

Forecasting Consumer Products Using Prediction Markets

By

Kai Trepte

B.A. Economics & Computer Science, Macalester College, St. Paul Minnesota, 1990

And

Rajaram Narayanaswamy

B.Tech., Electronics and Communication Engineering, National Institute of Technology,

Kurukshetra, India, 1997

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 2009

© 2009

Kai Trepte and Rajaram Narayanaswamy

All rights reserved.

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

Signature of Authors
Master of Engineering in Logistics Program, Engineering Systems Division
May 8, 2009

Certified by
Dr. Larry Lapide
Director, Demand Management
MIT Center For Transportation & Logistics
Thesis Supervisor

Certified by
Prof. Yossi Sheffi
Professor, Engineering Systems Division
Professor, Civil and Environmental Engineering Department
Director, Center for Transportation and Logistics
Director, Engineering Systems Division

Abstract

Prediction Markets hold the promise of improving the forecasting process. Research has shown that Prediction Markets can develop more accurate forecasts than polls or experts. Our research concentrated on analyzing Prediction Markets for business decision-making. We configured a Prediction Market to gather primary data, sent out surveys to gauge participant views and conducted in-depth interviews to explain trader behavior. Our research was conducted with 169 employees from General Mills who participated in Prediction Markets that lasted from two to ten weeks. Our research indicates that short term forecasting Prediction Markets are no more accurate than conventional forecasting methods. It also presents and addresses three interesting contradictions. First, the Sales Organization won the majority of the Prediction Markets, yet the overall performance of Sales as a group was worse than that of other groups. Second, Prediction Markets were able to gain access to more information than General Mills' current process, yet the impact on forecast accuracy was not significant. Third, with a MAPE of 11% for promotional Prediction Markets, it would seem that promotional demand was well understood up-front, yet when we dissected the promotional forecasts we discovered that participants changed their minds over time degrading overall forecast accuracy. We believe that we have extended the current body of work on Prediction Markets in ways that will increase the utilization in business environments.

Acknowledgements

We would like to thank General Mills for their unwavering support of this project. Each time we asked to expand the scope they were there to ensure our success. We would also like to thank Tim Weidenhaft for giving us the freedom to explore Prediction Markets while keeping us from wandering too far off the path and turning the research into a summary of many unrelated areas. We would like to acknowledge Kathy Rodriguez for her tireless support, and ability to navigate the General Mills organization and bring the team together, which made our thesis an experience to treasure. We would finally like to thank the 169 market participants; without them we would not have had a study.

We would also like to thank Dr. Larry Lapide for his ability to herd cats. Larry's guidance was instrumental in helping focus our thoughts, validate our research and make sure we were paying attention to the details. Larry without your guidance, this project would not have gotten done.

We would also like to thank Dr. Bill Haas, his focus on quality and clarity drove us to create a document that we hope will live up to his standards.

We would also like to thank our wives, Anna and Soumya for their tireless reading of our thesis and constant and untiring support to our demanding schedule. Without them we would have been lost.

We would like to thank Dave Perry and Ken Kittlitz and the entire ConsensusPoint organization. The late nights they spent and wisdom they imparted to us enriched our understanding and analysis.

Finally, we would like to thank Dr. Chris Caplice, MIT and the Center for Transportation and Logistics for providing us with this unique opportunity to work on an exciting topic for research.

