

100

The Application of Prediction Markets to Business

by

Jared M. Schrieber

B.S., Supply Chain Management (1998)

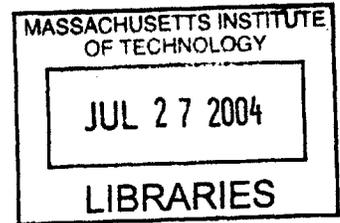
Arizona State University

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



© 2004 Jared Schrieber. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly
paper and electronic copies of this thesis document in whole or in part.

Signature of Autho.....
Engineering Systems Division
May 7th, 2004

Certified by
Abbott Weiss
Senior Lecturer
Thesis Supervisor

Accepted by.....
Yossi Sheffi
Director, MIT Center for Transportation & Logistics

The Application of Prediction Markets to Business

by

Jared M. Schrieber

Submitted to the Engineering Systems Division
on May 7th, 2004 in Partial Fulfillment of the
Requirements for the Degree of Master of Engineering in Logistics

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

Table of Contents

1. Introduction	4
2. Literature Review	8
2.1 Financial Markets & Economic Theory	8
2.2 Gambling & Betting Markets	9
2.3 Idea Futures & Web Games	10
2.4 Business Applications	12
3. Benefits of Prediction Markets	15
3.1 Accuracy	15
3.2 Immediacy	15
3.3 Insight	16
3.4 Scalability	17
3.5 Incentives	17
4. Prediction Market Implementation	19
4.1 Guidelines for When to Use a Prediction Market	19
4.2 Considerations for Prediction Market Design	22
4.2a Forecast Objectives	23
4.2b Participation	24
4.2c Trading Mechanisms	27
4.2d Claim Structure	28
4.2e Claim Definition	32
4.2f Incentives	37
4.2g Account Balances	38
4.2h Trading Interface	39
5. Applications Beyond Forecasting	41
5.1 Resource Allocation	41
5.2 Risk Hedging	43
6. Conclusion	45
7. Bibliography	47

Acknowledgements

I would like to acknowledge the generous support and guidance of Dr. Abbott Weiss and my wife, Reka Cseresnyes. This has been a truly rewarding experience. Thank you both.

1. 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. A single demand forecast alone may be used to drive important decisions throughout a company, including: sales quotas, marketing budgets, materials procurement contracts, production levels and factory hiring, capital investment for manufacturing capacity, and more. 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. Simply removing a single percentage point of error from a demand forecast can add millions of dollars to a company’s bottom-line (Kahn, 2003). In the extreme case of Cisco, poor demand forecasting resulted in a single-quarter loss of \$2 billion when the resulting excess inventory had to be written-off.

In order to improve forecasting methods, we must first evaluate the limitations of approaches most-commonly employed in businesses today. There are two broad categories of forecasting methods used in business today: statistical approaches and survey-like methods. Statistical forecasts apply algorithms to transform predefined data into predictions, while survey-like methods generate forecasts from human opinions (Spann and Skiera, 2003). A primary difficulty in using such methods is the time they require to incorporate new information into a forecast. 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*¹ 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 that 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. 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.

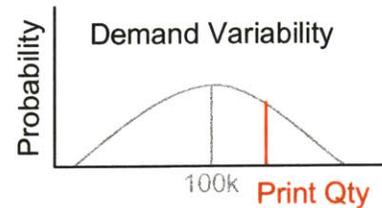
¹ In the literature, there is no common name used to describe what this paper will refer to as 'prediction markets'. Alternative names with similar meanings include: idea futures, idea stocks, artificial markets, decision markets, forecasting markets, information markets, information aggregation mechanisms, virtual stock markets, and even web-games.

Prediction markets have been proven to offer more than just accurate forecasts however. They have demonstrated an ability to support forecasts for almost any future event, quickly aggregate input from any number of sources, continuously update their predictions, and provide insight into the future's variability. A simple *newsvendor problem*² demonstrates the value of understanding the uncertainty within a forecast (see sidebar), yet most businesses base their most important decisions on a single 'best guess' (or 'best case') point-estimate, failing to quantify and plan for uncertainty. Fisher et al (1994) argue that a probabilistic model should underlie all supply chain plans: "Contrary to what many believe, market uncertainty is a manageable risk." When used properly, the probability distributions created by prediction markets fundamentally change critical decision making processes by guiding informed decision making in the face of uncertainty.

Simple Newsvendor Problem

USA News Inc. sells its daily newspaper for \$1.00 although each copy only costs \$0.35 to print and distribute. If its demand forecast for tomorrow's paper is 100,000, how many copies should USA News print?

The answer depends on the level of uncertainty in the forecast. Because these newspapers have a gross margin percentage of 65%, USA News can achieve optimal profits by printing enough copies to fully support 65% of all possible demand scenarios.



As demand variability increases, so does the optimal print quantity for USA News. This is because the cost of lost-sales exceeds the cost of excess inventory for these high-margin papers. If USA News were to only print the 100,000 copies in its demand forecast, it would experience some level of stock-outs 50% of the time. This would result in lower revenue, reduced profits and the potential for lost customers.

The objective of this thesis is to look beyond the questions of 'if' and 'why' prediction markets should be used in business and address the subsequent issues concerning 'when' and 'how' prediction markets should be deployed in a business context. Section 2, Literature Review, provides a comprehensive history of prediction markets from their roots in financial theory through the lessons learned from recent experiments in business

² For a more thorough understanding of the newsvendor problem and the mathematics behind it, see: <http://www.msb.edu/faculty/bios/schmidt/newvendor.pdf>

forecasting. Section 3, Benefits of Prediction Markets, draws upon both theory and empirical evidence to identify five distinct benefits of prediction markets over alternative forecasting methods. Section 4, Prediction Market Implementation, provides a framework for evaluating when a prediction market should be employed and a detailed analysis of the considerations for optimal market design in support of forecast-based decisions. Section 5, Applications Beyond Forecasting, proposes some additional uses for internal markets such as resource allocation and risk hedging. Section 6, Conclusions, summarizes some of the key findings presented in this thesis and offers some ideas for areas of further research.

Those interested in only a high-level understanding of prediction markets may find the following sections sufficient:

- Section 2.1 - Financial Markets & Economic Theory;
- Section 2.4 - Business Applications;
- Section 3 - Benefits of Prediction Markets;
- Section 4.1 - Guidelines for When to Use a Prediction Market;
- Section 4.2d - Claim Structure;
- Section 5.1 - Resource Allocation; and,
- Section 6 - Conclusion.

2. Literature Review

The roots of prediction markets can be found in two distinct fields: financial markets and gambling. Only recently have the lines between these practices blurred as speculative, prediction markets take on the form of web games, betting exchanges, and corporate forecasting tools.

2.1 Financial Markets & Economic Theory

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. Vernon Smith demonstrated that people behave rationally in experimental markets, and was awarded the 2003 Nobel Prize in Economics for using experimental markets to prove and refine important theories about market behavior.

In 1972, Richard Sandor introduced the concepts of financial (interest rate) futures and weather futures (Pomerantz, 2003). Although futures trading for tangible commodities had long existed, the idea of trading the 'untangible' was foreign to the financial community. These ideas have been remarkable successes, with over 200 million interest rate futures trades taking place each year and weather-derivatives trade approaching \$16 billion per annum. In fact, the weather futures market has even proven to be a valuable tool in the forecasting of weather (Kulkarni, 2003). In addition, guidelines for the trade of sulphur dioxide pollution credits (futures) were included in the 1990 Clean Air Act and are widely credited for the legislation's success. This demonstrates that a carefully structured futures market can be a powerful tool to motivate desirable behaviors from otherwise disinterested parties.

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).

2.2 Gambling & Betting Markets

While gambling has existed for centuries, only recently has published research appeared related to the predictive accuracy of betting markets. In 1971, Winkler found that published book-maker point spreads significantly outperformed a set of students and professors in predicting the outcomes of football games (Winkler, 1971). He also found that a weighted-average consensus of the students and professors had strong predictive abilities, although it was still outperformed by the point spreads from the bookies (which represent the consensus from a large betting market). Winkler also discovered that the use of scoring rules and wagers in his study caused participants to carefully assess their performance and decision making process. Finally, Winkler concluded that the use of real-money wagers drove higher participant interest levels and superior performance than simple scoring rules.

An in depth study of the predictive power of Nevada's NFL betting market found that the market outperformed newspaper sports editors, power scores based on probit regressions, and other forecasting methods in the prediction of game outcomes (Boulier and Stekler, 2001). In addition, research indicates that changes in published point-spreads greatly improve the likelihood that the betting market will accurately predict a game's outcome (Gandar et al, 1998). This demonstrates that changes in market sentiment reflect valuable information that usually leads to improved forecasts. Finally, there have been a number of papers that evaluate whether or not various betting markets are truly efficient. Although a number of anomalies have been discovered, the vast majority of them have already been disproved and the general consensus is that most betting markets are truly efficient.

2.3 Idea Futures & Web Games

The historical line between betting markets and financial futures blurred when Robin Hanson introduced the concept of 'idea futures' in 1988 (Hanson, 1996).

