Fall 2007 | Vol.50, No.1 | REPRINT SERIES

California Management Review

New Product Blockbusters: The Magic and Science of Prediction Markets

Teck-Hua Ho
Kay-Yut Chen

University of California
Berkeley
Haas School of Business

New Product Blockbusters:
The Magic and Science of Prediction Markets

Teck-Hua Ho
Kay-Yut Chen

New products are engines of growth for many firms. Successful new products are often sources of long-term competitive advantage. Indeed, some firms' survival depends on their ability to manage new product launches. In 2004, the average percent of firm's total sales attributable to new products developed within the last three years is about 33%.[1] An extensive longitudinal study suggests that new product introductions increase a firm's long-term financial performance and market value.[2] Thus, the importance of new product development cannot be overemphasized.

A significant portion of a firm's financial resources and managerial attention are spent to develop extensions to existing products and to create new products. More and more products are launched every year.[3] As much as 8% of a firm's sales are spent in the undertaking of these endeavors.[4] A survey by IEEE Spectrum shows that the top 100 R&D spenders, taken as a whole, spent $254 billion in 2004.[5] In fact, Forrester Research reports that more than 90% of CEOs feel that new product innovation is very or extremely important to growth.[6] The same research study indicates that more than 50% of senior executives are dissatisfied with the returns on new product innovation. There are significant differences across firms in new product performance. The PDMA study finds that the best performers in new product development achieve twice-as-high sales than the worse performers.[7] Hence, it is not only possible, but also critical to maximize the return of new products.

Two powerful ways of improving the return of new products are to invest in only the most promising new product ideas and to improve supply planning before products launch. Most firms fail to correctly pick new product winners—only one out of every five new product launches is successful.[8] Also, firms frequently are unable to capitalize on the successes of a new product blockbuster

because of poor demand forecasts. For example, Nintendo launched the new console Wii in November 2006 with huge product shortages and many frustrated fans. Similarly, Apple Computer had to push back their international launch date of iPod Mini because of supply constraints. Hence, it is important to discover and manage new product blockbusters.

There are two common approaches to predicting new product demand before a launch. The first is to survey target customers about their purchase intentions. Products that have high purchase intention are selected and launched. This approach has been used widely in the screening of new product ideas. It has two problems. First, surveys do not motivate customers to reveal their true purchase intentions. As a consequence, the collected data can be quite noisy. Hence, the link between stated purchase intention and ultimate purchase behavior is weak, and the associated demand forecast is inaccurate. Second, most people are imitators and rely on others to learn about the potential benefits of new products. Hence, when surveyed, and without learning from early adopters, they give biased response of their purchase intentions.

The second approach to predicting new product demand involves the pooling of experts' opinions. It is used widely in the fashion industry where new product demand is highly uncertain. Under this approach, a group of experts are asked to state their opinions and the average opinion is used to gauge the success of the new product. This approach has three problems. First, the pool of available experts is usually small. Hence there is considerable variance in the forecast. Second, opinions are typically weighted equally independent of expert's knowledge. Ideally, individuals with better knowledge should be assigned more weight. Third, experts' opinions may not be independent of each other because they may rely on same information source.

In this article, we describe a novel approach to screening new product ideas and predicting their demand. Under this approach, individuals are motivated financially to participate in an organized market with well-defined rules. The goal of a prediction market is to aggregate relevant information from multiple and diverse people. After the new product is launched, the market rewards participants based on their forecast accuracy.

The prediction market addresses the potential problems of the survey-based and expert-based approaches. First, participants are compensated for accuracy in forecasting. Second, everyone can learn from others about the potential of a new product idea through the markets. Such learning allows individuals to update their beliefs and develop a better forecast. Third, the price discovery process naturally weighs accurate information more heavily. The same price discovery process also removes redundant and dependent information sources appropriately. Fourth, prediction markets can accommodate many participants at a minimal incremental cost. Once the system is built, it can be used on a continuous basis.

Teck Ho is the William Halford Jr. Family Professor of Marketing at the Haas School of Business at UC Berkeley.

Kay-Yut Chen is a principal scientist at Hewlett-Packard Laboratories.

