Gambling With The Future:
Prediction Markets in the Corporate World

Prepared for: Dr. Michele Shauf, Instructor

Prepared by: Brian Cox

May 7, 2008

Contents

Executive Summary

In his book, *The Wisdom of Crowds,* James Surowiecki (2004) gives several examples of how groups of people who are asked questions that none of them can answer alone can answer if all of their answers are collected and a collective answer is provided. There is a group guessing the weight of an Ox or the number of jellybeans in a jar. Their group answer is usually better than the answer of the smartest of them.

This phenomena is partly what drives the stock market and since 1988, the electronic prediction markets. Prediction markets are markets where traders can bet on outcomes of certain events like Presidential elections, box office receipts and other world events. The surprising thing about these markets is how accurate they are.

Businesses have taken notice of the accuracy and utility of these markets and have started to experiment with markets in their companies. Google, Intel, Microsoft and other companies have all set up prediction markets within their companies to take advantage of the knowledge of their employees.

I have researched the various types of prediction markets and how they function. Then, I looked at how three businesses—Hewlett Packard, Yahoo, and Microsoft—used the markets. Hewlett Packard used the market to accurately predict financial events at their company; Yahoo used their market to determine new product development; and Microsoft used a market to predict a release date for new software.

After looking at these three companies and researching prediction markets, I am convinced that businesses are the next area where prediction markets will be used effectively. The use of prediction markets can give businesses an advantage over their competition by allowing them to make accurate predictions about earnings, product development and marketing strategies.

When the Space Shuttle *Challenger* exploded on January 28, 1986, the stock market reacted quickly by selling off shares of the four major contractors involved in building it. By the end of the day, the stocks had all gone back up slightly except for one company's—Morton Thiokol. Morton Thiokol had built the booster rockets for the shuttle and by the end of the day, its stocks had dropped 12%. It wasn't until six weeks later that the Presidential Commission on *Challenger* issued its report blaming the disaster on the O-ring seals made by Morton Thiokol. In their paper examining the stock market's reaction to the explosion, Maloney and Mulherin (1998) show that the stock market had predicted the cause of the explosion weeks before any extensive information was available.

How did the stock market predict the cause of the explosion so accurately? James Surowiecki (2004) calls it "the wisdom of crowds" in his book by the same name. He says, "under the right circumstances, groups are remarkably intelligent, and are often smarter than the smartest people in them." (p.XIII) This idea was first expressed by the economist and philosopher Friedrich A. Hayek (1945) in his essay, "The Use of Knowledge in Society." In the essay, he asserts that if one is given all of the information and all of the circumstances about a situation, for example the economy, then one can make logical and perfect decisions. However, we don't ever get all of the information. This presents a problem. The only way to get all of the information is to collect the scattered information in one place. Hayek sees the price of goods as a collection of this hidden knowledge, but the stock market functions in the same way—as a mechanism to collect dispersed information.

Now, with the development of the Internet, collecting information from diffuse sources has become much easier. Hidden knowledge from people around the world can be gathered into one place and that information can be used. Markets with mechanisms similar to the stock market can now be set up online to predict the outcome of almost any event from political races to movie box office receipts (Cherry, 2007). These markets are called Prediction Markets or Decision Markets, and when used carefully and set up properly, can provide accurate information about many topics.

Prediction markets have many untapped possibilities. Recently, the markets have been used to predict United States elections. In 2006, these markets correctly predicted every

Senate race in the United States, including the close ones. In the past few years, the markets have also been used to try to predict a wide variety of outcomes to things such as influenza outbreaks, world events and even weather events like the amount of snowfall in New York City. (Cherry, 2007)

One area of life that has been slow to use prediction markets, but is now experimenting with them is business. Companies with large numbers of well-informed employees or customers are ideal places to set up prediction markets and many companies, such as Hewlett Packard (Plott & Chen, 2002) and Google (Cowgill, Wolfers & Zitzewitz, 2008) are beginning to test them. If used properly, prediction markets can help companies can gain an advantage in today's business world. Markets can be used to forecast sales, develop new products or services, or predict product release dates.