“Such markets pool the information that is known to diverse individuals into a common resource, and have many advantages over standard institutions for information aggregation, such as news media, peer review, trials, and opinion polls. Speculative markets are decentralized and relatively egalitarian, and can offer direct, concise, timely, and precise estimates in answer to questions we pose. These estimates are self-consistent across a wide range of issues and respond quickly to new information. They also seem to be cheap and relatively accurate.” (Hanson, 1999)

The first web-based artificial market, Idea Futures, covering science, math and miscellaneous other categories was implemented in 1994 (Kittlitz, 1999). It was the first trading system on the internet, although it only used play-money. Over the last several years, the market (now called the Foresight Exchange) has been proven to accurately predict events as often as it would expect to. For example, when evaluating all trades made at a price of .75 (out of 1.00), one finds them to pay-out roughly 75% of the time. However, a 'favorite-longshot bias' did appear, such that events with just a 1% chance of occurring were usually traded at higher values. One explanation for this is that such low prices afford players with contrary opinions the ability to easily influence the consensus market price without much capital or risk. Ken Kittlitz, who administers the Foresight Exchange, has also found that the market rapidly incorporates new information into its prices. Finally, he has found that the claims attracting the greatest trading interest are usually the least ambiguous to judge (e.g., George W. Bush will be President of the United States of America on February 1st, 2005).

The use of markets to forecast the outcome of election results is quite common. For example, the Iowa Electronic Market has consistently outperformed traditional polls and expert opinions in predicting election results (Feder, 2002). In addition, this market has employed more complex claim structures that allow for greater insights into participant beliefs. It has been argued that Republican delegates using the Iowa Electronic Market as

a guide should have known that Dole would not defeat Clinton in the 1996 presidential election, but that selecting Colin Powell instead may have turned the odds in their favor (Berg and Rietz, 2003).

In 2001, Pennock et al found that the forecasts of play-money markets such as the Hollywood Stock Exchange, the Foresight Exchange, and the Formula One Pick Six competition were generally more reliable than those of other sources, including expert opinion (Pennock et al, 2001). In their report, the authors stated, “The games themselves appear to serve as a mechanism for collecting, merging, and cleaning data from human experts, naturally handling some of the more difficult steps in a typical data mining application.”

Research from these web-games also indicates that markets with fewer traders and/or lower financial incentives tend to incorporate new information a little more slowly than their larger counterparts (Wolfers and Zitzewitz, 2003). Wolfers and Zitzewitz also demonstrated that lower volume exchanges tend to exhibit a ‘favorite-longshot bias’ more than higher volume markets. Speculative bubbles, however, appear to be less likely in prediction markets because informed participants generally have sufficient leverage to keep prices in-check. Although attempts to manipulate market prices have occurred in these markets, their impact has been insignificant and the duration of their impact was quite short. The research of Wolfers and Zitzewitz also indicates that the performance of prediction markets versus the average (consensus) of a large group of expert forecasters is often quite similar. In such cases the prediction markets add value by more accurately quantifying the uncertainty in the forecast (e.g., standard error or the shape of probability distribution curve).

The previously mentioned Hollywood Stock Exchange is more than just a web game, it is one of the movie industry’s most trusted predictors of how new films will perform at the box office. Nearly 1-million registered users actively trade ‘movie-stocks’ and ‘star-bonds’ using play-money. Summarized data generated by their trades is then sold to entertainment companies seeking to understand public opinion and forecast revenues for

their upcoming movies (Feder, 2002). In addition, a number of pharmaceutical companies have recently turned to public web-games (NewsFutures.com) as an attempt to predict demand for their drugs. Although the jury is still out on the results of this experiment, there are several examples where prediction markets have been used successfully in industry.

2.4 Business Applications

In 1997, Gerhard Ortner developed and administered a prediction market used by Siemens to predict a large software project's completion date (Ortner, 1998). The project had roughly 200 employees involved over a 6 month period. The prediction market included 60 participants wagering 100 Austrian Schillings (ATS) of their own money (although they were guaranteed 200 ATS for participating). Although most of the participants were unfamiliar with trading stocks in real-life, they were able to learn with little effort and became active traders early in the life of the market.

Throughout the project, claim prices quickly adjusted to reflect new information (prior to the public announcement of that information), thus achieving one of the goals of the prediction market: to encourage immediate sharing of information that may impact the project's deadline. Almost 3-months before the scheduled deadline, the claim designed to predict the project's actual completion date had already predicted the 2-week delay that would inevitably occur. One month prior to the project's final due-date, the price of 'on-time' contracts fell to almost zero, although the management team still saw on-time completion as feasible based on their traditional planning methods. The study also found that the vast majority of orders were placed during break-times, and were not seen to have impacted productivity in any way. Although half of the traders lost-money in the market, almost 90% of them stated that they would participate again.

From 1996 to 1999, Kay-Yut Chen and Charles Plott administered prediction markets at Hewlett Packard on twelve occasions in order to predict the sales volume (dollars and/or units) of printers three months in advance (Chen and Plott, 2002). Participation was limited to marketing and finance, with the exception of a few uninformed traders from

HP Labs used to add liquidity. Each time, the market was open to between 7 and 24 participants during lunch and after-work-hours. Trading was done with apparently nominal amounts of real-money, fully subsidized by the company.

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).

In 2001, a large German telecommunications company experimented with a prediction market to forecast demand for five of its services. The market was found to be more accurate than the four statistical methods it was compared against, as well as the consensus of an expert panel (Spann and Skiera, 2003). The study had 20 participants (12 of which traded actively) from the marketing and planning departments trading claims for each of the five services. The market used play-money and the only remuneration was a \$500 mobile-phone set to be given the single trader with the highest final account balance.

In 1999 and 2000, three experiments were run to test a prediction market's ability to be used as a market research tool, in place of traditional mechanisms such as surveys, focus groups and conjoint analysis (Chan et al, 2002). Participants in the experiment bought and sold shares of up to 10 comparable product concepts (e.g., different types of bike pumps) based on their perceived value of each product. Some of the products were real,

while others were truly just concepts; and, the price of one product was fixed to help anchor the value of the market's currency. The relative product share prices proved to be consistent with the product preference results achieved by conventional methods. (However, no tests were performed to determine which method actually did the best job of predicting end-consumer preferences.) While more research is required, there is great promise for the use of prediction markets to remove time, costs and bias from market research efforts.

3. Benefits of Prediction Markets

As alluded to in the above literature review, prediction markets offer distinct benefits over alternative forecasting methods used by businesses today. This section describes how prediction markets derive their value in terms of accuracy, immediacy, insight, scalability and incentives.

3.1 Accuracy

Economic theory suggests efficient market prices are the best possible predictors of future events. Laboratory experiments have demonstrated that prediction markets tend to perform just as well as the single best individual (Chen, 2004), without requiring knowledge of who that individual is in advance. In practice, prediction markets have already been proven to outperform common forecasting practices such as polls, expert surveys, advanced statistical methods and hybrid techniques that allow experts to apply judgment to statistically-based forecasts. Such markets are not limited by the source or type of information they can aggregate and are capable of performing best under difficult circumstances, such as when bits of relevant information are highly dispersed among participants.

Prediction markets automatically penalize inaccurate forecasters through the depletion of account balances when they are wrong. The resulting changes in account balances give a capital advantage to proven forecasters that can be used to exert more influence over market prices. The natural allocation of forecasting weight, combined with participant anonymity, makes prediction markets particularly adept at removing bias by preventing dominant personalities or powerful individuals from consistently influencing consensus-based forecasts in the wrong direction.

3.2 Immediacy

Prediction markets continuously reflect the knowledge of their participants and quickly assimilate new information as it becomes available. In prediction markets, people are rewarded for quickly applying accurate ‘insider’ information in the form of trades. As a result, market-based forecasts do not have to wait for the next reporting cycle, staff

meeting or committee approval to reflect important information. Even ‘bad-news’ is immediately shared in such markets because anonymous participation bypasses the potential for stigmas and other negative repercussions.

Since prediction markets are capable of quickly assimilating new information into prices, it stands to reason that the market will continuously reflect the best possible prediction of the future. If business processes are engineered to take advantage of continuous forecasts, companies can be assured that they are always making decisions based upon the best possible information. Imagine a manufacturer that uses a monthly forecast to determine its production schedule. If the business environment changes between when an official forecast is published and production for a given product occurs, then a prediction market could enable the factory to respond appropriately, and with confidence.

3.3 Insight

There will always be error in forecasts. Therefore, every organization should understand the risks inherent within each forecast and make informed decisions that proactively mitigate the consequences of likely deviations. Beyond point forecasts (e.g., Demand in Q2 is expected to be 1250 units), prediction markets can provide valuable insights into risk through variances and the shape of probability distributions. When tied to demand forecasts, such information can be critical to those responsible for supplier contracts, manufacturing capacity and inventory planning.

Prediction markets provide managers with additional insights as well. For example, a market allows them to easily assess how much impact the introduction of particular information has on a forecast. If new information is released to market participants and the price for a given claim (stock) jumps by 20%, then one can infer the importance that such information has on the expected outcome. Managers could then focus their efforts on improving processes that would yield such important information as early as possible. In addition, account balances (or percent of correct trades) can be used to identify which individuals or organizations have the most relevant information regarding a particular

forecast or set of forecasts. This may lead managers to reengineer decision making processes to allow for more input from these proven experts.

3.4 Scalability

Most forecasting processes employed by businesses today lack the ability to efficiently adapt to changes in the type and sources of new information. Prediction markets are capable of transforming practically limitless amounts of qualitative information into accurate, quantitative assessments of the future. Unlike their traditional alternatives, these markets continue to operate efficiently as the number of inputs and/or participants increases. (In fact, prediction markets generally function better with more participants and higher trading volume.) Thanks to the internet, prediction markets also have the ability to aggregate information at any time, from any place. This enables large, multinational corporations to give all informed employees the opportunity to contribute to a collaborative forecasting process. The addition of participants from additional organizations and/or lower levels of responsibility is likely to improve forecast accuracy while transforming these people into valued forecast contributors and stakeholders.