The Basics of Prediction Markets

Intellectual History

The idea of soliciting inputs from diverse individuals to improve decision making dates back to the dawn of civilization. The Lord of Menchang, in the Period of Warring States in China (around 300 B.C.), housed three thousand guests in order to tap into advice and expertise from a diverse group. Over two thousand years later, in the fall of 1906, at the annual West of English Fat Stock and Poultry Exhibition, 800 people entered into a contest to guess the weight of a fat ox. The group consisted of a few experts and many laymen. To the surprise of everyone, the average guess (1197) was phenomenally close to the actual weight (1198).[9] Many other similar and amazing examples were documented showing that large groups of individuals consistently outperform experts. These examples share two common traits. First, the number of participants is large. Second, participants come from diverse backgrounds and have independent sources of information.

In the modern world, companies have diverse employees and hence possess the promise of tapping into this power. However, this potential is seldom realized. The most common way of gathering input is to conduct a meeting. This method is plagued by several problems. First, members in the meeting may not have incentives to provide unbiased information. Worse yet, they often have incentives to provide biased input. Second, members often yield to their superiors because of a hierarchical power structure. Third, there is no systematic way to assign relative importance to each input. As a result, whoever argues most eloquently usually has his or her input weighted significantly more. However, a person's ability to communicate may not have any direct bearing on whether they have relevant information.

Economists have long wrestled with this information aggregation problem. In 1948, Edward Chamberlin conducted the first economic experiment to determine whether the market can aggregate demand and supply information. Subjects were provided monetary incentives to buy and sell a fictitious item. Half the subjects were sellers who had different costs of production, and the remaining were buyers who had different values for the item. Sellers were paid in real money based on profit (*price* $\times$ *quantity sold* $-$ *cost*) and buyers were compensated based on net surplus (*values* $-$ *price* $\times$ *quantity bought*). The values and costs were designed to reflect both a linear demand and supply function. No one subject, however, was aware of the entire demand and supply functions and hence the predicted market price. The subjects were free to negotiate on a one-to-one basis in a decentralized fashion. If this decentralized market were able to aggregate demand and supply information, the price would be at the intersection of the aggregated demand and supply functions. Despite monetary incentives, the market failed to aggregate information and yielded the predicted price.

Vernon Smith was a subject in Chamberlin's experiment. Smith recognized that market rules can have a dramatic effect on a market's ability to aggregate information. In Chamberlin's experiment, the decentralized nature of the

market did not allow subjects to learn from each other. Consequently, in his 1962 experiment, Smith retained most of Chamberlin's design except for the use of a different trading rule.[10] In Smith's design, offers made by buyers and sellers as well as the transaction prices were posted publicly. This change, while simple conceptually, made a powerful impact. Subjects could now learn from the market and adjust their behavior accordingly. As a result, this centralized market successfully aggregated demand and supply information to yield the predicted price. This seminal work laid the scientific foundation for subsequent research on the design of markets. In 2002, Smith won the Nobel Prize in economics for the study of alternative market mechanisms using laboratory experiments.

Chamberlin's and Smith's experiments illustrate several important market design principles. First, incentives drive subject behavior. Subjects trade to make money but their trades reveal information to the market. Second, a common metric (i.e., transaction price) is necessary to capture this market information. The most updated information is always reflected in the current market price. Third, a market should be transparent to encourage learning. This is typically done by making all market activities public.

Building on these principles, Charles Plott and Shyam Sunder created the first prototype of the modern prediction market in 1982 and further enhanced it in 1988.[11] Instead of using the market to aggregate demand and supply information, their market was designed to disseminate and aggregate individual diverse information about the value of a stock. The market prices captured that aggregate information and could be used to forecast the value of a stock. The basic design is as follows. Subjects were asked to bet on three stocks. Only one of the three would pay cash at the end of the experimental session. Subjects were provided with partial information about the winner. For example, some might be told that the winner was one of the two stocks (hence the remaining one stock could never be the winner). Subjects were also given some initial allocation of these stocks and seed money. They were then asked to trade these stocks in a market. Subjects could post public offers to buy and sell any stock in real time. The prices of the three stocks captured all the aggregate information of the subjects. The market worked well and was able to predict the winning stock reliably. This finding established the potential of the innovative use of markets to predict future events. Consequently, their design became the gold standard for future prediction markets.