Betting on elections

The first prediction markets began appearing in the 1850's as wagers on the results of presidential contests. These markets were set up in many large cities, but the largest was in New York City. They were first run out of shops and hotels, but eventually ended up being run by Wall Street. The *New York Times* and other city papers printed the results

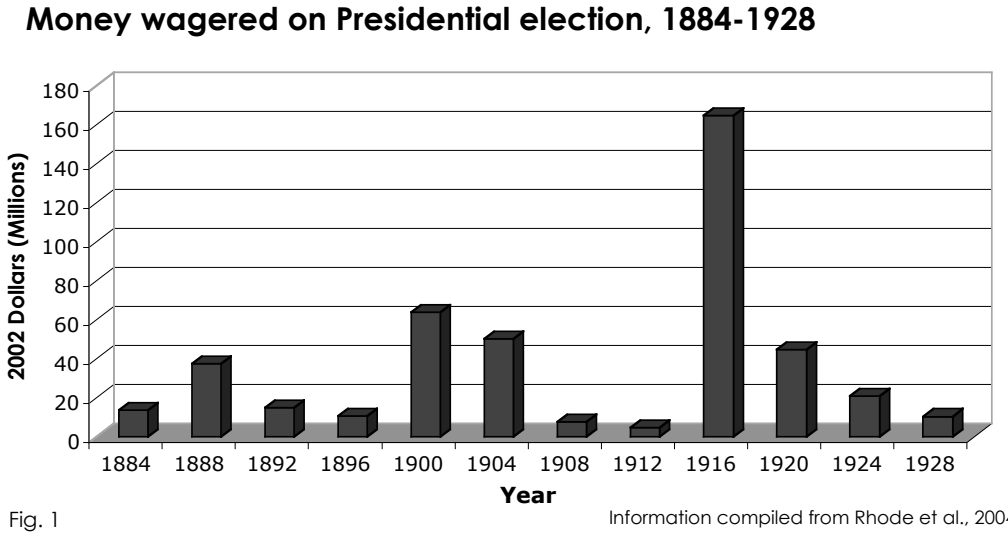

Fig. 1 Information compiled from Rhode et al., 2004

on a regular basis. Rhode and Strumpf (2004) studied these betting markets and found that they were popular, accurate and efficient. Between 1868 and 1940, millions of dollars were bet in these markets (see Fig.1). In 1916 election, over $165 million dollars

was bet on the Presidential contest. Despite the slower communications of the time, the betting market was wrong only once: in 1916 (Rhode et al.,2004).

Beginning in the 1920s and continuing through the 1950s, there was a great interest in group dynamics. Psychologists and sociologists set up numerous tests and experiments that dealt with how crowds functioned. Students were asked to guess the room temperature. Groups were asked to guess the number of jellybeans in a jar. Some groups were asked to guess the weight of cows at a fair. In all cases the group's guess was more accurate than most of the individual guesses (Surowiecki, 2004).

Friedrich Hayek, the economist and philosopher became interested in the nature of dispersed knowledge. He felt that a small group of economic planners, no matter how well informed would still be lacking information. If the knowledge that was dispersed through society could be collected and used, decisions could be made much more accurately. He believed that prices act as a mechanism that collects and reflects this dispersed knowledge in one place:

> "The whole acts as one market, not because any of its members survey the whole field, but because their limited individual fields of vision sufficiently overlap so that through many intermediaries the relevant information is communicated to all." (Hayek, 1945)

The next step toward prediction markets occurred in the late 1980s. George Mason University's Robin Hanson began to study prediction markets more seriously. He called them decision markets and developed what has since become a standard format for setting up a decision market. It is called Hanson's Market Maker and allows people to buy or sell contracts for whatever question is being asked in the market (Cherry, 2007).

At about the same time prediction markets for elections were restarted at the University of Iowa. In 1988, the University set up the first electronic prediction market, the Iowa Electronic Market, to allow traders to buy election shares. Since it's beginning, the Iowa Electronic Market has been accurate in its election predictions as well as in its other predictions. Today, the University continues to study the methodology and accuracy of different types of these prediction markets (Cherry, 2007).

With the success of the Iowa Electronic market, more commercial betting websites began to multiply on the Internet. Some of the more popular sites today are Intrade, based in Dublin and Tradesport. These markets let traders buy and sell shares in markets with subjects that run the gamut from sports to earnings of publicly traded companies to elections. Like their predecessors, these markets are also very accurate. In 2006, Tradesport's markets predicted all 33 of the US Senate races accurately. No public opinion polls were as perfect (Cherry, 2007).

The theory behind the market

Prediction markets work much like the stock market or placing a bet. In the case of prediction markets, what is being invested in or bet on is not the earnings or value of a company, but rather the outcome of certain events. Who will win the Oscar for Best Actress? When will a public official resign? Or, who will win the next election? The contract will pay off based on the outcome of the question.