3.5 Incentives

A primary reason for the superior performance of prediction markets is that they have built-in feedback mechanisms and incentive structures. The potential for rewards and losses (either intrinsic or extrinsic) based on forecasting performance has a number of positive effects on the forecasting process:

- Rewards those who share information through trades that lead to accurate predictions;
- Motivates people to share relevant information in a timely manner;
- Dampens personal and political motives for withholding information;
- Attracts informed individuals to participate in the forecasting process, while discouraging those that have bad information or intuition;
- Encourages people to carefully evaluate their decision making processes;
- Drives participants to find new ways of obtaining useful information; and,
- Forces individuals to honestly assess the confidence they have in their opinion.

In addition, a company can easily adjust a market's compensation mechanisms to encourage more or less diligence in the forecasting process. For example, forecasts that are more important to the company's performance may have higher remuneration attached to successful predictions. This in turn, would drive participants to focus more attention to gathering and applying information towards particular claims, naturally reprioritizing their relative efforts.

4. Prediction Market Implementation

The question surrounding prediction markets is no longer whether they can add value in a corporate setting, but should be: “When and how should prediction markets be used in business?”

4.1 Guidelines for When to Use a Prediction Market

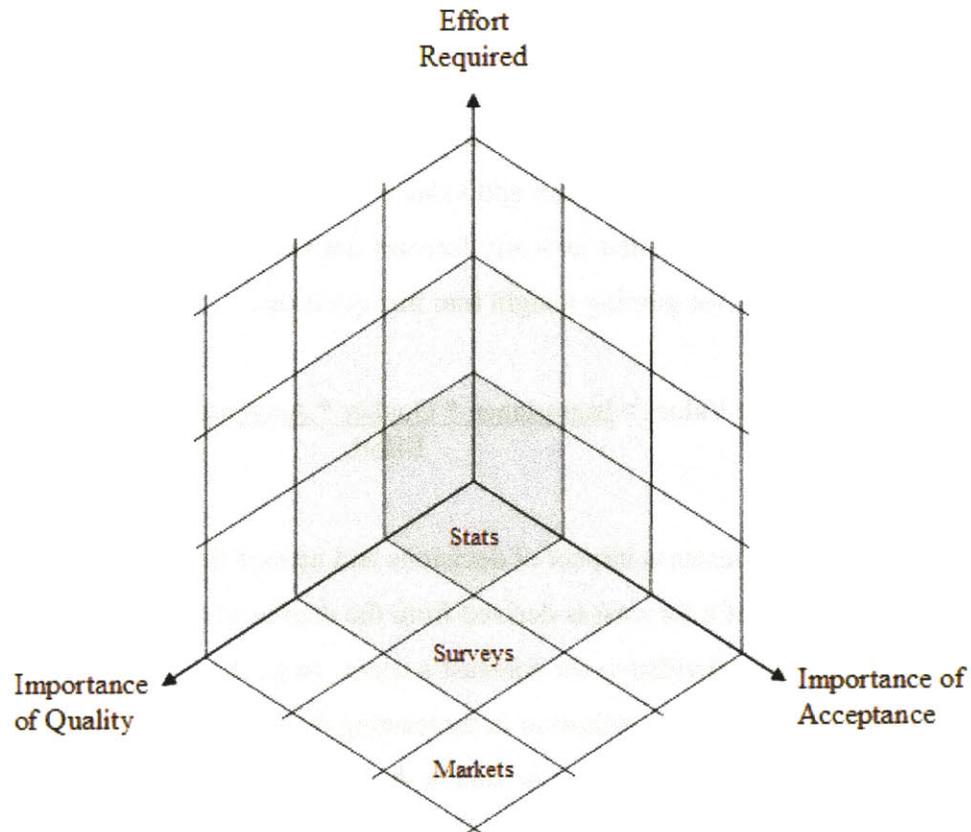
Determining when a prediction market can add value over alternative forecasting methods first requires one to consider how any forecast derives its value. The following formula is offered as a tool for gaining insight into this question:

$$\text{Forecast Value} = \frac{\text{Importance} * \text{Quality} * \text{Acceptance}}{\text{Effort}}$$

Importance relates to the business impact of decisions and actions that are influenced by the forecast. The *quality* of a forecast is derived from the degree of accuracy and the usefulness of information provided to the forecast’s users. (e.g., A point-estimate is of lower quality than a probability distribution in forecasting demand.) *Acceptance* measures the extent to which the forecast is used in the decision making process. If a manager looks at a forecast, and then applies additional judgment before taking action, the forecast is not fully accepted. More importantly, if a number of managers in different departments each apply their own judgments before making decisions, then a forecast’s acceptance is quite low. In addition, a forecast that is not timely may not be accepted at all, and could reduce the overall forecast value to zero. Finally, *effort* is a reflection of the amount of resources (people, time and money) invested in formulating a forecast.

Although the importance of a forecast does not change depending on the forecasting method employed, the above formula implies that quality and acceptance have a greater impact on the value of important decisions than less important ones. This is quite intuitive in practice and explains why greater effort is generally exerted on improving the quality and acceptance of the most important forecasts. Thus, the use of appropriate forecasting methodologies for various tasks can help optimize the use of scarce resources

to maximize forecast value. The diagram below depicts one method of identifying the general forecasting technique that is most appropriate for a particular decision making problem.



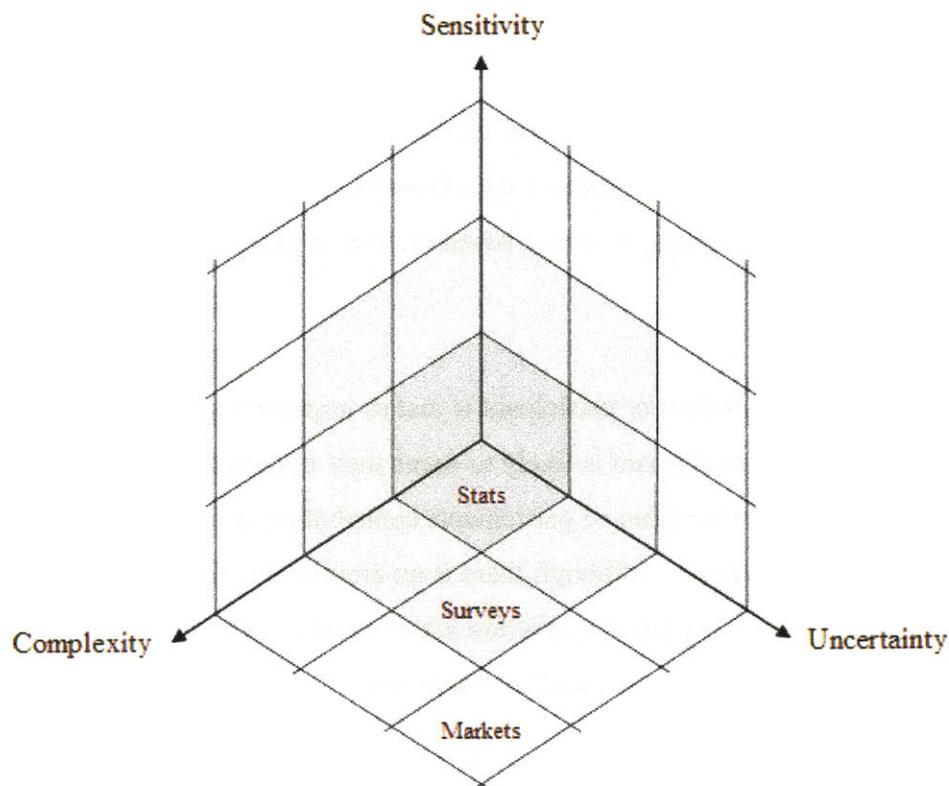
Under this depiction, prediction markets add the most value when:

- high levels of effort are required to aggregate predictive inputs; and/or,
- the impending decisions are highly sensitive to the quality of the information taken from the forecast; 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.

It is important to note that the use of survey methods does not preclude the use of statistical methods in deriving a particular forecast, just as the use of prediction markets does not exclude people from trading based on statistical analysis or expert judgment. In addition, many experts who participate in survey-type forecasting actually derive their knowledge from underlying statistical models and analysis. A prediction market that

provides meaningful incentives to its participants will actually encourage the use of these other techniques to justify informed trading decisions.

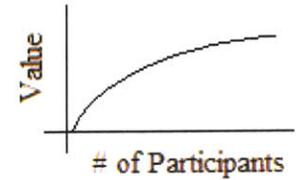
Another framework for considering when to use a prediction market over other forecasting methods is shown below.



In this diagram, prediction markets are best applied in 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,
- history is unlikely to repeat itself and expectations regarding potential outcomes may vary greatly as new information becomes available (e.g., new product introductions).

As stated above, the complexity of aggregating forecast information and the effort required to generate a meaningful forecast are important indicators of when a prediction market may be of significant value. However, prediction markets also require that useful information be dispersed across a number of informed individuals. If a very small number of people have a major informational advantage over other market participants, the prediction market is likely to fail. For example, if two sales representatives support clients that represent the vast majority of a company's demand, then their insider knowledge could discourage others from trading claims used for demand forecasting. In this case, a survey-based method is probably a better approach to forecasting.



