfaces that do not require people to know how to undertake complex trades, like selling short. At the same time, oversimplification and inadequate “pay out” descriptors run the risk of collapsing the experience to survey taking or multi-voting and may erode some of the distinct edge created by the prediction market over just such tools, useful elsewhere.

Scheduling the Markets

In a business environment, opening the market occasionally for short ‘bursts’ may be a good idea – it helps focus the attention of people (“the market is now open!”), and it also helps with time management. Of course, there is a problem if not everyone is available to trade when the markets are open. It is possible to create a schedule such that, for example, markets are run on a quarterly basis, and they are opened for one day every week, or for 3 hours every day, etc. Obviously, it is difficult to make a concrete suggestion here that would apply for every business and every market. This is one of the many parameters that should be tailored to your specific circumstances, needs and objectives.

Finally, although public markets, such as presidential election markets, run for months, prediction markets can actually serve to rapidly aggregate information from many people to create predictions. In lab experiments conducted at MIT, markets that ran for a few minutes only provided valuable predictions (Nagar and Malone 2011). Here again, the trading duration can be set for "spot results" or it can provide what amounts to a running external analysis, updated in real time over periods where the environment is in flux. Your business intelligence needs will dictate which to use and when, but most commercial market platforms readily accommodate both in the same overall marketplace.

Technical Aspects

Once you have decided to implement prediction markets, there are some technical questions, including: What software platform to use? What market mechanism to use (continuous double auction, pari-mutuel, whether to use a market maker, and if so, which?), etc.

We argue that these questions are minor and should not guide your choices nor paralyze your speed to adoption. Using a market maker can help simplify trading, and creating liquidity in the markets. Beyond that, to date, there is no research that directly compared those mechanisms and found advantages of one over the other. Nor is there a sound theoretical argument in indisputable favor of any mechanism. As explained above, the power of prediction markets comes from their ability to aggregate diverse, informed predictions, and to create candid responses based on "payouts," as opposed to personal wishes for certain outcomes, rather than from any specific "trick".

As for software platforms, there is no single platform that has acquired complete market dominance. There are several good packages out there, each with its own strengths and weaknesses. A good integrator should be able to offer some flexible options.

Management Support

Management support is often needed for two basic reasons: 1) where internal traders are used, they need to know that this is a valued part of their contribution and will not be considered a distraction from "their day job." And, 2) decision-makers may need to incorporate market predictions in their management choices and integrate the market data with other forms of analysis. We would never suggest that crucial business decisions be simply "outsourced" to the markets and this decision integration is a crucial part of organizational effectiveness. If prediction markets are new to your organization, you may want to identify consultants who can speak knowledgeably not only about the market implementation but the integration of market outcomes into the overall organizational processes. We have more to say on this subject below.

Integrating Markets into Organizational Decision-Making

What would you use prediction markets for? A recently published book, *Oracles: How Prediction Markets Turn Employees into Visionaries* (Thompson 2012), opens with an extreme example: A Rhode Island IT company called Rite-Solutions, (http://www.ritesolutions.com) founded in 2000, created a prediction-markets-based democracy in which all of its employees make most of its business decisions. This model, of course, is not suitable for most companies. So where should you focus and what can you gain?

Markets have been successfully used in companies to predict product demand, product performance (e.g. number of bugs), and deliverable timing, and to predict external events such as industry news, customer practices, competitor actions, regulatory changes or legislative policies. Integrating the predictions into the organizational decision-making is a significant challenge that requires deliberation and planning. In each of these specific domains there are specific challenges. For example: when predicting project schedules some managers may worry about self-fulfilling prophecies. When predicting sales, you might get market indications that are different from those created in established traditional ways, and this might get people worried (e.g. sales people whose bonuses depend on those predictions). What do you do when you get such a situation where the newly implemented markets and the established systems are at odds? Should you reflexively say the market was right and question your established system? Before embarking on a project, it may be a good idea to consider such scenarios. For example, in this case, although you are not likely to question your entire tried and true way of making predictions, finding big gaps can serve an appropriate trigger for some inquiry of facts and underlying assumptions (as described in the Microsoft anecdote in the beginning of this paper).

"Integrating the predictions into the organizational decision-making is a significant challenge that requires deliberation and planning."