Table of Contents

LIST OF FIGURES.....	6
LIST OF TABLES	7
LIST OF EQUATIONS	8
1 INTRODUCTION TO PREDICTION MARKETS.....	9
2 LITERATURE REVIEW	10
3 METHODS FOR CONFIGURING PREDICTION MARKETS	17
3.1 PREDICTION MARKET PILOT.....	17
3.1.1 <i>Contract Types within Prediction Markets</i>	18
3.1.2 <i>Selection of Prediction Market Pilot Participants</i>	22
3.1.3 <i>Testing for Prediction Market Manipulation</i>	24
3.2 PREDICTION MARKET CONFIGURATION	25
3.2.1 <i>Participant Selection</i>	25
3.2.2 <i>Prediction Markets Configured at General Mills</i>	27
3.2.3 <i>Prediction Market Contracts</i>	29
3.2.4 <i>Setting Stock Ranges in a Prediction Market</i>	30
3.2.5 <i>Setting the Number of Stocks in a Prediction Market</i>	34
3.2.6 <i>Determining the Winner of a Prediction Market</i>	36
3.2.7 <i>Setting Stock Prices in a Prediction Market</i>	42
3.2.8 <i>Setting Investment Caps in a Prediction Market</i>	45
3.2.9 <i>Setting Trading Hours in a Prediction Market</i>	46
3.2.10 <i>Providing Training for Prediction Market Participants</i>	47
4 METHODS FOR ANALYZING PREDICTION MARKET DATA	48
4.1 MAPPING CURRENT PLANNING PROCESS	49
4.1.1 <i>Operations Forecasting Process</i>	49
4.1.2 <i>Sales Forecasting Process</i>	50
4.2 <i>Quantitative Analysis of Data</i>	51
4.2.1 <i>Prediction Market Data Description</i>	52
4.2.2 <i>Prediction Market Price Interpretation</i>	55
4.2.3 <i>Definition of MAPE for Measuring Error</i>	56
4.2.4 <i>Prediction Market Groups for Analysis</i>	57
4.2.5 <i>Analyzing Overall Results</i>	58
4.2.6 <i>Analyzing Forecast Accuracy</i>	58
4.2.7 <i>Analyzing Market Activity</i>	59
4.2.8 <i>Analyzing Trader Behavior</i>	64
4.3 <i>QUALITATIVE ANALYSIS</i>	66
4.3.1 <i>Analyzing Participants Using Surveys</i>	66
4.3.2 <i>Analyzing Participants Using Interviews</i>	67
5 FINDINGS	68
5.1 <i>ANALYZING OVERALL RESULTS</i>	68
5.2 <i>ANALYZING FORECAST ACCURACY</i>	70
5.2.1 <i>Market Type Forecast Accuracy</i>	70
5.2.2 <i>Forecast Type Forecast Accuracy</i>	72
5.2.3 <i>Functional Forecast Accuracy</i>	75

- 5.2.4 Location Forecast Accuracy..... 77
- 5.2.5 Forecast Type Transaction Forecast Accuracy 78
- 5.2.6 March Forecast Accuracy..... 79
- 5.3 ANALYZING MARKET ACTIVITY 80
 - 5.3.1 Flow for Forecast Type Prediction Market Groups..... 80
 - 5.3.2 Flow for Market Type Prediction Market Groups..... 81
 - 5.3.3 Flow for March Prediction Markets 82
 - 5.3.4 Convergence and Confidence in Volume Prediction Markets 82
 - 5.3.5 Convergence and Confidence in Promotional Prediction Markets 84
 - 5.3.6 Convergence and Confidence in Category Prediction Markets 86
 - 5.3.7 Confidence and Convergence of March Markets 88
- 5.4 Analyzing Trader Behavior 90
 - 5.4.1 Trader Behavior in Volume Prediction Markets 90
 - 5.4.2 Trader Behavior in Promotional Prediction Markets..... 91
 - 5.4.3 Trader Behavior in Category Prediction Markets 92
 - 5.4.4 Trader Behavior in March Markets 92
- 5.5 ANALYZING PARTICIPANTS USING SURVEYS 93
 - 5.5.1 Analysis of First Market Survey 93
 - 5.5.2 Analysis of Final Market Survey 95
- 5.6 ANALYZING PARTICIPANTS USING INTERVIEWS 96
- 6 PREDICTION MARKET CONCLUSIONS 98**
 - 6.1 PREDICTION MARKET FORECAST ACCURACY 99
 - 6.2 PREDICTION MARKET INFORMATION AGGREGATION 100
 - 6.3 PARTICIPANT SELF SELECTION 100
 - 6.4 PREDICTION MARKET MANIPULATION 100
 - 6.5 PARTICIPANT BEHAVIOR..... 101
 - 6.6 NEXT STEPS FOR GENERAL MILLS 102
 - 6.7 ADDITIONAL RESEARCH 103
- REFERENCE LIST 105**