The number of tradable claims per participant is just as important as the number of participants. Since each participant is likely to focus their trading on just a few claims, this effectively reduces the number of participants contributing their knowledge to the accuracy of each claim's price. Although there is no empirical evidence that quantifies the effects of having too few traders and/or too many claims, most authors recommend at least five informed participants per claim (or claim-set).

4.2 Considerations for Prediction Market Design

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. We have already seen how important the natures of the decisions being made from a forecast are to determining the appropriate methodology. This understanding is doubly important when actually designing a prediction market. While the work of Hanson (1999), Spann and Skiera (2003), and Wolfers and Zitzewitz (2003) offer some guidelines for market design, this section offers a more comprehensive and detailed framework. Specifically, this section 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.

4.2a Forecast Objectives

The primary objective of a forecast is to support important decisions; therefore it is essential to understand exactly how these decisions are influenced by the forecasting process. Mapping the flow of forecast information throughout the enterprise (and potentially beyond) is a valuable exercise that may help achieve this understanding. The following questions may serve as a useful guide in this process:

- What is being forecasted and at what level(s) of detail? (e.g., Forecast Horizon, Time Buckets, Product Levels, and Location Levels)
- Who currently creates the forecast and what information and procedures do they use to create it?
- Who provides inputs to the forecast and what information do they provide?
- Who uses the forecast, and what do they use it for?
- How sensitive are the decisions these people make to a range of potential forecast outcomes? What is a significant change in the forecast according to them? What increment of change in a forecast could lead to a different decision on their part? What is the difference in cost or risk between those decisions?
- What is the impact of forecast error and how does this impact change in relation to the size of the forecast error?
- Do the users judge the forecasts in anyway or feel that they could do a better job of forecasting?
- What additional information would they like to see incorporated into the forecast, either as an input or output?
- Who else uses this forecast or a variation of it?

Even if a prediction market is not used, the above process questions may be useful in discovering ways to improve upon existing forecast methodologies. In addition, it may lead to some fruitful discussions regarding whether or not the forecast is even necessary. For example, some forecasts may be replaced by improved information sharing or the reduction of lead-times (e.g. - supplier, production or decision).

4.2b Participation

Mapping the forecast process should help identify individuals who have unique perspectives or relevant information pertaining to the future. Ideally, these individuals will comprise the set of market participants. In addition, there may be other informed individuals in adjacent organizations, or even outside the company, who are not currently associated with the forecasting process, but could add value to the collective wisdom of a prediction market. Unlike traditional forecasting methods, the costs of involving additional people in the creation of a prediction market's forecast are quite small. Therefore, additional participants should be able to join-in at any time if they are likely to add new information to the market (Ortner, 1998). If the amount of time participants may spend trading is a productivity concern, then it may make sense to restrict trading to certain time-periods (e.g., lunch-time), rather than preventing people with distinct information from participating. A reduction in market hours can actually lead to more intense trading because participants are likely to encounter a market full of activity on a consistent basis.

There are very few reasons why an employee who desires to be a part of a prediction market should be turned away. Legitimate concern regarding the confidentiality of forecast information is one such reason. Another reason to restrict participation is the potential for 'moral hazard' situations, where an employee has the power to profit from causing a negative outcome to come about. For example, in the Siemens project management case described earlier, first-line managers were not allowed to participate in the market because they could easily manipulate outcomes to achieve trading profits. Although such employee gamesmanship with quotas and budgets is commonplace, it should be avoided whenever possible. However, this does not mean that decision makers should be prevented from trading just because they can influence the outcome in question. These people often have the best information to share with the market, so it is important to ensure that potential incentives for hazardous behavior are mitigated (e.g., by reducing market payouts and/or prizes). Such proactive steps will also remedy concerns over potential insider-trading by those in power that would otherwise deter others from participating (Wolfers and Zitzewitz, 2003).

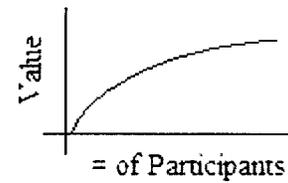
If there is a significant risk of bias formed by having too many participants from a particular organization represented, there are other alternatives to limiting their influence than removing potentially knowledgeable traders. One option is to simply reduce/limit the initial endowments given to any one group of people. Over time, those individuals who are adding the most information to the market will accumulate market wealth which gives them the influence they deserve, regardless of their rank or organizational affiliation. In addition, uninformed participants are known to increase market liquidity, which can lead to improvements in accuracy because knowledgeable traders have more opportunities to share their insights for profit (Chen and Plott, 2002).

In some cases, informed traders may be added to the marketplace without increasing the number of human participants involved. When relevant statistical methods may be used to support the forecasting process, they can be transformed into automated agents that buy and sell claims on their own accord. For example, a particular agent may be designed to execute trades whenever claim prices deviate significantly from the forecast generated by a predefined moving average forecast. This creates a naturally occurring methodology for weighting particularly useful statistical forecasting methods while allowing people to focus on the aggregation of more qualitative information. If a small number of statistical models and/or participants significantly outperform the rest of the market over a period of time, then one should reasonably question whether a prediction market is the most appropriate forecasting solution.

Another important consideration is the degree of trader anonymity to be maintained in the marketplace. When information regarding the ‘owners’ of open orders, actual trades, and/or account balances (performance) is made visible in a marketplace, then some degree of participant anonymity may be desirable. The range of alternatives includes: completely anonymous trading, randomly assigned user IDs, user defined IDs (e.g. – like a person’s Yahoo e-mail ID), and full disclosure. The selection of one alternative over another is highly dependent on the behavioral dynamics involved in the forecasting process. The following table offers some insight into when each of these options may be the preferred choice.

<u>Degree of Anonymity</u>	<u>Reason for Use</u>
Full Disclosure	If an important objective is to encourage communication of information, methods and ideas across organizations.
User Defined IDs	If some participants wish to remain anonymous, while others would like to share their identities with trading partners.
Randomly Assigned IDs	If the anonymity of peoples' names is important, but the ability to 'read' the trading patterns of individuals may provide some traders or observers with meaningful insights.
Completely Anonymous	If there is potential for people not to share information that could lead to negative feedback, and it is likely that the ID could be decoded based on historical trading patterns or performance.

As previously mentioned, prediction markets add value when they are able to aggregate information that is irregularly distributed across a number of traders. Having too few traders can diminish the benefits of prediction markets altogether. A 'thin-market' problem occurs when participants are unlikely to trade claims that they believe have a low probability of resulting in a near-term transaction.



This creates a powerful negative feedback loop that can practically halt trading. At a minimum, thin-markets tend to aggregate information more slowly, are more subject to price manipulation, and exhibit a more pronounced 'favorite-longshot bias'. While it is unknown how many different informed traders are required for a prediction market to function properly, we have seen successful examples where as few as seven people participated. One trick for overcoming a thin-market is to intensify trading by requiring all traders to participate in the same, short timeframe; although this may be especially challenging when participants are in different geographies (Chen, 2004). In addition, there are trading mechanisms (discussed in the following section) that appreciably reduce the adverse affects of a thin-market³.

³ The problems with thin-markets are well documented, yet thick-markets may offer a unique problem of their own. If the most useful information is focused in a few people's hands and they don't have sufficient capital to move the entire market, this can cause 'drift' and reduce forecast accuracy (Chen, 2004).

4.2c Trading Mechanisms

There are two primary mechanisms by which trades occur in prediction markets: continuous double-auctions and ‘market-makers’. Continuous double-auctions require a buyer and seller to agree upon a specified price and quantity for a trade to be made. While this is perfectly acceptable for most active markets, in thin market situations, it can lead to illiquidity and large spreads between the highest ‘bid’ and lowest ‘ask’. The resulting low trading volumes and high price volatility reduce the quality of the information generated by a predictive market.

An automated ‘market-maker’ can be added to a continuous double-auction to reduce the bid-ask spread and increase liquidity by placing buy and sell orders near the last traded price for a claim (Kittlitz, 1999). While this may eliminate the need for ‘dummy’ participants (e.g., HP’s market example), it comes at a price. The ‘market-maker’ assumes positions in various claims and takes on risk for the administrator if the participants’ remuneration is tied to their absolute performance. In addition, the rapid liquidity a ‘market-maker’ provides could facilitate short-term price manipulations (especially before the close of trading) before other traders are able to step-in and restore prices to a point of equilibrium. An automated ‘market-maker’ may also be used to test a market’s equilibrium when trade stalls. This is achieved by creating an order (or complete trade) outside of the current range to see if informed participants drive the market price back to its previous consensus level.

Unfortunately, there is no empirical evidence to demonstrate the accuracy benefits of one trading mechanism over another. As a rule of thumb, continuous double-auctions are the easiest to implement and are perfectly suited to markets with a large number of participants or high trading volume. In the event that a thin-market problem (few traders per claim and / or slow trading over long periods of time) is likely to occur, the addition of an automated ‘market-maker’ is probably worth the risk of higher payouts. When more complex types of trades are necessary (e.g., conditional or grouped trades), the use of Hanson’s market scoring rule or Pennock’s dynamic pari-mutual mechanism should be considered (Hanson, 2004; Pennock, 2004). While these mechanisms are conceptually

brilliant, they are likely too complex for use in most corporate prediction markets. Finally, trading commissions and interest rates for unspent balances should be avoided in all circumstances as they both unnecessarily discourage important trading activity.