The first field application that received widespread recognition is the Iowa Electronic Market (IEM), which is still running today. The IEM is not-for-profit and operated by the University of Iowa Tippie College of Business. The IEM has been focusing on predicting the outcomes of political events, such as presidential elections. Its participants were drawn from the general public and were allowed to use their own money in the market. However, a speculator can only bet up to $500 in the IEM. The following is an example of how it works. The 2008 U.S. Presidential Election prediction market has three outcomes: Democratic Party nominee, Republican Party nominee, and an independent candidate. Spectators can bet on who will become the president. The winning stock pays one dollar

while the losing ones pay nothing. When a spectator enters a market, he or she can exchange one dollar with one share in each of the outcomes. Since only one outcome will eventually pay, the market maker does not make any profits. IEM predictions are impressive. *Business Week* reported, "It [IEM] predicted the vote totals of the past two Presidential elections within two-tenths of a percentage point, outperforming national polls."[12] This success is the beginning of a widespread application of prediction markets. For example, IEM has recently used prediction markets to forecast influenza activities.

A Standard Prediction Market Setup

Following is our step-by-step approach to setting up a prediction market.

Determine the Schedule of Market Sessions

It is important to schedule the market in a way that its generated forecasts can influence resource allocation decisions. Since the major resource decisions may occur at multiple points in time (e.g., which product idea to invest, how large the plant should be), it is important to start the market before the first major decision. The market is left open all the way until product launch and is frozen on the day of the product launch.

Recruit Participants and Determine the Budget

New product development spans multiple business functions and hence it is important to receive inputs from all involved parties. Thus, participants should be recruited from different parts of the organization. For the market to function properly, the number of participants N must be sufficiently large. We recommend at least 50 participants. A market that has fewer participants is likely to encounter a liquidity problem.

Active participation from all is necessary for the market to function well. Hence each participant must be motivated sufficiently with financial incentives. We recommend an average compensation (C) for a participant of at least \$500. Hence a modest budget of $50 \times \$500 = \$25{,}000$ is enough to make this work.

In general, the total budget B is the product of number of participants and the average compensation. That is $B = N \times C$.

Determine Number of Competing Ideas

A firm may start with many new product ideas. However, prediction markets work best when this initial list is narrowed down to a manageable set of 2 to 4 of the most promising ideas. Each new product idea has its own prediction market so that we will have as many prediction markets as new product ideas. Participants are asked to trade in all markets and the product idea that wins the horserace (in terms of expected sales and profits) will be launched.

Define Demand Outcome Scenarios

A prediction market has a pre-determined number of demand outcome scenarios. A demand scenario is a range of sale volume and the actual eventual sales volume must fall into one of the scenarios. We typically divide the demand outcome into ten scenarios (D_1, D_2, . . . , D_{10} in the order of increasing demand). This design provides fine enough resolution and makes it simple enough to run. To determine the range of each demand scenario, we first solicit anonymously from each relevant employee his or her initial estimates for the most optimistic and pessimistic demand forecast. There are two numbers of particular interest. Let the highest of all optimistic estimates be H and the lowest of all pessimistic estimates be L. We define D_1 to cover the range below L and D_{10} to cover the range above H. The intermediate sales volume (D_2–D_9) is evenly spaced between L and H. The width of the interval is $w = (H - L)/8$. Hence, the demand scenarios are depicted as follows:

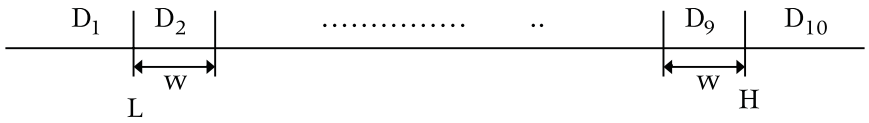

Design Market Parameters

Assume a total budget of B dollars, N participants, and M product ideas (prediction markets). It is a customary practice to stop participants from using their own money in trading. Therefore, they are typically provided with initial money and shares. To facilitate trading, the total budget is often divided evenly into cash and shares. Hence, the total cash given out is $B/2$. We divide the total cash evenly so each participant will receive $B/(2 \cdot N)$ in cash initially. Only shares in the correct demand scenario for the winning product idea pay and the worth of each share is set to $1 dollar. Note that shares in the incorrect demand scenarios for the winner product idea pay nothing. Also, shares in all demand scenarios of unselected new product ideas pay zero. As a result, total number of shares in each demand scenario is $B/2$. Consequently, each participant is provided with $B/(2 \cdot N)$ shares in each demand scenario in each prediction market.