List of Figures

FIGURE 3.1.1 - SPREAD CONTRACT	19
FIGURE 3.1.2 - INDEX CONTRACT	20
FIGURE 3.1.3 - WINNER-TAKE-ALL CONTRACTS.....	22
FIGURE 3.2.1 - RANGE SELECTION APPROACH	32
FIGURE 3.2.2 – EXPANSION SELECTION APPROACH.....	33
FIGURE 3.2.3 - JAN PREDICTION MARKET CONFIGURATION.....	35
FIGURE 3.2.4 - BUY AND HOLD STRATEGY	36
FIGURE 3.2.5 - BUY AND SELL STRATEGY	39
FIGURE 3.2.6 - SHORT SELLING STRATEGY.....	40
FIGURE 3.2.7 - NORMALIZED PRICING APPROACH	44
FIGURE 4.1.1 - FORECASTING PROCESS AT GENERAL MILLS.....	50
FIGURE 4.1.2 - ORGANIZATIONAL STRUCTURE FOR GENERAL MILLS SALES REGIONS	51
FIGURE 4.2.1 – PREDICTION MARKET CONFIDENCE	61
FIGURE 4.2.2 - GRAPH OF CONVERGENCE COEFFICIENT OF VARIATION AND MAPE	63
FIGURE 5.3.1 CONFIDENCE AND CONVERGENCE FOR VOLUME PREDICTION MARKETS	83
FIGURE 5.3.2 - CONFIDENCE AND CONVERGENCE IN PROMOTIONAL PREDICTION MARKETS	85
FIGURE 5.3.3 - CONFIDENCE AND CONVERGENCE IN CATEGORY PREDICTION MARKETS	87
FIGURE 5.4.1 - TRADER BEHAVIOR IN VOLUME PREDICTION MARKETS	90
FIGURE 5.4.2 - TRADER BEHAVIOR IN PROMOTIONAL PREDICTION MARKETS.....	91
FIGURE 5.4.3 - TRADER BEHAVIOR IN CATEGORY PREDICTION MARKETS.....	92
FIGURE 5.4.4 - TRADER BEHAVIOR IN MARCH PREDICTION MARKETS	93

List of Tables

TABLE 3.2.1 - PREDICTION MARKET PARTICIPATION	26
TABLE 3.2.2 - SUMMARY OF PARTICIPANTS BY JOB FUNCTION	26
TABLE 3.2.3 - PREDICTION MARKETS CONFIGURED AT GENERAL MILLS	28
TABLE 3.2.4 - BUY AND HOLD STRATEGY OUTCOMES.....	38
TABLE 3.2.5 - BUY AND SELL STRATEGY	40
TABLE 3.2.6 - OUTCOMES FOR VALUES OUTSIDE THE 48,100,000 AND 49,200,000 DELIVERY RANGE	41
TABLE 4.2.1 - TABLE OF DATA CATEGORIES USED IN ANALYSIS	52
TABLE 4.2.2 - DESCRIPTION AND SAMPLE OF TRADING HISTORY.....	54
TABLE 4.2.3 - DESCRIPTION AND SAMPLE OF DAILY STOCK SUMMARY.....	55
TABLE 4.2.4 - ILLUSTRATION TO DERIVE PREDICTION MARKET FORECAST FROM PRICES.....	56
TABLE 4.2.5 - MAPE CALCULATION ILLUSTRATION.....	57
TABLE 4.2.6 - DESCRIPTION AND EXAMPLE OF OVERALL RESULTS FROM PREDICTION MARKETS	58
TABLE 4.2.7 - WEEKLY MAPE DATA FOR PREDICTION MARKET “HOW MANY CASES OF PRODUCT A WILL BE SOLD IN MARCH 2009”	59
TABLE 4.2.8 – TRANSACTION FLOW FOR PREDICTION MARKETS	60
TABLE 4.2.9 - TABLE OF COEFFICIENT OF VARIATION AND MAPE.....	63
TABLE 4.2.10 - JOB FUNCTION FORECAST SUMMARY.....	66
TABLE 4.3.1 - CATEGORIES OF PARTICIPANTS FOR INTERVIEWS	68
TABLE 5.1.1 – OVERALL RESULTS FROM PREDICTION MARKETS.....	69
TABLE 5.2.1 – MARKET TYPE DEFINITIONS.....	71
TABLE 5.2.2 – MARKET TYPE MAPE COMPARISON.....	71
TABLE 5.2.3 - STANDARD DEVIATION OF MARKET TYPE MAPE	72
TABLE 5.2.4 – FORECAST TYPE PREDICTION MARKETS	73
TABLE 5.2.5 – FORECAST TYPE PREDICTION MARKET MAPE	74
TABLE 5.2.6 – FORECAST TYPE STANDARD DEVIATION OF MAPE	75
TABLE 5.2.7 – JOB FUNCTION PREDICTION MARKET MAPE.....	76
TABLE 5.2.8 – LOCATION PREDICTION MARKET MAPE.....	78
TABLE 5.2.9 – FORECAST TYPE PREDICTION MARKET MAPE	79
TABLE 5.2.10 – MARCH PREDICTION MARKET MAPE	79
TABLE 5.3.1 - FLOW FOR FORECAST TYPE PREDICTION MARKETS	81
TABLE 5.3.2 - FLOW FOR MARKET TYPE PREDICTION MARKETS	82
TABLE 5.3.3 - TABLE FLOW FOR MARCH PREDICTION MARKETS.....	82
TABLE 5.3.4 - CONFIDENCE AND CONVERGENCE IN MARCH PREDICTION MARKETS.....	89