4.2d Claim Structure

A claim is a statement about the future that is traded like a stock within a prediction market. The type of claims (or families of claims) that participants can trade may vary widely depending on the type of information being sought from a prediction market.

The **Table of Claim Types** (opposite), based primarily on the work of Wolfers and Zitzewitz (2003), defines the claim structures that are best suited for particular forecasting objectives. Each of these claim types – All-or-Nothing, Index, and Spread – are explored in further detail throughout this section.

All-or-Nothing Claims

All-or-Nothing claims (and sets of them) are the most common type of claims used in practice. Not only are they simple to comprehend, but a set of these claims creates a probability distribution that can yield insights about the expected mean, variance and percentile values for the event in question. For example, the opposite table's example for the 'Probability Distribution' structure would lead to the following conclusions:

- The market believes there is a 50% chance that demand will be less than 100k, and an 85% (50% + 35%) chance that it will be less than 200k;
- The expected value (mean) for the outcome is 115k, assuming that Claims A, B, & C have average values of 50k, 150k, & 250k respectively;
- The variance of the distribution is roughly 5329k, for a standard deviation of 73k.

Table of Claim Types

<u>Objective</u>	<u>Type</u>	<u>Formula</u>	<u>Scenario</u>	<u>Result</u>																
Probability of an Event Occurring	All-or-Nothing	Pays \$1 if event occurs, else \$0	Will the FDA approve Drug ABC for the treatment of Disease XYZ before January 1 st , 200X?	Trading price of \$0.09 represents a 9% chance of the event occurring.																
Probability Distribution	Set of All-or-Nothing Claims	Winning claim pays \$1, others pay \$0	How many units of Product X will be sold in Q2 of 200X? Each claim represents the probability of sales falling within a particular range. (This can be particularly useful for capacity and contract negotiations.)	<table border="1"> <thead> <tr> <th>Claim #</th> <th>Range</th> <th>Price</th> <th>Odds</th> </tr> </thead> <tbody> <tr> <td>A</td> <td>< 100k</td> <td>\$0.50</td> <td>50%</td> </tr> <tr> <td>B</td> <td>100k - 200k</td> <td>\$0.35</td> <td>35%</td> </tr> <tr> <td>C</td> <td>> 200k</td> <td>\$0.15</td> <td>15%</td> </tr> </tbody> </table> <p>Claim C represents the probability of sales exceeding 200k. Its \$0.15 trading price equates to 15% odds, and will pay \$1 if ruled true.</p>	Claim #	Range	Price	Odds	A	< 100k	\$0.50	50%	B	100k - 200k	\$0.35	35%	C	> 200k	\$0.15	15%
Claim #	Range	Price	Odds																	
A	< 100k	\$0.50	50%																	
B	100k - 200k	\$0.35	35%																	
C	> 200k	\$0.15	15%																	
Expected Value (Mean)	Index	Pays \$x, where 'x' is the outcome	What will the new manufacturing process' yield be in its first full quarter of operation? Every penny of a claim's price equals 1% in yield.	Trading price of \$0.79 represents an expected yield of 79%. If actual yield proved to be 81%, the claim would result in a payout of \$0.81 per share.																
Variance Estimate	Index	Pays \$x ² , where 'x' is the outcome	What will the new manufacturing process' yield be in its first full quarter of operation? The claim's price equals the square-root of the expected yield. If actual yield proved to be 81%, the claim would result in a payout of \$0.90 per share. (The square-root of .90 is .81)	Trading price of \$0.64 represents an expected outcome of 80%. (The square-root of .80 is .64) Subtracting the square of the Expected Value Index of 79% (above) from the Variance Index yields the market expectation of the variance. So, \$.64-(\$.79 ²) = .159. Thus, 15.9% is the variance.																
Contingent Expectation (Mean)	Index	If 'Y' occurs, then it pays \$x, where 'x' is the outcome. Else, all bets are off.	If Product X is released to market at a price of \$99, how many units will sell in the first 12-months? Every penny of a claim's price equals 1,000 units of sales.	Claim price of \$0.85 represents an expected outcome of 85,000 units. Actual sales of 82,000 would result in a payout of \$0.82 per share. If the product is not released to market at \$99, then all trades are reversed.																
Median Estimate	Spread	Costs \$1. Pays \$2 if 'price' paid exceeds outcome. Else, pays \$0.	What will the difference in sales be between Product X and Product Y in Q2 of 200x?	Trading 'price' of 1,300 represents a 50% chance that the outcome will be above or below 1,300. (Note: the cost of a share is fixed at \$1; the 'price' reflects the expected median of the outcome, not a dollar cost of the share.)																
Percentiles	Spread	Costs \$x. Pays \$1 if 'price' paid exceeds the outcome's value.	What will average weekly demand for Product X be in 200x? The trading 'price' (in units of demand) reflects a 90% confidence in demand exceeding that value.	A contract costs \$.90 and pays \$1. A trading 'price' of 850 reflects a 90% chance that the outcome will be less than 850. (Adjusting the cost of the contract allows the selection of a particular percentile.)																

Using wide ranges to define claims, however, can lead to a number of issues. First, we have no insight into whether or not the mid-point of a claim's range fairly represents the market's expected value (mean) for that claim.

Secondly, unbounded ranges provide very little information if the claim trades at a high value (probability). In the adjacent chart, forecast users have no way to discern the probability of sales exceeding \$3.0 billion. Therefore it is important when using



Source: NewsFutures.com

claim-sets to select appropriate increments for each claim and to use as many claims as necessary to obtain meaningful information from the upper and lower bounds of the distribution curve. Obviously, this could result in a large number of claims (ranges) that may lead to trading difficulties such as a thin-market problem (i.e., by asking too many questions, one may not receive a single meaningful answer). In this case, it may be important to re-evaluate what information or level of detail is required from the forecast, and whether a set of 'all-or-nothing' claims is the best approach.

The considerations described above also lead one to question what the appropriate increments of claim ranges or trading prices are. This question ties back to the original forecast process mapping exercise (Section 4.2a) that was designed to identify the sensitivities of various decisions to increments of change in the forecast. In addition, the following points may serve as an additional guide for determining the appropriate level of detail built into each claim:

- Trading increments should be small enough to allow decision makers to spot important trends that may influence key decisions or strategy.
- The output of forecasts must be at a level of detail that allows for the direct allocation of resources to potential system constraints. These constraints may include people (by skillset), process lead-times, or equipment.
- Alternatively, some situations may dictate that forecasts provide visibility into the utilization of potentially divestible resources (e.g., excess factory equipment).
- Finally, the ability to forecast at a particular level of detail may be limited by the participants' knowledge and/or interest.

Index Claims

Index claims are particularly useful for obtaining the market's expected value (mean estimate) for a future outcome. Common uses include demand forecasts and project completion dates, but index claims can be used to create almost any kind of point forecast. For example, index claims may be particularly useful in estimating demand curves and establishing prices for new product introductions. (e.g., What will demand be if the price is \$X? What if the price is \$Y?) The primary draw-back of index claims is that they do not offer visibility into the uncertainty of a forecast. Wolfers and Zitzewitz (2003) proposed using an index claim that pays according to the square of an outcome as a tool for predicting forecast variance (see table of claim types above for details). While this method is incredibly clever, it could be a bit confusing for business users because it requires participants to calculate squares and square-roots in order to place trades. If insight into forecast variance is required, a set of all-or-nothing claims will generally suffice.

Spread Claims

Although spread claims are the most commonly used claim type in the context of gambling (e.g., Patriots by 7 points over the Panthers), their applicability to business is quite limited. One is hard-pressed to identify examples of important business decisions that require knowledge of 50-50 odds. However, some break-even analyses for major investments could require assumptions regarding a 90% (or other) confidence interval. For example, a factory expansion analysis might reveal that as long as there is a 90% chance that average weekly demand for Product X will exceed 850 units, then the expansion project should be approved. By creating a claim with a fixed cost of \$0.90 and a payout of \$1.00, the claim's trading 'price' fluctuates to represent an average weekly demand that should be exceeded 90% of the time. If the market reaches equilibrium at a trading 'price' of 600 units, then this sends a clear signal to the decision maker(s) to postpone the expansion. In this example, it is probably not wise to inform participants of the 850 units per week break-point that will determine whether or not the factory expansion should move forward. This number may 'anchor' participant's expectations and introduce bias to their trading.

All Claim Types

A final consideration regarding claim structure is that of price, especially when participants are expected to trade a number of different types of claims. When claim prices are low relative to available funds or alternative investments, one can easily reflect their relative conviction by issuing orders for smaller or larger share quantities. However, if some claims cost exponentially more than others, or a single share represents a significant investment of one's account balance, then it requires the trader to have more confidence in their opinion before entering an order at all. If the goal is to encourage the continuous sharing of information in the form of market prices, then the cost of shares should be normalized to a consistent level, regardless of which claim is being traded. This can be achieved through the use of basic addition, subtraction, multiplication and division to transform potential outcomes into appropriate share prices. For example, if an index claim (designed to estimate the mean outcome of a future event) is used to forecast demand for a product that will sell roughly 1,000 to 3,000 units, then share prices could reflect expected demand divided by 1,000. Therefore, shares priced at \$1.35 reflect a demand expectation of 1,350 units. Note that this implies that forecasting in tens-of-units instead of single units is inconsequential.