Let us illustrate the above with a concrete example. Assume we have a budget of $25,000 with 50 participants. Here, each participant is provided with $250 cash and 250 shares in each of the demand scenarios for each of the prediction markets. Note that the total number of shares in each demand scenario is 12,500. Note that each participant cannot corner the market because they cannot use their own money to trade.

Participants can buy and sell shares in any demand scenario in a manner similar to trading in the stock market. In this respect, we have M stock markets and each has 10 companies with symbols (D_1, . . . , D_{10}). Individuals can post their buy and sell offers for a stock (e.g., demand scenario) and a trade occurs when a posted offer is accepted. All information is publicly posted and visible to all participants. Participants can trade any time after the market is open and the market will remain open until the new product is launched.

Discover the Winner and Generate Demand Forecasts

Each prediction market generates two kinds of demand forecasts. The market provides information about the most likely demand scenario and the chance of each demand scenario occurring. At any point in time, the probability of demand scenario $i\,(\theta_i)$ occurring is determined by dividing the price of demand scenario $i(P_i)$ by the sum of prices of all demand scenarios as follows:

$$\theta_i = \frac{P_i}{P_1 + P_2 + \ldots + P_9 + P_{10}}$$

The most likely demand scenario is the demand scenario that has the highest chance of occurring. This information can then be used to compute financial metrics such as expected profits and return-on-investment.

To discover the winner, one must first decide on appropriate financial metrics. These metrics are then calculated for each new product idea based on generated demand forecasts from the corresponding prediction market. A selection is made based on these financial metrics.

Once a product idea is selected, only that particular product prediction market continues to operate. The demand forecasts generated from this market can be useful before major supply decisions are made. For example, one can use the expected demand distribution to size the capacity of a plant.

Determine Participants' Earnings

Shares in unselected product ideas receive zero payout. The payoff of shares in the launched product is determines as follows. After the product is launched and the sales volume is observed, each share of the demand scenario that contains the actual sales volume is worth one dollar. All shares of other demand scenarios pay nothing. For example, consider "Hewlett" who started with 250 shares of every demand scenario and a cash pot of $250. At the end of trading session, Hewlett owned 500 shares of D_5 and 500 shares of D_8 in selected new product idea and had a balance in cash of $80. If the actual sales are within D_8, then Hewlett will receive a payment of $500 for his shares of D_8, giving him a total earning of $580.

Encourage Active Trading

The prediction market will only function well if participants actively trade. Here are some practical tips for ensuring participation.

- Senior management should encourage active participation. They should emphasize the importance of such activity because it provides useful inputs to the planning process.
- We can impose a rule so that only participants who transact above a certain minimum level (e.g., an average of one transaction per day) can receive their earnings after the market is closed.

If the number of participants is a multiple of ten, we can provide each participant with only shares of one demand scenario to encourage trading. For

example, if there are 50 participants and a budget of $25,000, we can assign the first participant 12,500 shares of D_1, the second 12,500 shares of D_2, and so on. Note that this pattern repeats so that the first, eleventh, twenty-first, thirty-first, and forty-first receive 12,500 shares of D_1.

Scientific Foundation

There is a scientific foundation for why prediction markets work. Prediction markets work on five principles: incentive, indicator, improvement, independence, and crowd.

Incentive

Prediction markets must provide strong incentives for good use of market information. They should neither reward status nor dominance, which are common in organizations. The principle of rewarding solely based on information use creates an environment that is conducive for opinions to aggregate and emerge.

Incentives also provide a natural way to weigh opinions. Unlike meetings, prediction markets do not weight opinions based on influence. Unlike market research, prediction markets motivate participants to reveal their true beliefs. Unlike polling, prediction markets treat individuals differently based on their knowledge. As a result, participants who are more confident will place a larger bet and hence their information will be given more weight in the aggregation process.

Indicator

Prediction markets employ a clear information indicator. In particular, price is used to convey aggregation information to all participants. This solves two fundamental problems of information pooling and dissemination (for learning). First, individuals have different mental models of demand. The use of price forces participants to express their thinking in a precise and common metric, which is paramount to the market's ability to merge information. For example, there is no scientific method to combine a news story about a trend in the market and an expert opinion about a new product into one single demand forecast. Prediction markets fill this gap.