List of Equations

EQUATION 4.2.1 - EQUATION FOR DERIVING PREDICTION MARKET FORECASTS55
EQUATION 4.2.2 - MAPE CALCULATION EQUATION.....57
EQUATION 4.2.3 - EQUATION FOR COMPUTING FLOW60
EQUATION 4.2.4 - EQUATION OF COV FOR PREDICTION MARKET FORECAST.....62

1 Introduction to Prediction Markets

In the past, two of the most popular forecasting approaches used for improving forecast accuracy have been statistical and collaborative forecasting (Berger). Statistical forecasting has not solved the problem because it relies on historical data, which may or may not reflect current markets conditions, to predict future shipments. Collaborative forecasting has been used to forecast shipments, but has not solved the problem either because dispersed information is difficult to compile and integrate.

The challenges of forecasting in the consumer packaged goods industry is further complicated because the industry spends over 75 billion dollars per year on promotions (IBM & SAP) accounting for 15% to 25% of its total revenue; it is estimated only 30% of promotions run have a positive return on investment. The investment in promotions to increase consumer demand causes forecasts to become unreliable by shifting demand patterns.

To manage the complexities of forecasting in the consumer packaged goods industry, many companies have made large investments in planning systems. While these systems have stabilized forecast accuracy, it is still not uncommon for error rates to be 30% or more (Berger). Forecast error requires consumer packaged goods companies to carry extra inventory to prevent stock outs, potentially leading to excess inventory. As a result of these and other issues, consumer packaged goods companies are always seeking to improve their forecasting results. One method for improving forecast accuracy is Prediction Markets.

Prediction Markets forecast use a market mechanism; participants buy and sell shares in quantity ranges to forecast demand (Berg et al. 2008). In the area of election results (Berg et al. 2004), new product launches (Ho & Chen) and technology adoption (Mangold et al.) Prediction Markets have developed more accurate predictions than traditional forecasting approaches.

There are four elements that are consistently credited for the forecasting success of Prediction Markets:

1. Prediction Markets aggregate information from multiple disparate sources.
2. Prediction Markets are relatively immune to coercion and manipulation.
3. Prediction Markets offer incentives and rewards for consistent good performance.
4. Prediction Markets have a market maker that enables them to function in low participation settings.

These elements have been extensively documented in an academic setting. Unfortunately, there has been little published research presenting Prediction Markets in a corporate setting; we performed our research with General Mills.