4.2e Claim Definition

Perhaps the single most important feature of a prediction market is the claim definition. All claims should be defined in such a way that they can be easily interpreted, unambiguously judged, and clearly associated with the objectives of the forecasting process. No matter how well a market has been designed thus far, a poorly worded claim can lead to participant confusion and market prices that do not reflect their expectations about the future. In addition, claims whose outcomes are left open to subjective interpretation are less likely to invite trading activity (Kittlitz, 1999). Clear rules for judging are very important and should include: judging criteria, a date or event that triggers judging, and who the (preferably impartial) judge(s) will be. For example, if participants are trading claims to forecast unit demand for the next quarter and stock-outs occur, then the exclusive use of sales/shipment data may not be an appropriate way to measure the outcome.

A good claim definition should contain at least the following information:

Field	Example
Trading Symbol	<u>WidgetSales.June2004.200-299</u>
Description	Each share of this claim will pay \$1 if Widget sales for the entire month of June fall between 200 and 299 inclusive. Therefore, a share price of \$0.35 would reflect a 35% chance of this claim being judged 'true'. If the claim is judged 'false', it will pay \$0.
Objective	The purpose of this claim (and its related claims) is to help forecast demand for Widgets in June 2004. The increments of 100 units were chosen to support decisions related to factory scheduling. This forecast is used by: <ul style="list-style-type: none"> • Purchasing for materials procurement decisions • Product Planning for inventory management • Factory Planning for shift scheduling and load balancing • Logistics for distribution requirements planning • Finance for P&L budgeting • Sales & Marketing for demand management and quotas
Background	Historically, demand forecasts have been generated by statistical models maintained by Sales & Marketing. Unfortunately, there are a number of qualitative variables impacting demand that cannot be incorporated into a simple statistical model.
Information Links	Current information related to widget sales can be found through the following sources: <u>Trading Room Chat</u> , <u>CorpDataWarehouse</u> , <u>Field Sales Reports</u>
Trading Begins	March 15 th , 2004 (Inventory planning decisions made April 1 st)
Trading Hours	12-1pm & 5-6pm daily
Trading Ends	June 15 th , 2004 (Final scheduling decision made shortly after)
Judge	Jared Schrieber (schrieber@alum.mit.edu)
Judge's Statement	I will judge this claim on July 1 st , 2004 by querying all orders that contain a Requested Delivery Date (RDD) within the month of June 2004. (See query labeled "June Demand" in CorpDW)
Judge's Disclaimer	If for some reason, this method becomes unviable for determining customer demand for the period (e.g. – system crash, supply constraints prevent customers from requesting June RDDs), and the winning claim cannot be obviously inferred from available information, all trades will be reversed. (i.e. – All accounts will be restored to their original values as if no trading had occurred.) The judge's decision is final.
Related Claims	<u>WidgetSales.June2004.0-99</u> ; <u>WidgetSales.June2004.100-199</u> ; <u>WidgetSales.June2004.200-299</u> ; <u>WidgetSales.June2004.300-399</u> ; <u>WidgetSales.June2004.400-499</u> ; <u>WidgetSales.June2004.500+</u>

Claim Description

The claim description is designed to quickly tell participants exactly what it is they are trading. The basic elements include the trading symbol, an unambiguous statement about the future, and the claim's background and purpose (objective). Before a claim description is made final, several market participants should have the opportunity to seek clarity regarding any potentially ambiguous statements or outcomes. In addition, the description must clearly relate the actual cost, trading prices, and payouts of the claim to all potential outcomes. Remember that spread claims have fixed costs and payouts, but the trading 'price' is not measured in currency. For example, a spread claim could have a set cost of \$4 and payout \$5 when the 'price' paid (e.g, 1100 units) is less than the actual outcome (e.g., 1250 units). In this case, the trading 'price' fluctuates to represent a unit volume that the market believes there is an 80% chance ($\$4/\$5 = 80\%$) of actual unit volume exceeding. Had someone purchased a claim for a 'price' of 1300 units, and the actual outcome was 1250 units, they would lose all \$4 of their investment.

Trading Horizon

As mentioned previously, condensed trading hours can be used to influence the amount of time employees spend trading, and even the intensity of trade. Perhaps even more important is the time horizon over which trade is to occur. Since corporate prediction markets exist solely to support decision making processes, claims should only be open for trade when decisions can be tied to their market prices. For example, trading should begin shortly before price information can be put to good use, any sooner would be a wasteful distraction. In addition, trading should only begin after some information exists by which participants can make informed judgments. Trading should continue until the final decision based upon the market-based forecast is made, or the event in question occurs.

When predicting a longer-term event, market participants may lack an incentive to actively trade in the early stages of the market, especially if there are nearer-term alternatives to trade. One way to overcome this issue is to provide only a short, intense trading window just prior to an important decision based on the market's price. This is a

similar approach to that taken in the Chan et al (2002) market research experiments which lasted less than an hour and awarded participants with the greatest increases in their account value during the session. A similar alternative is to create a separate market with new accounts for participants to use exclusively for the trade of longer-term claims. If trading is to continue over a long period of time (in support of continuous decisions), then a scoring/payout structure may need to be implemented that rewards early, correct predictions more than later, correct predictions. The down-side to this approach is that a trader may not update their portfolio to reflect new information because the expected value (payout * probability) of their current position(s) is higher than it would be if they acted on their new beliefs.

Judgment Criteria

Who will judge the claim's outcome? Questions about how a claim will be judged under various circumstances may arise at any time during trading and should be addressed by this judge. In addition, participants should have confidence in this person's (or panel's) ability to accurately judge the claim in an unbiased manner. This is particularly important when there is the potential for discrepancies between what actually happens and what gets measured or reported (Kittlitz, 2004). For example, time-lags in the complete reporting of events could impact how claims are judged. If sales numbers for the current quarter are not fully reported for several months, this may impact when and how related claims should be judged.

In addition to having rules for the judgment of claims, procedures are also required for transforming market prices into official forecasts. Although these procedures do not change how outcomes are judged or the awarding of points among traders, they are still important. Since low-volume claims are likely to experience greater volatility towards the close of a market, the median or volume-weighted average price over the last 20-50% of trades may be more appropriate to use as a forecast (Chan et al, 2002). HP's market used the volume-averaged transaction price over the last 50% of trades, since laboratory experiments had demonstrated that trading tended to reach equilibrium towards the end of trading. Although the closing price of a claim might be slightly skewed as the result of a

short-term ‘bubble’, it is quite likely that the late change in price reflects the incorporation of new information. In higher-volume markets, the closing price of a stock (or the average of the bid/ask spread) should be sufficient for determining the market’s prediction. Finally, when claims are defined by an outcome range (e.g., Sales will fall between 200 and 300), they should be assigned an average value according to a predefined method. These methods include using the range’s mid-point, or the application of a best-fit curve to the entire probability distribution created by a set of ‘all-or-nothing’ claims.

Exception Handling

Claims are a kind of contract that payout if a certain condition is met. Since most authors of claim descriptions are not contract lawyers, one should expect situations where the final judgment of a claim is not a black-and-white decision (Kittlitz, 2004). In the Siemens project management example mentioned earlier, participants were trading all-or-nothing claims based on whether or not the project would be completed ‘on-time’. When the project’s customer pushed out the deadline by a significant margin and management reshuffled resources accordingly, the holders of ‘on-time’ shares were awarded the payout. This frustrated many of the participants who argued that all funds should have either been refunded back to their original levels or their most recent portfolio values before the announcement. Had participants known up-front how such an exception was to be judged, they would have adjusted their trading strategies accordingly and not fought the ruling.

The creator of a claim and its judge should proactively consider potential disruptions/exceptions to claim trading and publicly state how these situations will be handled. The ‘Judge’s Disclaimer’ should inform traders how unexpected or potentially ambiguous situations will be handled. Will the judge try to follow the intent of the claim, the literal wording of the claim, or some other interpretation of the claim? If the claim cannot be judged due to an unexpected or unknown outcome, will the judge simply reverse all prior trades, rule the outcome as equal to the forecast derived by market prices, or use their own best judgment to estimate the actual outcome? Finally, additional

'escape clauses' can be added at any time as circumstances change and questions arise - as long as these clauses are communicated to all participants.

4.2f Incentives

All markets fail when traders are insufficiently motivated by the potential for profit, risk mitigation, or intrinsic reward, thus the use of incentives can be critical to a prediction market's success. Although many informed individuals are motivated to trade simply out of an intrinsic enjoyment of trading based on their knowledge, remuneration may be required to extract timely and accurate information from others. The importance of the forecasts, as well as the amount of time and effort expected of participants should be the focal point of such compensation decisions. In addition, rewards may be given at the conclusion of each claim, a set of related claims, and/or at set time intervals to adjust the influence that incentives have on trading (forecasting) effort.

The use of predetermined awards based solely on the relative rank of participants creates a simple, effective, and fixed-cost incentive structure. In this case, the participants with the largest increases in their account balances are each given a prize. A drawback to this method is that it changes peoples risk profiles and can theoretically lead to irrational trades by less informed participants who, with nothing to lose, know that they have to be lucky to win (Chen, 2004). An alternative approach is to provide cash payouts proportional to the participants' account balances. (Let us assume that the market is fully subsidized by the company and does not require the cash contributions of participants.) As long as the scoring rules for all claims ensure that the marketplace is a zero-sum game, then the administrative costs are fixed and each participant is encouraged to trade rationally. (In the event that an automated 'market-maker' is used, it may lose money, and thus increase the risk of larger payouts to human participants.) However, some participants may choose not to trade at all, in favor of simply collecting their initial endowment at no risk. This can be mitigated if all participants are required to invest at least 50% of their endowment in order to receive a payout.