Secondly, different participants have different levels of accuracy of information. The use of price allows the market to give more weight to more informed individuals. These individuals are more likely to trade and hence influence the market. This is so because the individuals can increase their profit by trading with their personal information. Thus, they have huge incentives to quickly "pump" information into the market before others do. As a consequence, the market forecast of demand becomes more accurate.

In fact, the existing stock markets are live and long-standing examples of prediction markets. They provide forecasts of firms' values. In these markets,

prices summarize all investors' mental models about firms' worth. Informed traders move the markets. For example, an individual who has a piece of positive news about a firm can benefit by buying shares and hence drive up its price. The updated price now captures this positive news.

Improvement

Prediction markets encourage individuals to improve their knowledge. As noted, prices capture the latest information about the demand. The price discovery process, in effect, allows the uninformed individuals to learn from the informed ones. Hence, all participants will become knowledgeable of the demand through trading.

This adaptation process allows individuals to piggyback their personal learning on others' information. This is analogous to the common scientific discovery process of riding on the shoulders of giants. Put differently, this process equalizes the information bases of all individuals, because prices are common knowledge.

Consequently, the markets become smarter through this continuous process of learning by all traders.

Independence

Prediction markets benefit from independent information sources. For example, Hewlett thinks that demand for a printer is either low or medium. "Packard" thinks that demand for the same printer is either medium or high. Prediction markets can pool Hewlett and Packard's information together to yield a demand forecast of medium.

Moreover, pooling can even be useful when there is an expert. The ancient Chinese saying, "Three humble shoemakers [laymen], brainstorming together, can defeat one great statesman [expert]," reflects the importance of pooling independent thinking. Now we understand that this insight is firmly rooted in the mathematical principle of Bayes' Law, a well-known law in statistics for aggregating information. The following simple example demonstrates this idea.

Hewlett concludes from a marketing research study that the demand for the new audio oscillator model 200A is either low or medium. Packard, after talking to distributors, deduces that the demand for the same product is either medium or high. Pooling both sources of independent information will yield the correct answer that the demand is medium.

In most organizations, there are rigid hierarchies. These layers of structures impede free flow of independent information. For example, a manager may filter off opinions incompatible with their own thinking, thereby stopping the organization from employing all useful information. Prediction markets are designed to remove these barriers.

Crowd

Prediction markets work best in a large crowd. The surprising fact that groups consistently outperform individuals is well documented.[13] This observation is grounded on the statistical principle of the Law of Large Number. This principle further states that a large group of laymen can even beat a small number of experts.

The following mathematical logic illustrates this insight. Let the margins of error be x and y for an expert and a layman respectively. Naturally $x < y$. If n laymen work together, and they have independent sources of knowledge, then the Law of Large Number states that the margin of error of the group is $y/\sqrt{n}$. So, if a layman makes four times more error, then the Law of Large Number states that you only need 16 laymen to perform as well as the expert (since $y/\sqrt{n} = 4x/\sqrt{16} = x$). Hence, a group of 64 laymen can beat the expert by a factor of two.

Summary

By this logic, there is no surprise that the Iowa Electronic Markets consistently predict election outcomes better than experienced political observers.

In sum, prediction markets motivate people to share information clearly and freely through the price discovery process. They encourage participants to learn from each other, and they promote pooling information from a large group of diverse individuals.

Application Sweet Spots

To show how prediction markets are used in practice, we have chosen examples from three industries where forecasting is hard because history is often not a good predictor of the future.

Movie Industry

Movie releases are new product launches. A movie is like a fashion product for three reasons. First, the life of a movie is short, lasting typically from a few weeks to several months. Second, the demand for a movie is highly uncertain. For example, "My Big Fat Greek Wedding" had a meager budget of $5 million with little marketing. To the surprise of the industry, it grossed about $240 million in the United States. On the flip side, "Basic Instinct 2," a sequel with a well-known star, cost $70 million and grossed only $6 million. Third, the early signs of demand are highly indicative of eventual success in the fashion industry. Similarly, the first weekend box-office sales of a movie are indeed highly predictive of its eventual gross receipts. The Hollywood Stock Exchange (HSX, <www.hsx.com>) is an online prediction market that has been used to forecast box-office receipts of movies.

