The Application of Prediction Markets to Business

by

Jared M. Schrieber

B.S., Supply Chain Management (1998)

Arizona State University

Submitted to the Engineering Systems Division
in Partial Fulfillment of the Requirements for the Degree of
Master of Engineering in Logistics

at the

Massachusetts Institute of Technology

June 2004

ABSTRACT

According to financial theory, open markets efficiently and effectively aggregate all available information about future events into their prices. Recent empirical evidence has shown that speculative markets, from gambling to web-games, are better at predicting the future than more commonly used statistical or survey-based forecasting methods. As a result, a number of companies have conducted experiments to evaluate the use of *prediction markets* as an alternative forecasting methodology. This paper offers a comprehensive framework for determining *when* and *how* prediction markets should be employed in a business context.

Thesis Supervisor: Abbott Weiss
Title: Senior Lecturer, Engineering Systems Division

Executive Summary

Introduction

Businesses rely on forecasts, whether formal or informal, to guide almost every important decision they make. Expectations related to revenue, margins, unit demand, capacity requirements, production yields, material prices, project budgets and timelines, competitor actions and customer behaviors are what drive corporate behavior. However, a recent survey of nearly 250 senior financial executives reported that 'difficulty accurately forecasting demand' was by-far the most common, significant problem within their companies' supply chains (CFO Research Services, 2003). It seems that no matter how much time and effort goes into predicting the future, it is practically assured that the prediction will still be wrong.

Given that accurate expectations about the future are critical to the successful management of a business, and knowing that those expectations will invariably err to some degree, companies must strive to find efficient methods for the systemic improvement of their predictive abilities. In order to improve forecasting methods, we must first evaluate the limitations of the statistical and survey-based approaches most-commonly employed in businesses today.

Statistical forecasts apply algorithms to transform predefined data into predictions, while survey-like methods generate forecasts from human opinions (Spann and Skiera, 2003). Most statistical models, such as time-series, regression, and multivariate approaches, have strict limitations regarding the timing and types of information they can assimilate. Once the inputs to these models are determined, the resulting forecasts are based upon a narrow set of assumptions about why the future will behave like the past. This is a primary reason why people tend to add their own judgments to statistically-based forecasts before making decisions.

While humans are more adept at assimilating new information from a multitude of sources at any time, it is quite challenging to create a single forecast from the opinions of group members. This task is often time-consuming and highly prone to biases caused by wishful thinking, politicized viewpoints, dominant personalities, and the influence of powerful members. Techniques used to generate forecasts from peoples' opinions include committees, surveys and Delphi analysis. Although these methods are theoretically capable of providing forecasts that represent the average participant's viewpoint, this is quite different than a true consensus that appropriately weights the knowledge and conviction of all informed individuals.

The recent development of *prediction markets* offers a promising new approach to forecasting. These speculative markets allow informed individuals to trade shares of claims (stocks) about the future. For example, a sample claim might state that, "Demand for Product A will fall between 1,000 and 1,499 units in Q3 of this year." If the claim proves to be true, then the claim pays $1, else it pays nothing. When the claim trades at a price of 30-cents, then the market is said to believe that there is a 30% chance that the claim will hold true. Of course, the person who bought shares of the claim at 30-cents believes that there is at least a 30% chance that the claim is true. The person who sold the claim believes that there is less than a 30% chance that the claim will become true.

The number of shares exchanged in the transaction serves as an indicator of the relative confidence that these traders have in their opinions. Non-action from other traders thus represents a lack of significant dissent regarding the claim's value (forecast) and demonstrates that the market has reached a consensus.

Theoretic and Empirical Evidence

The hypothesis that such markets have remarkable predictive powers is more than mere conjecture; it is well founded in economic theory, laboratory research and empirical studies. In 1945, Friedrich Hayek first suggested that open markets efficiently and effectively facilitate the aggregation and transmission of information through prices (Hayek 1945). Twenty years later, Eugene Fama offered the efficient market hypothesis which states that an efficient market continuously reflects all available information about future events into security prices (Fama, 1965). This implies that security prices reflect their true expected value and that no additional, available information can be combined with efficient market prices to improve the market's forecast accuracy. Finally, economic theory also explains three primary sources from which contingent commodities (claims that pay out depending on the outcome of an uncertain future event) derive social utility: they allow for efficient risk sharing and pooling, quickly assimilate information useful in making predictions, and can be enjoyable to trade as in the case of gambling (Wolfers and Zitzewitz, 2003).