The benefit of this kind of market is that it collects dispersed information from its various investors and reflects this knowledge in the form of contract prices—just as the stock market reflects information about a company in its stock prices. The advantage in using markets to make a prediction is that it can accommodate changing situations using continuously changing information. It collects a wide variety of information from a wide variety of sources and reduces them to guesses, rendering them unbiased and difficult to manipulate (Wolfers & Zitzewitz, 2004). Someone with an agenda is only one of many sources of knowledge in the market.

However, there are necessary parameters that must be followed. Surowiecki (2004) points out four general conditions that must be met to produce "wise crowds" rather than just mobs (p.10):

1. Diversity of opinion
2. Independence
3. Decentralization
4. Aggregation

There must be a diversity of opinion to make the market accurate. A group of people with the same job or same knowledge base would generally come to the same conclusion and the market's prediction would be useless. A diversity of people, with differing opinions will balance the market. In a prediction market, even the wrong answers contribute to the accuracy of the market by encouraging those who think they know the right answer to participate and by pulling the data toward the median.

As Surowiecki (2004) explains, traders must be independent from each other. For the market to be accurate, it is necessary for the users to buy contracts based on their own knowledge and opinions. If others influence them or tell them how to invest, the accuracy of the market will suffer.

Also, the crowd must be decentralized. Too much of one kind of knowledge will skew the market. The market needs to draw on a variety of knowledge and not a group of people with the same knowledge.

Aggregation means that there has to be some way to collect this knowledge. The stock market or a prediction market can collect this dispersed knowledge, concentrate it and reflect its conclusions.

After meeting those parameters, one has the potential traders for a successful prediction market. However, in order to actually set up a market, several other conditions must be met. Hanson (1999) lists six steps to make sure the prediction market is accurate.

First, the claim must be worded clearly. (Hanson, 1999, p. 17) For example, if one wanted to predict the month for a hurricane to hit the United States, she would set up a series of specific choices possibly starting with July and ending in September. This would allow traders to buy contracts in whichever month they felt was most likely to produce a hurricane.

Secondly, there must be some way to determine the outcome. Hanson (1999, p.17) suggests picking a trusted third party to declare the results. However, the third party may be any trusted source. Elections and events such as the hurricane mentioned above can be easily verified.

There must be an incentive to participate. It may be money or some other form of reward for getting the right answer. The return on the trader's time and investment must be sufficient enough to encourage knowledgeable involvement. (Hanson, 1999, p.17) Because of gambling laws in the United States, some markets have gone to play money exchanges. There are differences in outcomes between real money and play money exchanges, but there isn't enough evidence at this time to make reach any conclusions about whether the type of incentive is affecting the markets (Wolfers et al.,2004). Also, there must be a financial institution to facilitate these exchanges.

Hanson (1999, p. 17) also recommends using markets that are interactive. He suggests allowing participants to trade contracts or assets with each other. By comparing the market prices and the amount of trades, one can determine the expectations of the traders.

If the market is about a topic that is hot at the moment, the topic itself might keep people interested. However, in most cases, the final step would be to subsidize the market. This means deciding how much money or other assets will be involved in the market to encourage participation (Hanson, 1999, p.17). Sometimes it may be necessary to add some money or a contract to the market to increase incentive for activity (Wolfers et al., 2004). The goal is to keep a wide variety of people trading on the market.

When it won't work

In his book, *Infotopia*, author Cass Sunstein (2006) looks at some of the possible errors or downsides of prediction markets that would cause them to be inaccurate. If the investors in a market are biased, then the market will be incorrect. Traders will all bet on their favorite outcome. For example, if a market asks a baseball question and the trader base is in New York, most of the trades will tend to bet for the New York Yankees regardless of whether the team is actually the favorite to win or not (Sunstein, 2006, p. 138).

Another factor that will skew prediction markets is too little information. When President Bush was nominating a Supreme Court Justice to replace Justice Sandra Day O'Conner, the markets were very wrong. The eventual nominee John Roberts was not a

favorite, or even much in the running until late in the evening "about the same time President Bush was contacting Senate leaders to inform them of his choice." (Sunstein, 2006, p.133) The market simply did not have enough information to be accurate.

Sunstein (2006, 137) also looks at the possibility of manipulating the market. Is it possible for someone with a lot of money to artificially influence a market? It is possible, however, the market tends to take care of itself. Those with more knowledge will take advantage of the person doing the manipulating and the market will quickly return to accurate levels while the smarter investors will make some money.