Over 175 people were trained and given access to our Prediction Markets; the participants were selected from customer service, finance, marketing, operations and sales departments and given \$100,000 play money to invest. The result of their trading activity was compared against the Operations Forecast and shipment results in order to determine forecast accuracy and bias. We have worked to extend existing research and apply it to the consumer packaged goods industry through our work with General Mills.

2 Literature Review

This literature review is designed to cover five major questions that have been addressed as Prediction Markets have moved from academic theory to market acceptance. First, can Prediction Markets develop more accurate forecasts than traditional approaches? Second, do Prediction Markets aggregate multiple points of view more efficiently than current techniques? Third, what metrics should be used to evaluate the success or failure of Prediction Markets?

Fourth, how well do Prediction Markets resist manipulation? Finally, what levels of success have Prediction Markets had within the business community?

Early Prediction Markets were created by the University of Iowa in 1989. The markets allowed people to buy and sell shares in political candidates thereby quantifying who would win a given election. It was not until 1996 that research about the Iowa election markets began to emerge; this research concentrated on describing Prediction Markets rather than analyzing results. This literature review begins in 2003 when dissection of Prediction Market Forecast accuracy began in earnest.

In 2003, with the publication of “Results from a Dozen Years of Election Futures Markets Research,” Berg et al. examined the accuracy of Prediction Markets. They described the ability of the Iowa election markets to aggregate information from multiple participants. In this paper they showed that the average market error was 1.41%, versus 1.91% for polls; this result held true over time. Berg et al. compared actual results to the Prediction Market Forecasts and calculated the mean absolute percentage error (MAPE) to judge the results. The research concluded, based on forecast accuracy, that Prediction Markets would be applied to other areas in the future. We adopted, as have many others, the same measurement approach for the Prediction Markets at General Mills.

In *The Wisdom of Crowds* (2004), Surowiecki explained how groups of people were able to deliver better results than experts or individuals. He began in 1884, when Francis Galton observed that individual fairgoers could not guess the weight of a steer, but were able to guess the weight correctly when their responses were averaged together. Surowiecki concluded his book with the 1999 show, “Who Wants to be a Millionaire,” where group voting helped predict the answer to difficult questions. He described cases where group dynamics delivered better

solutions and where they did not; in doing so Surowiecki raised awareness of the benefits of Prediction Markets in the business community. This book provided us with insights to understand how the General Mills' culture would react to Prediction Markets.

In her 2004 Time article, Kiviat chronicled the move of Prediction Markets from university experiments to business pilot projects. Ms. Kiviat explained that companies such as Microsoft, Eli Lilly and Hewlett-Packard used Prediction Markets to distill information from employees and gain insights into project completions and future forecasts. She cited specific results at HP, where Prediction Markets outperformed marketing managers 75% of the time. She showed that Prediction Markets could answer general questions and, as the results have shown, aggregate multiple points of view. We used this article to understand how Prediction Markets solve business forecasting problems.

In 2004, Wolfers and Zitzewitz wrote the seminal compilation of Prediction Markets research; this paper has been referenced by forty-one other authors, four times more than any other paper. Their paper provided an overview of the benefits and risks that accrue to Prediction Market users. In pulling this information together Wolfers and Zitzewitz laid the foundation for others to begin drilling deeper into the inner workings of Prediction Markets. We used this paper to identify the major Prediction Markets that have been used in the majority of the research studies to date.

In his 2004 master's thesis, Schrieber, described the uses of Prediction Markets within a business context. He classified three benefits companies could derive from Prediction Markets: accuracy, immediacy and insight. According to Schrieber, Prediction Markets provided more accurate forecasts than traditional methods. He argued that Prediction Markets were more immediate because they were able to aggregate the opinions of multiple people in real-time. He

concluded Prediction Markets provided additional insight by providing a range based forecast rather than a single number. This thesis provided the metrics that we used to measure the effectiveness of the Prediction Markets at General Mills.