The HSX has 1.7 million registered users. New users are provided with two million "Hollywood dollars" (fake money) and can increase the value of their portfolio by trading. Elberse and Eliashberg[14] and Elberse and Anand[15]

FIGURE 1

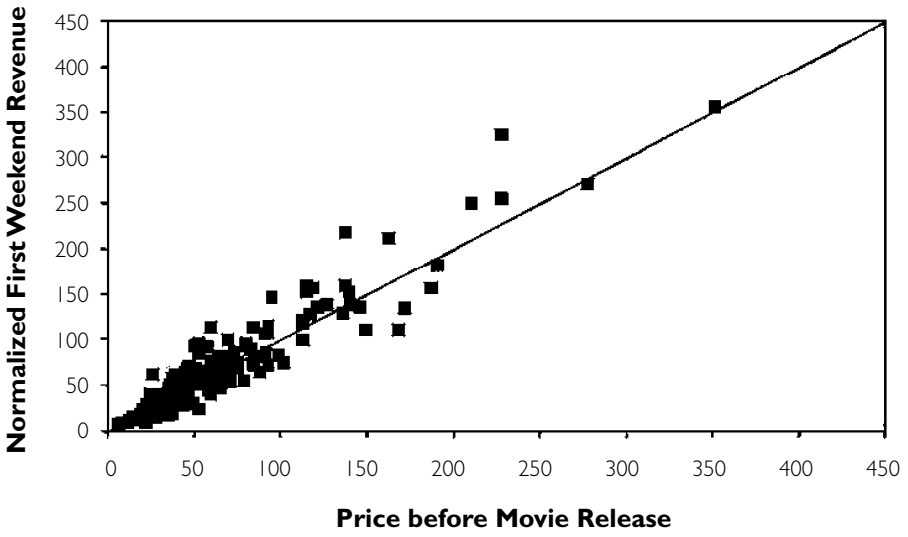

provide empirical evidence to show that HSX traders collectively produce excellent forecasts of actual box office returns.

More recently, Elberse investigated 192 movies that had been widely released as of January 2005, and finds the correlation between the prices before movie releases and normalized first weekend box office sales to have a correlation coefficient of 0.94.[16] Figure 1 shows this highly predictive relationship.

In the figure, if HSX were perfect, all the points would have been on the indicated straight line. As one can see, the points are remarkably close suggesting that HSX is highly accurate.

Information Technology Industries

Many information technology (IT) products, such as personal computers, have short life cycles, ranging from 12 to 18 months. New models with better feature sets frequently replace existing models. Historical sales data from old models are not useful for forecasting demand for new models because the improvements are significant and the customer tastes change rapidly. Companies in the IT industry, such as Hewlett-Packard (HP), have long understood the difficulties in forecasting demand when life cycles are short and demand variability is high. As a result, the firm actively seeks input from involved individuals through meetings and phone conversations. However, the process has been informal and the number of participants is severely limited by the schedule of the group.

In 1996, HP conducted its first field application of prediction markets to forecast sales of several families of products, including both new and existing

models.[17] The pilot application focused on gathering opinions from groups of involved executives (group size ranging from 7 to 26) in the HP marketing and finance organization. The group size was smaller than one would like because it was then a radical idea for many. A total of eleven tests were conducted. Eight of those eleven tests were conducted alongside of official forecasts, generated by a designated manager. Prediction markets were more accurate in six out of these eight tests. Table 1 reports these results.

TABLE 1

Event	Official Forecast Error	Prediction Market Error
1	13%	5%
2	60%	57%
3	9%	8%
4	32%	31%
5	30%	24%
6	4%	7%
7	0%	2%
8	28%	24%

These results show the potential of the prediction markets. We believe the initial results can be improved by increasing the number of participants and encouraging active participation.

Health Care Industry

One of the central challenges for health care officials is to predict an outbreak in seasonal influenza. An accurate prediction, even 1-2 weeks in advance, would allow the officials to install measures to contain the severity of an outbreak. Until recently, there is no method for predicting seasonal influenza activity accurately. Iowa Electronic Market started an influenza prediction market that attracted a lot of attention. The traders in this market are health care workers, which include microbiologists, epidemiologists, physicians, emergency room personnel, nurses, pharmacists, and administrators. The severity of the influenza activity is divided into five color-coded levels (Yellow, Green, Purple, Blue, and Red—as increasing level of severity) determined by Centers for Disease Control and Prevention. Participants are asked to trade which level of severity is most likely in each week in several geographical regions. The market has provided a two-to-four week advanced warning of when the flu season would hit a region. Table 2 shows its typical performance in 2004-2005 for predicting influenza activity in Iowa.