There is also the question of herd mentality. Fads and hyped markets may trigger a "bubble" if the trading base is swayed by unreliable, but highly publicized information. People have a tendency to follow and if their fellow traders are all buying into one contract there is a likelihood they will follow erroneous information as well (Sunstein, 2006, 140).

One final drawback to be wary of is the type of question asked in the market. Prediction markets are not good at evaluating what Sunstein calls "judgements of value." (2006, p.143) Moral or ethical questions that have no definitive answer or outcome are not objective questions. Prediction markets must have some verifiable answer in order to be effective. A market question about whether Congress should pass a bill limiting Carbon Dioxide emissions wouldn't give any accurate predictions, but a question asking if they will pass a bill would likely be very accurate.

Market Methods

Several market models exist to collect information for prediction markets and they all can be used to make accurate predictions. In all of them, the amount of trades or contracts bought at a particular price will show what the market's predictions are. As contracts are bought—or bets placed—the market will rise or fall giving a "probability, mean or median, respectively" that will show the expected outcome of the market (Wolfers et al., 2004). Wolfers and Zitzewitz (2004) list three main methods in use that either separately or jointly can be used to make predictions. These are "winner-take-all" contracts, "index " contracts, and spread betting.

In a “winner-take-all” market, a trader buys a contract for some amount and the contracts pay off if a certain event occurs. As more contracts are bought, the price for the contract goes up reflecting the likelihood of the event occurring. Several choices of contracts would also be given. For example, in a “winner-take-all” market for the 2004 Presidential election, the market would offer contracts for George W. Bush, John Kerry and Ralph Nader. If a trader bought a contract for 23 cents betting that George W. Bush would win, this would mean that at the time of the sale Bush had only a 23 percent chance of winning and the trader would be paid $1 on the contract if Bush won. The price of Bush’s contracts would go up as more of his contracts were purchased and the market price would show the predicted chances of Bush’s election.

In an “index” market, contracts would be sold based on a percentage of the total. Traders buy shares in the market and the winners are paid based on the percentage of the outcome. In the Presidential election example mentioned above, a trader would buy shares in his favorite candidate. As shares were bought, the market would rise to reflect the confidence of the investors in their chosen candidate and would show the likelihood of each candidate’s winning. However, the payoff would be a percentage of the vote total the winning candidate gets. Since Bush won, the payoff to Bush investors would be based on the percentage of the vote total he received in the election.

The third kind of market is the spread. This is the same as betting on a sporting event. Traders would purchase contracts, basically betting that the prediction will cover the spread. In the 2004 election example, if the trader bought a contract betting that Kerry would get at least 49.5% of the vote, the trader would be paid only if Kerry covered the spread or exceeded 49.5% of the vote total.

The uses of these prediction markets are still being explored. Although, the examples used by Wolfers and Zitzewitz (2004) are political markets, any prediction market can use some variation or combination of these three types of markets. When analyzed these markets can give a great deal of insight into an almost limitless amount of topics. Effects of events on each other, changes in public opinion, and how information is processed by the public are all being studied through the use of prediction markets. One area where prediction markets have recently begun to be studied is in business.

Prediction markets in the business world

Until recently, prediction markets were used mainly for public events, such as politics, sports and economics. They are now, however, starting to be tested by several large companies to tap into the knowledge of their employees. Executives and researchers see that companies with lots of well-informed employees are ideal for prediction markets and increasingly, more prediction market studies are being done at companies like General Electric, Best Buy and Intel (Cherry, 2007) to determine the best uses of these markets for their companies.

Businesses are seeing benefits in using prediction markets. They allow employers to access the untapped resources of their employees by giving a voice to employees that otherwise would not speak up. They are flexible and can reflect change in areas where change is rapid and quick adaptation is necessary. Markets also allow companies to detect trouble early. If something is not going as planned, the information is sure to show up in the prediction market (King, 2006)

Prediction markets also save time and money. Rather than hire experts to make the decisions and test the products, companies can rely on a market set up internally to get the answers they need. Monthly meetings of hired experts are no longer needed.

Finally, the information gathered by prediction markets is accurate. At Hewlett Packard, Kay-Yut Chen who has been researching the use of prediction markets says, “In six out of eight cases, it was more accurate than corporate forecasts.” (King, 2006) This accuracy is useful for companies to be able to predict product release dates, determine which products to develop, and forecast sales figures.