Remember that even classical business intelligence gathering often produces data in support of or in conflict with any given course of action. These data sources are no different. But we'd argue that to avoid or ignore them entirely is at the peril of improved decision-making. The knowledge atomized within a crowd when properly aggregated has been consistently robust as opposed to solo pundits hampered by biases, agendas and very deep, but simultaneously, narrow perspectives.

Interpreting Prices in Prediction Markets

In theory, contract prices in prediction markets are accurate estimations (of probabilities, numbers, etc., depending on the type of contract). Page's Diversity Prediction Theorem is often correct in the assumption that independent participants' errors are not correlated. In practice, however, many participants are correlated, whether that is due to information cascades (for people in the same department), or due to common biases such as optimism. These biases can accumulate and affect the predictions. For instance, although markets at Google were generally well calibrated and diversified, they did reveal several types of bias, including an optimistic bias, and a reverse favorite-longshot bias (i.e. people overestimated the chances of likely events, and underestimated the chances of unlikely events). These biases were more common among new employees and people new to the markets, and declined over time. Even so, all markets, as we are all too familiar with in 2012, create bubbles. Prediction markets are no different and decision-makers must remain alert to such factors while extracting the value we argue for in this paper.

At the bottom line, market output is a signal for management, not a management decision. There may be things management knows which are not known to market participants. It may be the other way round.

Expect it to take some time for management to learn how to integrate markets into effective decision-making. This is somewhat analogous to doctors who have worked with X-rays, and now receive a CT scanner. The CT scanner has a higher resolution and can find things you won't see with X-rays. But it takes some time to learn how to use it and interpret its output. Therefore, it is recommended to start the process with some guidance from experienced consultants.

Are Prediction Markets Prone to Manipulation and Bubbles?

Theory argues, and experiments have shown that prediction markets are quite robust against manipulation (in fact one study even argued that a manipulator can aid prediction markets accuracy (Hanson and Oprea 2009). The collective intelligence of market participants is quick to identify opportunities for gains, and so, fluctuations that are based on manipulation rather than on real information tend to be quickly corrected. Manipulation is of course much less of an issue in organizational markets to begin with.

Bubbles can happen, theoretically, due to information cascades, as they do in financial markets. But in organization markets that run for relatively short times, and with a relatively small number of people, again, the problem is less likely. Bubbles are also much less likely when the participant population is diverse – though management must always acknowledge their possibility.

Combining Models and Markets

At a more advanced stage, once the organization is more familiar with markets, it is also possible to consider the market as a mechanism to aggregate predictions not only from people, but also from models (robotic trading). Experiments at MIT (Nagar and Malone 2011) showed that this has the potential to improve the accuracy and robustness of predictions. Google also reported that some employees developed bots that traded in its markets.

CONCLUSION

Companies today face fierce global competition, constant scrutiny from investors and regulators, and quickly shifting consumer demands. To contend with these challenges and succeed in the marketplace, executives need better business intelligence than ever before.

What is the potential market size for a new product? How will competitors react? What will happen to raw material costs? Or, as was the case with Microsoft at the beginning of this paper, will our products ship on time?

This paper highlighted how prediction markets offer an accurate way to inform critical business decisions. But designing and managing the market, recruiting participants, distilling the results and integrating them into decision making brings a new set of challenges.

With a respected team of experts in prediction markets, business intelligence and knowledge management, integration partners can offer companies an efficient way to launch a prediction market. They can draw on an external network of subject matter experts to identify participants who are both objective and knowledgeable in the field. This approach saves companies the cost and complications of choosing and implementing a software platform.

Predicting the future will always be an imperfect science. However, a thoughtful approach to prediction markets may be the most accurate and efficient way to inform key business decisions and achieve marketplace success.

References

Armstrong, J. Scott. 2001. “Combining Forecasts.” In *Principles of Forecasting: A Handbook for Researchers and Practitioners*, edited by J. Scott Armstrong, 417–439. New York: Kluwer Academic Publishers.

Berg, Joyce, Robert Forsythe, Forrest Nelson, and Thomas Rietz. 2008. “Results from a Dozen Years of Election Futures Markets Research.” In *Handbook of Experimental Economics Results*, vol. 1, edited by Charles R. Plott and Vernon L. Smith, 742–751. Amsterdam: North-Holland / Elsevier.