In her 2004 master's thesis, studying the Iowa election markets, McCabe compared the accuracy and bias of Prediction Markets to econometric models and industry experts. She concluded that Prediction Market Forecasts were as accurate as the experts, but noted that markets adjusted to new information more quickly and suffered less bias. McCabe implied that because markets did not have "reputations" to uphold, they could adjust to new information much more rapidly than models or experts who had clients to report to. In determining that Prediction Markets reduced bias, McCabe showed that Prediction Markets can add value to the forecasting process even if they are not able to generate additional accuracy. This thesis confirmed Shrieber's metrics of forecast accuracy, bias and insight were the correct measures for evaluating Prediction Market results.

In their 2004 paper studying the Iowa election markets, Wolfers and Zitzewitz described how contract construction and form related to proper market function. They presented two major types of contracts: winner-take-all and index. Winner-take-all contracts, contracts where the stock either pays money out or does not, are the easiest contracts to administer because they have a single outcome: win or loss. Index contracts, contracts where the stock pays out based on how close the actual outcome was to the purchase price, are challenging to administer because potentially all participants can receive a payout. Wolfers and Zitzewitz concluded with examples illustrating that Prediction Markets are more stable, but tend to overvalue low probability events. According to Wolfers and Zitzewitz, Prediction Markets are well suited to helping managers

make decisions if contracts are well thought out and low probability events are accounted for.

We used these principles to configure the Prediction Markets at General Mills.

In 2005 Yahoo! Research Labs and O'Reilly Media created a Prediction Market called the "Tech Buzz Game." This project used Prediction Markets to anticipate the future of technology. In the market, Mangold et al. encountered many players who subverted the rules by setting up multiple accounts enabling them to inflate prices; this issue was quickly resolved through email verification limiting participants to a single sign-on. Another issue occurred when a pair of seventeen year old students uncovered a flaw in the pricing logic enabling participants to drive all stock prices within a market to zero; this issue was corrected by changing the market pricing algorithm. Based on these experiences, we set trading limits for each Prediction Market and used linked pricing algorithms within the General Mills Prediction Markets.

In "Information aggregation and manipulation in an experimental market" (2006), Hanson et al. explored the effects of price manipulators on Prediction Market stock prices. The experiment selected groups of twelve market participants and offered to pay two of them a bonus if they were able to push the price of a stock above fifty dollars. The participants were then allowed to trade for 12 timed trading sessions. Although the would-be manipulators consistently bid above market prices for assets, they were not able to drive the price above fifty dollars. Unlike the Yahoo! and O'Reilly experience, each market participant was allowed one account, preventing one person from dominating the market through nefarious means. The authors concluded that Prediction Markets were robust and could resist price manipulation attempts in many circumstances. This paper provided the background we used to set trading hours and initial funding levels for the Prediction Markets at General Mills.

In 2006, Guo et al. proposed a framework for incorporating Prediction Markets into the Demand Planning. They proposed that Prediction Markets offered a mechanism for bringing retailers and suppliers together to optimize channel inventory and maximize profits. They cited the ability of markets to aggregate input and maintain anonymity as essential ingredients for allowing Prediction Markets to help retailers and suppliers share pricing information. In proposing this use of Prediction Markets, Guo et al. moved the benefits of Prediction Markets beyond accuracy and bias to strategy and supply chain integration. This paper helped us formulate additional uses for Prediction Markets at General Mills.

In 2007, Ho and Chen compared and contrasted new product forecasts developed using Prediction Markets with survey results and expert opinions. They laid out parameters for recruiting participants, setting market budgets and developing incentives to overcome the issues experienced in the Tech Buzz Game. They concluded that in successful markets, incentives must align with corporate goals, investment levels must prevent participants from cornering the market, and there must be a large number of participants. If these criteria are met, markets will succeed in delivering better results than surveys or experts. This paper helped us determine incentives for participants who participated in the Prediction Markets.

In 2008, Berg et al. delved into the long term results delivered by the Iowa election market. They showed Prediction Markets were able to deliver more accurate forecasts in the 66 to 100 day range when compared with polls. In this time period, Berg et al. found that Prediction Markets outperformed polls 68% to 84% of the time; consistent with the HP results presented by Kiviat. This result, according to Berg et al., was largely driven by the self selecting nature of Prediction Markets. Self selection is a crucial difference between Prediction Markets and polls;

polls reach out to a group of people who may or may not be interested in the poll. This paper guided us to expand the number of participants in the Prediction Market pilot.