Over the last several years, some leading companies have successfully employed prediction markets in an empirical context. For example, from 1996 to 1999, prediction markets were experimentally used at Hewlett Packard to predict the sales volume (dollars and/or units) of printers three months in advance (Chen and Plott, 2002). The prediction markets outperformed HP's official forecast 75% of the time and had a significantly lower absolute percent error, even though the markets' predictions were known prior to when the official forecasts were made. Since the prediction markets offered a complete set of claims that were related to specified ranges of demand (e.g. – Claim A is used for demand less than or equal to 1000, Claim B for demand between 1000 and 2000, and etc.), they provided additional information beyond a point-estimate. The result was an accurate representation of the probability distributions surrounding each forecast, which would likely be of significant benefit to individuals responsible for materials procurement, production capacity or inventory planning. In fact, it is known that HP secured materials from its suppliers at a lower cost than competitor Compaq simply because HP had a better ability to forecast and quantify the risk of demand fluctuations (CFO Research Services, 2003).

Prediction markets have been proven to offer more than just accurate forecasts however; they also offer distinct advantages over alternative forecasting methods such as: the immediacy of forecast updates, the provision of additional insights about the future, scalability to a diverse set of topics and participants, and built-in incentives that encourage more accurate and timely forecasts. When used properly, prediction markets can fundamentally change critical decision making processes by facilitating informed decision making in the face of uncertainty.

Guidelines for When and How to Use a Prediction Market

Prediction markets, however, are not a panacea for a company's forecasting woes. In particular, prediction markets seem best suited only to those situations where:

- there are a number of significantly different actions that can be taken depending on the specific value(s) of the forecast (e.g, capacity expansions to support demand through new processes, shifts, equipment, lines or factories); and/or,
- the complexity of inputs that could influence a forecast or decision making process make it cumbersome to aggregate the information in a meaningful and timely manner (e.g., impacts of regional promotions and competitor pricing on aggregate demand for CPG products); and/or,
- there is high risk of biased decision making and/or the need to make decisions prior to an updated and approved forecast becoming available; and/or,
- history is unlikely to repeat itself and expectations regarding potential outcomes may vary greatly as new information becomes available (e.g., new product introductions).

It is also important to note that the use of prediction markets does not exclude the use of statistical or survey-based forecasting methods. In fact, markets encourage their participants to trade based on statistical analysis and expert judgment.

Once the decision to implement a prediction market is made, it is important to consider a number of design variables that can impact the type and quality of information revealed by the market. Specifically, this thesis offers an in-depth analysis of the key considerations involving the following design elements: Forecast Objectives, Participation, Trading Mechanisms, Claim Structures, Claim Definitions, Incentives, Account Management, and the Trading Interface.

Applications Beyond Forecasting

Markets are clearly adept at the efficient allocation of scarce resources to those who are most willing to pay. For example, when several business groups compete for shared production capacity, an internal 'capacity market' may offer significant benefits over traditional methods for capacity planning and allocation. In such a market, a factory could sell future shifts at their marginal operating cost and allow the business groups to buy and sell these shifts amongst each other as their forecasts update and requirements change. Each 'shift' would represent a pre-determined amount of production volume for each business group's products. (For example, Group A could expect 500 units of Product Y or 750 units of Product Z for each shift of production it buys.) When capacity runs tight, the market price of a production shift will grow and may drive the factory to consider selling more shifts (e.g., overtime shifts) or cause it to re-evaluate a capacity expansion project. In addition, the 'capacity market' should help ensure that higher-margin products are given priority in a capacity crunch because these products can better afford the higher production (shift) costs. Finally, the 'capacity market' prices should give managers strong and objective insights into whether or not outsourcing production would be a more viable alternative.

Conclusion

Prediction markets offer decision makers a novel approach for evaluating an uncertain future. The accuracy, immediacy, and insights of information provided by these markets can far surpass the output of traditional forecasting approaches in use today and fundamentally change the nature of critical decision making processes. Prediction markets also contain built-in incentive mechanisms that encourage the creation and sharing of tradable information used in forecasts, further increasing their value. Finally, prediction markets can successfully aggregate information from a wide range of sources and individuals, performing best in support of the most complex of decision making scenarios.