Companies like Hewlett Packard, Yahoo, and Microsoft have all tried to use prediction markets in various ways to improve their businesses. By looking at these three companies, we can see examples of how prediction markets may be used in the future to make businesses more efficient and productive.

First tests with Hewlett Packard

Hewlett Packard was one of the first companies to test prediction markets. Starting in 1996, California Institute of Technology and Hewlett Packard began a series of

prediction markets or what they called an Information Aggregation Mechanism. They used Hewlett Packard employees to test the utility of prediction markets in a business environment. Charles R. Plott of Cal. Tech and Kay-Yut Chen of HP designed a series of markets to test 12 events at HP. These events ranged from the amount of profit sharing to be announced by upper management to monthly sales. They then asked a group of employees to invest in a market and try to predict the outcome of these events. A comparison between the markets' predictions and the Hewlett Packard's official forecasts would then show the accuracy of the market (Plott et al., 2002).

The results (Fig. 2) were encouraging. They reported, "Not only did the IAM market predictions consistently beat the official HP forecasts; the outcomes predicted are consistent with the probabilistic predictions of the IAM." (Plott et al., 2002, p.17)

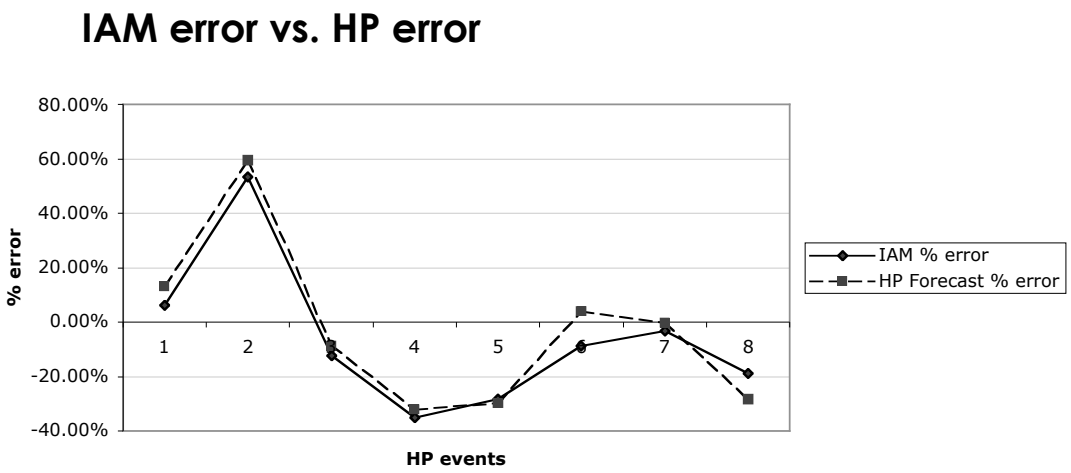

Fig. 2 Data compiled from Plott et al., 2006, p.20

Other conclusions that Plott and Chen(2006) arrived at were that the markets were flexible—they can be used to collect "any type of information by different people," their system was scalable and could be used for any number of people, and incentives keep people from hiding or concealing information (p.17). The Hewlett Packard test led the way in using prediction markets to help forecast earnings and monthly sales.

Yahoo employees get their say

In 2006, Yahoo had their annual hack day, which consists of teams of Yahoo employees giving their advice on products they would like to see Yahoo bring to market. In the past a small group of executives would decide which products to develop. That year, however, Yahoo research scientist David Pennock set up a market to allow the employees to bet on

which products they thought would do well with consumers. Employees were given money and allowed to invest in which products they thought were best. Prabhakar Raghavan, head of Yahoo research explained, “As a product organization, you hope to remove the burden of predicting from one product manager.” (King, 2006)

Yahoo used the predictions to determine which products to make into prototypes. By not having to hire consultants to market-test the ideas, Yahoo saved the company money. At the same time, using prediction markets took the pressure off the management and gave the employees, many who may use the product, some input into what would go to market.

Microsoft sets the date

Todd Proebsting at Microsoft had been asked to use prediction markets to help determine the release date for software the company was working on for internal use. The initial release date was set for November 2004, and Proebsting, a researcher at Microsoft was asked to create a market to predict when the software would actually be completed.