TABLE 2

Time Period	Mostly Likely Severity	Actual Severity
10/03/2004 - 10/09/2004	Yellow	Yellow
10/17/2004 - 10/23/2004	Yellow	Yellow
10/31/2004 - 11/06/2004	Yellow	Yellow
11/14/2004 - 11/20/2004	Yellow	Yellow
11/28/2004 - 12/04/2004	Yellow	Yellow
12/12/2004 - 12/18/2004	Green	Green
12/26/2004 - 01/01/2005	Green	Blue
01/09/2004 - 01/15/2004	Blue	Blue
01/23/2004 - 01/29/2004	Red	Red
02/06/2004 - 02/12/2004	Red	Red
02/20/2004 - 02/26/2004	Red	Red
03/06/2004 - 03/12/2004	Blue	Purple
03/20/2004 - 03/26/2004	Purple	Green
04/03/2004 - 04/09/2004	Green	Green

During this period, the market was correct in 11 out of 14 times. When the market missed a forecast, the error was never more than one level of severity away. This success enables health officials to better contain an influenza outbreak.

Summary

The above examples show that prediction markets can work in a wide range of industries even in cases where forecasting is challenging. These markets work because they could effectively aggregate information from a diverse group of individuals.

Potential Pitfalls

While the prediction market can be a powerful tool, it is not a panacea for all forecasting situations. There are many common pitfalls.

Prediction markets only aggregate information well if the number of participants is large. A small pool of participants can limit the amount of liquidity in the market and hence the market's ability to aggregate information. This runs counter to the principle of having a large crowd, as noted above. The success of a prediction market relies on active participation. If participants are not sufficiently motivated, they will have a low level of trading activities. Consequently, prices of the markets will not capture information well. This violates the principle of incentives. In combination, these problems lead to low market liquidity.

A recent application at a computer manufacturer is an illustration of this liquidity problem. This particular manufacturer experimented using a prediction market to forecast sales of computing products. This particular prediction market size had a small pool of participants (about 15). Most participants were executives from marketing and finance organizations. Since the size of the stakes was small ($50 per person on average) and these executives were busy, they did not pay enough attention to the prediction market. As a result, the market had low liquidity. Since information aggregation stems from trading activities, this low level of liquidity stopped active participants from executing trades that they wanted and this further exacerbated the problem. Despite these problems, the results were somewhat encouraging. The generated forecasts slightly beat the company's official forecasts in six out of eight times. However, the improvements were not substantial enough to justify the cost and time spent by the participants. As a consequence, this prediction market failed to be adopted as an ongoing forecasting tool.

The principles of crowd and independence imply that people who have relevant information should be included in the market. Otherwise, the aggregate information may be coarse and consequently the forecast may not be accurate. Put differently, the prediction market is a system of "garbage in, garbage out." In 2003, after the U.S. forces took Bagdad, Saddam Hussein was on the run and managed to avoid capture for several months. During that time, www.trades-

ports.com conducted a prediction market on whether Saddam Hussein would be captured before a certain deadline. The price of the scenario, saying Saddam would be captured before the deadline, was hovering around 9 cents over the dollar for a long time and only jumped up to 30 cents two days before his actual capture. This rapid adjustment is typical when new information is assimilated into the market. However, even at 30 cents, the market was predicting a strong chance (70%) that Saddam would not have been captured before the deadline. Obviously, the market was inaccurate, predicting only a 9% probability of capture until two days before his capture. This inability of the market to forecast accurately underscores the importance of the existence of "wisdom" in the crowd. Therefore, the importance of recruiting individuals who have information cannot be overemphasized. When the size of the group is limited by circumstances to be small, one can use related methods that are developed for this specific purpose.[18]

Prediction markets can only work well if players do not use other incentives to trade. That is, players trade solely to make money in the market. The market can fail if players have other incentives. For example, consider a prediction market for product demand. A product manager, who will gain a higher level of resources if the generated demand forecast is high, may want to mislead the market by pushing up prices of shares of high demand scenarios. Similarly, a Republican might want to depress prices of Democratic Party nominee in a political market in order to influence the ultimate political outcome (e.g., affect voter turnout). That is the reason why it is important to limit the stake of individuals in the market. The IEM limits the stake to $500 for the same reason. This size limit of the stakes also minimizes incentives for the participants to manipulate ultimate outcome. For example, a sales representative, who has bought shares of low demand scenarios, would not have deliberately reduced his own sales effort to make the prediction come true.