Despite all of the advantages of prediction markets, their adoption has been somewhat slow. IT managers are often overly concerned with proving that the technology will work within the context of their business before making a decision to implement the technology (Gebert, 2004). While business managers typically like the idea, they frequently believe it is too complex for their employees to use. Although this misperception is usually overcome through a simple market simulation game, many of these managers also believe that their employees don't have sufficient knowledge to serve as 'informed' participants. Therefore, they are apprehensive about the amount of time their employees will spend in training and trading activities. Finally, the CEO of Incentive Markets, Carol Gebert, has found that the single biggest obstacles to implementing a corporate prediction market are legal and political (e.g., moral hazard concerns) conflicts with a company's existing compensation system. If prediction markets are not given an appropriate budget for rewarding participants, they can lose their incentive value. This exacerbates the thin-market problem that already plagues their application to most business scenarios (Chen, 2004).

For all that is known about prediction markets, many important questions remain unanswered, thus hindering their wide-spread corporate adoption. In addition, software and training costs can make implementing prediction markets an expensive proposition relative to most statistical and survey-based forecasting methods. Therefore, businesses should focus their deployments on forecasting problems that can achieve significant gains from the unique benefits offered by prediction markets. These include the ability to perform well under high-degrees of forecast sensitivity, forecasting complexity, future uncertainty, and political bias in the decision making process. In the case of demand planning, this suggests that forecasts used for New Product Introductions (NPI) and longer-term capacity requirements planning may be the most appropriate uses of prediction markets.

While prediction markets are certainly adept as forecasting tools, their 'killer application' may be in the realm of resource allocation. These markets seem perfectly suited for helping companies prioritize projects and product concepts for funding, allocate production capacity to business units, and ensure that budget dollars and highly skilled employees are put to work in those areas with the greatest need. Resource allocation decisions are often among the most subjective and contentious decisions made in business today, yet without structured methodologies for carefully weighting opinions, an honest consensus that participants trust is rarely achieved. Rather than resolving such issues with an 'iron fist', prediction markets offer senior decision makers an 'invisible hand' that gently guides employees to unwittingly make sacrifices for the common good of the company.

Bibliography

Berg, J., Rietz, T., (2003) “Prediction Markets as Decision Support Systems”, *Information Systems Frontiers*, 5(1), pp79-93, 2003.

CFO Research Services, (2003) “CFOs and the Supply Chain,” www.cfo-research.com, 2003.

Chase Jr., C., (2000) “Composite Forecasting: Combining Forecasts for Improved Accuracy”, *The Journal of Business Forecasting Methods & Systems*, Summer 2000; 19, 2, pp. 2-6, 2000.

Chen, K.Y., (2004) Experimental Economist at HP Labs and administrator of prediction market pilot at HP. [Interview] April 22nd, 2004.

Chen, K.Y., Plott, C., (2002) “Information Aggregation Mechanisms:
Concept, Design and Field Implementation for a Sales Forecasting Problem.”, California Institute of Technology, Social Science Working Paper 1131, 2002.

Fama, E. F., (1965) "Random Walks in Stock Market Prices," *Financial Analysts Journal*, September/October 1965.

Gebert, C., (2004) CEO of Incentive Markets. [Interview] April 20th, 2004.

Hayek, F. A., (1945) “The use of knowledge in society.”, *American Economic Review*, 35, pp. 519-530, 1945.

Kittlitz, K., (2004) Founder and administrator of Foresight Exchange. [Interview] April 13th, 2004.

O’Hara, M., (2001) “Overview: market structure issues in market liquidity”, April 2001. [On-line] http://www.bis.org/publ/bispap02a.pdf, March 8th, 2004.

Spann, M., Skiera, B., (2003) “Internet-Based Virtual Stock Markets for Business Forecasting”, *Management Science*, INFORMS, Vol. 49, No. 10, October 2003, pp. 1310-1326.

Wolfers, J., Zitzewitz, E., (2003) “Prediction Markets”, November 19th, 2003. [On-line] http://faculty-gsb.stanford.edu/wolfers/Papers/Predictionmarkets.pdf, February 2nd, 2004.