Berg, Joyce E., Forrest D. Nelson, and Thomas A. Rietz. 2008. “Prediction Market Accuracy in the Long Run.” *International Journal of Forecasting,* 24 (2): 283–298. doi:10.1016/j.ijforecast.2008.03.007.

Brodkin, Jon. “Google Bets on Value of Prediction Markets.” *Network World*, March 5, 2008. http://www.networkworld.com/news/2008/030508-google-prediction-markets.html.

Chen, Kay-Yut, and Charles R. Plott. 2002. “Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem.” Social Science Working Paper 1131, Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA.

Chen, Yiling, Chao-Hsien Chu, Tracy Mullen, and David M. Pennock. 2005. “Information Markets vs. Opinion Pools: An Empirical Comparison.” Paper presented at the Proceedings of the 6th ACM Conference on Electronic Commerce, Vancouver, BC, June 8.

Clemen, Robert T. 1989. “Combining Forecasts: A Review and Annotated Bibliography.” *International Journal of Forecasting* 5 (4): 559–583.

Cowgill, Bo, Justin Wolfers, and Eric Zitzewitz. 2009. “Using Prediction Markets to Track Information Flows: Evidence from Google.” In *Auctions, Market Mechanisms and Their Applications: First International ICST Conference, AMMA 2009, Boston, MA, May 8–9, 2009, Revised Selected Papers*, edited by Sanmay Das, Michael Ostrovsky, David Pennock, and Boleslaw Szymanski. New York: Springer.

Fama, Eugene F. 1970. “Efficient Capital Markets: A Review of Theory and Empirical Work.” *Journal of Finance* 25 (2): 383–417.

Forsythe, Robert, Forrest Nelson, George R. Neumann, and Jack Wright. 1992. “Anatomy of an Experimental Political Stock Market.” *American Economic Review* 82 (5): 1142–1161.

Hanson, Robin, and Ryan Oprea. 2009. “A Manipulator Can Aid Prediction Market Accuracy.” *Economica* 76 (302): 304–314.

Hopman, Jay W. 2007. “Using Forecasting Markets to Manage Demand Risk. *Intel Technology Journal* 11 (2): 127–136.

King, Rachael. 2006. “Workers: Place Your Bets.” *Business Week*, August 3. http://www.businessweek.com/stories/2006-08-02/workers-place-your-bets.

Malone, Thomas W. 2004. “Bringing the Market Inside.” *Harvard Business Review* 82 (4): 106–114.

Nagar, Yiftach, and Thomas W. Malone. 2011. “Making Business Predictions by Combining Human and Machine Intelligence in Prediction Markets.” Paper presented at the Thirty Second International Conference on Information Systems (ICIS 2011), Shanghai, China, December 5.

Page, Scott E. 2007. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton, NJ: Princeton University Press.

Pennock, David M., Steve Lawrence, C. Lee Giles, and Finn Årup Nielsen. 2001. “The Real Power of Artificial Markets.” Letters, *Science* 291 (5506): 987–988.

Servan-Schreiber, Emile. 2012. “Prediction Markets: Trading Uncertainty for Collective Wisdom.” In *Collective Wisdom: Principles and Mechanisms*, edited by Hélène Landemore and Jon Elster, 21–37. New York: Cambridge University Press.

Spann, Martin, and Bernd Skiera. 2009. “Sports Forecasting: A Comparison of the Forecast Accuracy of Prediction Markets, Betting Odds and Tipsters.” *Journal of Forecasting* 28: 55–72.

Sunstein, Cass R. 2005. “Group Judgments: Statistical Means, Deliberation, and Information Markets.” *New York University Law Review* 80 (3): 962–1049.

———. 2006. *Infotopia: How Many Minds Produce Knowledge*. New York: Oxford University Press.

Surowiecki, James. 2004. *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. New York: Doubleday.

Thompson, Donald N. 2012. *Oracles: How Prediction Markets Turn Employees into Visionaries*. Boston, MA: Harvard Business Review Press.

Wolfers, Justin, and Eric Zitzewitz. 2004. “Prediction Markets.” *Journal of Economic Perspectives* 18 (2): 107–126.