Conclusions

Prediction markets are "smart" markets that are capable of accurately predicting outcomes of uncertain future events. They have been successfully used in a wide range of settings and industries. Prediction markets work well if the underlying scientific principles (incentive, indicator, improvement, independence, and crowd) are adhered to. When a prediction market fails to predict future outcomes, it is often the case that one of the underlying principles does not hold.

As long as prediction markets are active, they always contain the most current "wisdom of crowds." Hence, they are better than a one-time survey in aggregating information from individuals because these same individuals have strong incentives to learn from each other through the price discovery process. Thus, the participants become smarter over time. In effect, a well-functioned market will eventually contain a large group of experts who freely share knowledge through their trading behaviors.

Prediction markets should be part of the forecasting tool kit for firms. A firm can use prediction markets to select the most promising new product ideas and to forecast demand for the selected new products before they are launched. The former use allows a firm to bet on the right product ideas and the latter ensures that the firm can better manage new product launches through better supply planning. Firms who invest in prediction markets are likely to gain competitive advantage because they have a powerful way to solicit input from employees (and perhaps others) and allow them to learn from each other. This collective wisdom at any time is then channeled to figuring out a stream of new products to fuel firms' revenue growth.

Notes

1. Booz Allen Hamilton, "New Product Development Study," 2004.
2. K. Pauwels, J. Silva-Risso, S. Srinivasan, and D. Hanssens, "New Products, Sales Promotions, and Firm Value: The Case of the Automobile Industry," *Journal of Marketing*, 68/4 (October 2004): 142-156.
3. V. Mahajan and J. Wind, "New Product Models: Practice, Shortcomings and Desired Improvements," *Journal of Product Innovation Management*, 9/2 (June 1992): 128-139.
4. M.F. Wolf, "Perspectives: News and Views of the Current Research Technology Management Scene," *Research Technology Management*, 33 (September/October 1990): 2-5.
5. <www.spectrum.ieee.org/print/2395>.
6. Forrester Research, Report on New Product Development, 2006.
7. PDMA, "The Product Development and Management Association Foundation's 2004 Comparative Performance Assessment Study," 2004.
8. AC Neilson, Report on New Product Development, 2003.
9. J. Surowiecki, *The Wisdom of Crowds* (New York, NY: Doubleday, 2004).
10. V. Smith, "An Experimental Study of Competitive Market Behavior," *Journal of Political Economy*, 70/3 (June 1962): 111-137.
11. Charles R. Plott and Shyam Sunder, "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models," *Journal of Political Economy*, 90/4 (August 1982): 663-698; Charles R. Plott and Shyam Sunder, "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," *Econometrica*, 56/5 (September 1988): 1085-1118.
12. Greg Burns, "The 'Election Futures' Market: More Accurate than Polls? As the U. of Iowa Goes, So Goes the Nation?" *Business Week*, November 11, 1996.
13. James Surowiecki, *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business* (New York, NY: Random House, 2004).
14. A. Elberse and Jehoshua Eliashberg, "Demand and Supply Dynamics for Sequentially Released Products in International Markets: The Case of Motion Pictures," *Marketing Science*, 22/3 (Summer 2003): 329-354.
15. A. Elberse and Bharat Anand, "The Effectiveness of Pre-Release Advertising for Motion Pictures," working paper, Harvard Business School, 2005.
16. A. Elberse, "The Power of Stars: Do Star Actors Drive the Success of Movies," working paper, Harvard Business School, 2006.
17. K. Chen, and Charles R. Plott, "Information Aggregation Mechanisms: Concept, Design, and Implementation for a Sales Forecasting Problem," working paper, March 2002.
18. K. Chen, Leslie Fine, and Bernardo Huberman, "Eliminating Public Knowledge Biases in Information-Aggregation Mechanisms," *Management Science*, 50/7 (July 2004): 983-994.