Although it is a widely-held view that real-money prediction markets are more accurate than their play-money equivalents, there is little evidence to support this claim. Recent research into this question yielded no significant differences between the accuracy of the two methods (Wolfers, 2003). It seems plausible that awarding prizes to a fixed number of high performing participants could result in a more pronounced favorite-longshot bias caused by low-balance participants betting on a low-probability jackpot to put them in contention. However, real-money markets have the potential to create procedural, technical, legal and ethical dilemmas that can swiftly increase the costs of this approach. In particular, real-money prediction markets have been found to encounter significant challenges from existing compensation systems and run a higher risk of moral hazard situations (Chen, 2004 & Gebert, 2004). In the end, the company must create an incentive system that motivates knowledgeable individuals to continuously share their information in a timely manner. This may require prizes or payouts that will serve as a meaningful bonus to employees responsible for generating accurate forecasts.

4.2g Account Balances

The initial endowment and ongoing allowance allocated to traders are important aspects of stimulating trade and potentially weighting the opinions of some individuals or groups over others. In statistical forecasting where different forecasting methods are combined, it is recommended that equal weights be given to all methods until sufficient historical information exists to warrant the weighting of particular methods over others (Chase, 2000). Applying the same logic to prediction markets implies that all participants should receive equal funding when their account is created, and that their wealth accumulation over time will effectively weight their opinions relative to others when predicting a particular type of event.

To initiate trade, it is recommended that all participants be given allotments of shares in each claim and a sum of currency equivalent in value (i.e., \$100 in currency and \$100 in shares). This makes it easy for traders to both buy and sell shares from the moment trade begins. If the market allows traders to sell-short, then a pre-established rule for

determining the amount of margin placed on hold may be required. In most cases, freezing an amount equal to the total risk of the short-sale should suffice.

If the same participants will continue to trade new, related contracts over a period of time, then an allowance system may be in order. Small or no allowances serve as a mechanism for giving greater weight to those that have a history of accurate predictions, while large allowances have the effect of creating a more-level playing field at the beginning of trade for each new claim. Limited allowances have also been argued to drive traders to make better decisions than they would under a system that provides them with a continuous source of funds. Allowances may be given for each new claim, every month/quarter, or when a participant's balance falls below a minimum threshold. Completely unrelated claims should be set up in different markets, with players receiving a separate account of cash and shares to manage exclusively within that market.

If a 'real-money' market is used, then one must deal with the issue of balance withdrawals. Performing payouts that reduce the participants' accumulated market wealth will also decrease the influence of the best performers over the market. This, in turn, can lead to degradations in the market's overall forecasting performance.

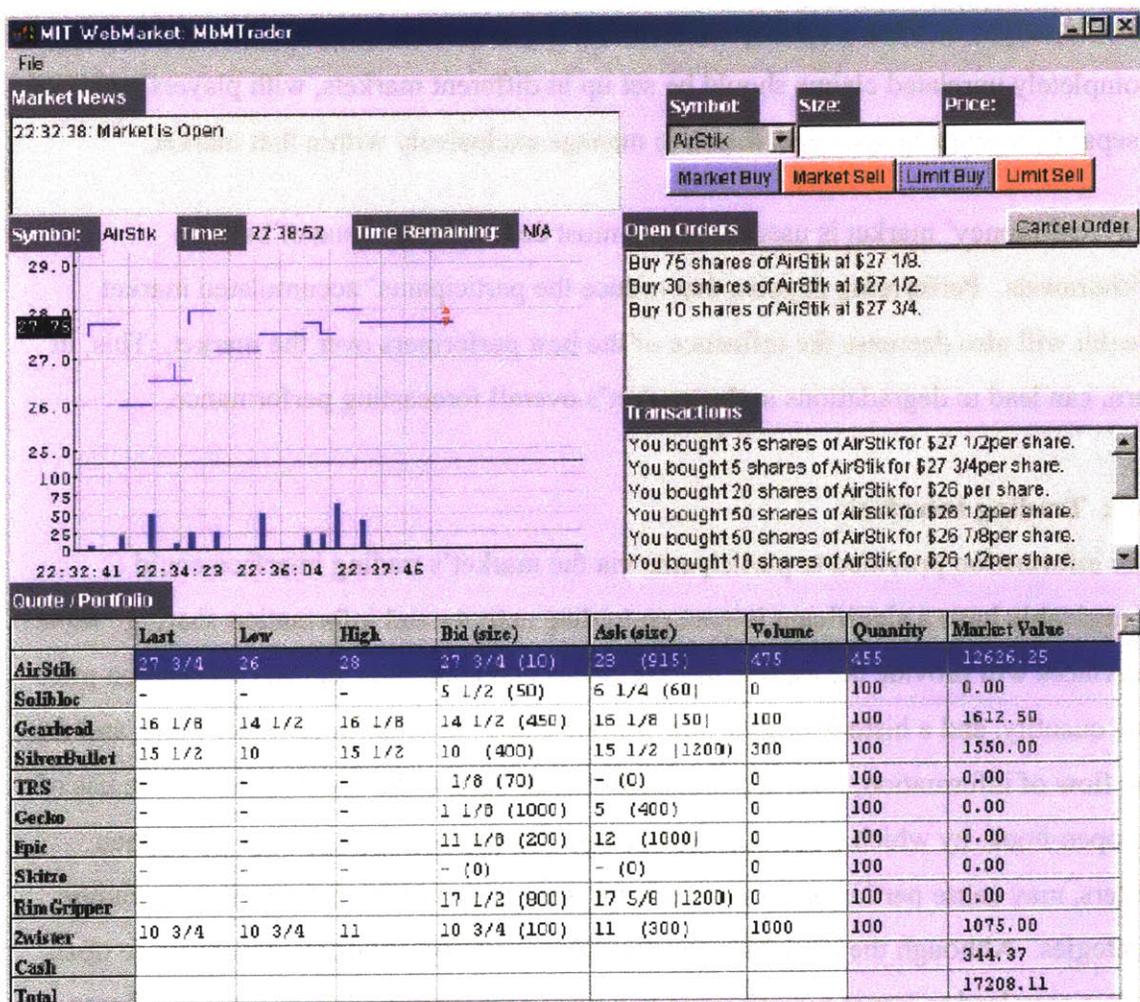
4.2h Trading Interface

The information provided to participants via the market's trading interface could conceivably have a significant impact on trading activity and information sharing. Most interfaces will provide the current bid/ask prices and quantities, the last transaction price and quantity, and a historical price and volume chart. In order to facilitate trade, and thus the flow of information, it is recommended that an open order book be used. The use of an open-book, by which all participants can see the prices and quantities of all open orders, may cause participants to re-evaluate their expectations and adjust their trading strategies. Although the use of limit orders adds liquidity, greater transparency to open orders may lead to larger bid/ask spreads (unless a 'market maker' is used) and fewer information revealing trades (O'Hara, 2001). Other information that may be incorporated into the trading interface include up-to-date data about the claim outcome in question

(e.g., current sales numbers for a sales forecast claim), or a chat-room by which participants can discuss the claim and share their insights.

It is probably wise to prevent each participant's balance (or ranking) from being seen by others, because it may encourage irrational trading decisions by those that are not currently in contention to win prizes. Additionally, someone that is clearly in the lead may choose to discontinue trading in order to avoid the risk of losing their assured prize. This would effectively eliminate the influence of the single most accurate participant.

Sample Trading Interface (Chan et al, 2002)



5. Applications Beyond Forecasting

The application of internal markets to business is not new. Most large companies already operate internal markets that enable business groups to buy and sell services from one-another. This section attempts to highlight opportunities for innovative uses of internal markets beyond what has been discussed thus far.

5.1 Resource Allocation

Markets are adept at the efficient allocation of scarce resources to those who are most willing to pay. For example, the highly resource-constrained IT departments of many companies charge internal business units for their services. This naturally applies a value to the services provided by the IT organization while reducing the politics involved with negotiating which projects the IT group takes-on. If the IT group's rates get too expensive relative to the cost of external consultants, then either the IT group should be given more resources, or its business partners should be allowed to contract with the lower-cost consultants.

Additional constrained resources that could be more effectively allocated through the use of markets may include budgets, specially-skilled employees, shifts of a production facility, or warehouse space. It is not difficult to imagine a scenario by which a division manager seeking to fund a major cost-saving project is willing to sell \$60k of a future quarter's budget for a \$50k increase in an upcoming quarter's budget figure. Under such a scenario, the parent company's quarterly budgets would remain constant, but its expenditures would likely be directed to more fruitful activities. When several cross-organizational projects compete for a common specially-skilled employee, the employee is generally assigned to the project with the highest net present value (NPV), regardless of how critical that employee is to the success of the program. If each project were given a 'human resources budget' from senior management upon approval, then project managers could bid for key resources according to the relative value that each person adds to the project. 'Pet projects' of little value could no longer compete with highly valued projects for key resources. This becomes an effective way to allocate peoples'

time, while objectively identifying those employees and skill-sets that are adding value in the company and ensuring that they are put to good use.

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.