Proebsting set up a market with six possible bets as to the outcome. He asked a group of employees from across the company to bet on the release date and gave them each $50 to bet with. All six possibilities started at 16 2/3 cents on the dollar. If they bought a contract for 16 2/3 cents and were correct about the release date, they would be paid a dollar for each contract they had purchased. The November release date immediately fell to zero. As features were taken off and added to the software in the development stage, the prediction of the release date fluctuated, but the market settled on February. As Proebsting says, “In the end, the product shipped in February, which is what the market predicted.” (Cherry, 2007)

Although the software that was being developed was for internal use, Microsoft could use the markets in the future to accurately predict other release dates. This ability to accurately predict deadlines would allow Microsoft not only to plan ahead for release, but also to flag any problems that may arise when they see the product is behind schedule. Prediction markets can add efficiency and accuracy to Microsoft’s organization.

Why not use it now?

While many companies are starting to experiment with using prediction markets in their businesses to help them make decisions, there still is an overall resistance to using them. Several factors contribute to this hesitation. First, most executives are reluctant to give up the decision-making authority they have. Executives and management have worked hard to get to where they are and they don't want to give up their authority. This top-down structure is typical in corporate America and does have some advantages. Someone has to set the schedule and make the decisions. For example, by setting deadlines, whether the deadlines are realistic or not, managers can motivate employees. (Cherry, 2007)

The subversive nature of using these markets in any business is that they have the possibility of undermining executive authority and damaging morale. A market letting the employees express opinions can have the effect of postponing deadlines and showing the company in a bad light. If a market shows a lack of confidence in a product or an executive, the morale or effectiveness of a leader may be damaged. (Cherry, 2007)

It is also a hard concept to grasp. Companies, after all, do hire executives and experts for a reason. It seems hard to believe that a group of uninformed employees can predict outcomes better than the professional forecasters.

Conclusion

While prediction markets have been around in some form since the 1860s, they didn't start becoming easy to manage and efficient until the growth of the Internet. With the prevalence of the Internet, it is easier to gather dispersed information into one place. This information can be compiled and used to predict events. Websites like Tradesport or the Iowa Electronic Markets have been used for years to predict the outcome of elections, results of the Oscar Awards, and high profile world events.

The accuracy of these markets has recently brought prediction markets to the attention of business. Businesses such as Hewlett Packard, Yahoo, Microsoft, Best Buy and Google have all recently started experimenting with prediction markets to help forecast

earnings, develop and test products and predict release dates of new products. This allows companies to save money and time and to plan their marketing strategies better.

The use of prediction markets to strengthen companies and gain advantages in the marketplace is still new to business, and there is an initial resistance to using them. However, once businesses see the advantages of implementing prediction markets in their companies, it is only a matter of time before more start tapping into one of their most valuable sources of knowledge: their employees.

Bibliography

Berg, J. & Rietz, T. (2003). Prediction markets as decision support systems. *Information Systems Frontiers,* 5(1), 79–93.

Cherry, S. (2007, Sept.). Bet on it [Electronic version]. *IEEE Spectrum*. Retrieved April 13, 2008, from http://www.spectrum.ieee.org/sep07/5488/2

Cowgill, B., Wolfers, J., & Zitzewitz, E. (2008, January 6). *Using prediction markets to track information flows: evidence from Google*. Retrieved April 21, 2008, from http://bocowgill.com/GooglePredictionMarketPaper.pdf

Hanson, R. (1999). Decision markets. *IEEE Intelligent Systems,* 14(3), 16-19.

Hayek, F.A. (1945) The use of knowledge in society. *American Economic Review,* 35(4), 519-30. Retrieved April 14, 2008, from http://www.econlib.org/Library/Essays/hykKnw1.html

King, R. (2006, August 3). Workers, place your bets [Electronic version]. *Businessweek.com*. Retrieved April 14, 2008, from http://www.businessweek.com/technology/content/aug2006/tc20060803_012437.htm

Maloney, M., & Mulherin, H. (1998, December 7). *The stock price reaction to the Challenger Crash*. Retrieved April 30, 2008, from http://ssrn.com/abstract=141971

Plott, C., & Chen, K. (2002, March). *Information aggregation mechanisms: concept, design and implementation for a sales forecasting problem* (Social Science Working Paper No.1131). California Institute of Technology. Retrieved March 18, 2008, from http://www.hpl.hp.com/personal/KayYut_Chen/paper/ms020408.pdf

Rhode, P., & Strumpf, K. (2004,). *Historical Presidential betting markets*. Retrieved

April 28, 2008, from http://www.unc.edu/~cigar/papers/BettingPaper_final (JEP_Resubmit).pdf

Sunstein, C. (2006). *Infotopia*. New York, Oxford University Press.

Surowiecki, J. (2004). *The Wisdom of Crowds*. New York, Anchor Books.

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