While some might argue that Activity Based Costing combined with cost allocations to the resource consuming group already achieves this very purpose, there are some important differences. Activity Based Costing can be an expensive and crude method for 'valuing' the resources consumed by a set of activities. These costs are calculated after resource allocation decisions have already been made, and do not account for the opportunity costs of using those resources for other purposes. As a result, there is a tremendous delay involved with using Activity Based Costs to make important trade-off decisions.

The examples above highlight situations where resources are to be allocated based on the expenditures of previously distributed budgets (e.g., cash or a 'human resources budget'). However, some of the most important decisions companies make are related to the initial

allocation of resources to projects that would otherwise have no budget of their own. How should these early stage projects get their initial funding? (See *New Product Development Decisions* sidebar at the end of this section for a case study on how prediction markets can be applied in this context.) While Discounted Cash Flow (DCF or NPV) and Real-Option approaches offer structured methodologies for making such decisions, they require valid assumptions regarding the probability of events occurring and a complete representation of the complex relationships among variables. In highly ambiguous situations, these models can take weeks or months to develop, populate with assumption data and still achieve a low level of stakeholder buy-in. Perhaps this is one reason why venture capitalists developed the simplistic “VC Method” to value start-up companies in search of capital. Under this method, venture capitalists estimate the future value of a company and discount that value to today, using a required rate of return. But, how does one determine the value of a company without a track-record? Prediction markets may be the answer.

5.2 Risk Hedging

While internal markets cannot be used to change a company’s risk profile, many public markets exist solely for the purpose of allowing participants to hedge their risk against fluctuations in interest rates, currency exchange rates, weather and more. Several researchers have proposed that public prediction markets focused on politics, world events and scientific developments may provide new venues for risk hedging. Of course, such markets would have to be legalized (they currently fall under gaming law in the US) and reach significant levels of liquidity before they could become a viable option for hedging risk. The sports betting market in Nevada is a perfect example of a prediction market with sufficient liquidity for corporate hedging. Sports merchandisers and retail businesses whose profitability is tied to the success of particular teams would be wise to use this tool to stabilize earnings and reduce their cost of capital. For example, if an apparel manufacturer has exclusive contracts with a few major teams, and sales of its apparel are driven by each team’s performance, then the company could lock-in minimum levels of profit by betting against those teams. Failure to hedge risks associated with team performances is one reason why smaller licensed apparel

manufacturers continue to go bankrupt (Anonymous, 2004). This helps explain the massive consolidation taking place in the licensed apparel industry, allowing major manufacturers to avoid reliance upon the performance of just a few licensed teams.

New Product Development Decisions

FashionWare is a leading provider of fashion footwear worldwide. Only 10% of the roughly 75 products it releases to the market each year are even moderately 'successful', while the remaining 90% are clear 'busts'. The Product Marketing Team (PMT) responsible for creating product concepts and driving them through to market believes that this is just a part of doing business in the fashion industry.

Their counterparts in the regional sales offices and at key customer accounts offer a different opinion. "I don't even know about new products until they are brought to market. Although many of these products clearly have zero chance at selling in my region (Germany), I am forced to push the product on my key accounts." says one Regional Sales Manager. FashionWare's design and development teams share similar concerns: "Marketing runs the show, and the only thing we're allowed to do is act as their servants. There have been hundreds of products that everybody knew were going to be duds, but the PMT just pushed them forward anyway."

The Product Marketing Team claims that if they had to cater to the wishes of these other groups, no product would ever ship. "This is an emotional business and everyone has their own opinion. We don't have time to collect and incorporate everyone's feedback. If we start catering to Germany, then France and Japan will want to be next. It's a no-win situation." says the PMT manager.

A prediction market where each product is traded like a stock from its initial conception through to its market release, combined with a Stage-Gate® type process, may be the answer. Every month, product ideas submitted by anyone go through an 'Idea Screen' trading session. Participant traders from marketing, design, development, sales and even key accounts spend just one-hour evaluating, buying and selling shares of these product concepts. Those products that exceed a minimum threshold in market value are passed on to the design team. Next month (or quarter), once product sketches have been designed, the same participants engage in another trading session to determine if a prototype should be developed. This process continues through the final decision to launch new products in the market.

The benefits to FashionWare are numerous. The company obtains more product ideas from a more diverse population of employees that have market-based incentives for submitting good ideas. These ideas are quickly and fairly screened by a wider-range of people who are more representative of what end consumers want. Extreme opinions held by different stakeholders are resolved objectively and anonymously so that emotions and politics do not bias decisions and hinder working relations. Perhaps most importantly, while FashionWare will release just slightly fewer shoes to market each year, the number of successful product launches is expected to increase dramatically.

6. 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:

- Can someone completely unfamiliar with equity markets learn how to rationally express their opinions through trading in a short period of time?
- What is the minimum number of informed participants required, utilizing various trading mechanisms, to achieve accurate results? Is there a minimum ratio of informed to uninformed participants required to achieve accurate results?
- Which trading mechanisms perform best, and under what circumstances?
- Are there experiments that can demonstrate the kinds of circumstances that cause prediction markets to outperform survey-based methods (and vice-versa)?
- Since it is known that the accuracy of survey-based methods decreases when participants hold significantly different opinions, how much better do market-based methods perform in particularly ambiguous forecasting challenges?

With so many unanswered questions, it is clear that prediction markets are not yet ready for wide-spread adoption in a corporate context. 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. ■

7. Bibliography

- Anonymous, (2004) Former founder and president of a major licensed apparel manufacturer. [Interview] January 5th, 2004.
- Berg, J., Rietz, T., (2003) "Prediction Markets as Decision Support Systems", *Information Systems Frontiers*, 5(1), pp79-93, 2003.
- Boulier, B., Stekler, H.O. (2001), "Predicting the outcomes of National Football League games," *The International Journal of Forecasting*, 19, pp 257-270, 2001.
- CFO Research Services, (2003) "CFOs and the Supply Chain," www.cfo-research.com, 2003.
- Chan, N., Dahan, E., Kim, A., Lo, A., Poggio, T., (2002) "Securities Trading of Concepts (STOC)", Jan 23rd, 2002. [On-line] http://www.anderson.ucla.edu/faculty/ely.dahan/content/chan_dahan_lopoggio.pdf May 1, 2004.
- 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.
- Feder, B., (2002) "To Learn What People Want, Trade 'Idea Stocks' ", The New York Times Web version, Jan 27th, 2002.
- Fisher, M., Hammond, J., Obermeyer, W., Raman, A., (1994) "Making Supply Meet Demand in an Uncertain World", *Harvard Business Review*, May-June 1994.
- Gandar, J., Dare, W., Brown, C., Zuber, R. (1998), "Informed Traders and Price Variations in the Betting Market for Professional Basketball Games," *The Journal of Finance*, Volume LIII, NO. 1, p 385, 1998.
- Gebert, C., (2004) CEO of Incentive Markets. [Interview] April 20th, 2004.
- Hanson, R. (1996), "Idea Futures," [On-line] <http://hanson.gmu.edu/ideafutures.html>, May 1st, 2004.
- Hanson, R. (1999), "Decision Markets," *IEEE Intelligent Systems*, May/June 1999, pp 16-19.
- Hanson, R., (2003) "Combinatorial Information Market Design", *Information Systems Frontiers*, 5, 1; ABI/INFORM Global, pp 107-119, 2003.

- Hanson, R. (2004), "Issues in Information Market Design," [On-line] <http://hanson.gmu.edu/infomkts.html>, March 15th, 2004.
- Hayek, F. A., (1945) "The use of knowledge in society.", *American Economic Review*, 35, pp. 519-530, 1945.
- Kahn, K., (2003) "How to measure the impact of a forecast error on an enterprise?", *The Journal of Business Forecasting Methods & Systems*, Spring 2003; 22, 1, pp. 21-25, 2003.
- Kittlitz, K., (1999) "Experiences with the Foresight Exchange", Extropy Online, 1999, [On-line] <http://www.yeti.com/articles/fx.htm>, December 7th, 2004.
- Kittlitz, K., (2004) Founder and administrator of Foresight Exchange. [Interview] April 13th, 2004.
- Kulkarni, V. (2003), "Using weather futures as weather forecasts," *Chicago Mercantile Exchange*, [On-line] <http://www.cme.com/files/EPSept03.pdf>, February 16th, 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.
- Ortner, G., (1998) "Forecasting Markets – An Industrial Application," *mimeo*, Technical University of Vienna, 1998.
- Pennock, D., Lawrence, S., Nielsen, F., Giles, C. (2001), "Extracting Collective Probabilistic Forecasts from Web Games," *Proceedings of the 7th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-2001)*, August 2001. [On-line] <http://artificialmarkets.com/am/pennock-kdd-2001-games/>, December 13th, 2003.
- Pomerantz, D. (2003), "Can This Man Save the World?," *Forbes.com* [On-line] <http://www.forbes.com/forbes/2003/0811/054.html>, May 1st, 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.
- Winkler, R. (1971), "Probabilistic Prediction: Some Experimental Results," *Journal of the American Statistical Association*, Vol. 66 No. 336 (Dec., 1971), pp675-685.
- Wolfers, J., (2003) "Prediction Markets: Does Money Matter?," *draft submitted to Electronic Markets – the International Journal of Electronic Commerce & Business Media*, 2003 [Personal e-mail] February 26th, 2004.
- Wolfers, J., Zitzewitz, E., (2003) "Prediction Markets", November 19th, 2003. [On-line] <http://faculty-gsb.stanford.edu/wolfers/Papers/Predictionmarkets.pdf>, February 2nd, 2004.