In 2008, Gartner initiated coverage of Prediction Markets, listing them as emerging technology with moderate business impact and less than 1% penetration. Cain and Drakos noted that while a number of vendors have emerged to service this market, many of the early adopters have entered the “Trough of Disillusionment” because they overestimated the impact the technology would have on their organizations. In Gartner’s opinion, Prediction Markets can best be applied to estimating sales volumes, product delivery dates and capacity needs. The coverage concluded with the statement, “Potential users should start with pilot programs so they can compare the results with traditional forecasting mechanisms.” This report prompted us to run a mini-pilot with eleven participants from Demand Planning and seven from Sales to identify issues we might encounter when running the full scale pilot with General Mills.

The final review performed was with Intel (Hopman) which has been using Prediction Markets for over three years. Intel has made Prediction Markets a core component of its planning process by gathering real time feedback in areas where there is disagreement or uncertainty. Intel set up markets that reach out to informed participants who have knowledge to contribute. This approach stems from the self selecting nature of markets and legal requirements surrounding SEC financial disclosure rules. In addition to addressing participation issues, Intel has integrated Prediction Markets into their compensation system; market participants receive their winnings in their paychecks. According to Hopman, Intel has seen improvements in forecast accuracy, immediacy and insight from Prediction Markets. The interview helped us avoid mistakes in setting up our Prediction Markets and caused us to pay close attention to those participants who had the most accurate input into the Prediction Market Forecasts.

This thesis seeks to extend the academic and business research that has been done by examining the applicability of Prediction Markets to the consumer packaged goods industry. Our goal will be to learn from past experiments and add more support to the results that other researchers have seen. The research will examine the effectiveness of Prediction Markets under three forecasting scenarios: predicting corporate sales volume, predicting customer specific sales volume and predicting promotional sales volume thereby extending the understanding of how Prediction Markets operate in a business setting.

3 Methods for Configuring Prediction Markets

As discussed in the literature review, Prediction Markets have been used for forecasting both numbers and events. The majority of Prediction Market research has been conducted in controlled environments or with large public sites such as the Iowa Election Markets or Betfair.com. The evidence shows Prediction Markets provide better results than experts or models (Wolfers & Zitzewitz 2006). Given this research, we wanted to extend the application of Prediction Markets to a consumer packaged goods company: General Mills. In this section we examine how Prediction Markets would operate in a corporate environment using a pilot; how we configured the Prediction Market used to gather our results; and finally, the calculations and surveys we used to assess the effectiveness of Prediction Markets at General Mills.

3.1 Prediction Market Pilot

Our literature review helped us isolate three necessary conditions that help Prediction Markets generate better results than experts or models.

1. There must be appropriate contract types for participants to trade in the Prediction Markets.

2. There must be a mix of informed and uninformed participants with access to the Prediction Markets.
3. There must be safeguards to prevent participants from gaining unfair advantage by manipulating the Prediction Markets.

The literature review showed that if we did not meet the three conditions, our Prediction Markets would fail to provide the data we needed in order to answer General Mills' question regarding forecast accuracy improvements.

As the project sponsor, General Mills asked us to research if Prediction Markets could improve the accuracy of its current forecasting process. As with many new initiatives, we decided to conduct a small pilot before moving to our study to build our knowledge and avoid pitfalls that often kill new projects in corporate environments. In hindsight, running a Prediction Market pilot was the best decision we made throughout the course of our research; we learned that participants would manipulate markets and required training to participate.

3.1.1 Contract Types within Prediction Markets

There are three different types of contracts that can be used to gather forecasts using Prediction Markets (Wolfers & Zitzewitz 2004): spread, index and winner-take-all. Each contract can contain one or more stocks that participants can buy or sell. A common attribute among Prediction Market contracts is having a price between \$0 and \$100. Sometimes the price can represent a probability and sometimes it can represent a forecast value; the difference between the contracts is how the range is interpreted.

In their seminal paper Wolfers and Zitzewitz explain the different Prediction Market contract types. They explain that spread contracts pay out like sports bets where a team must win by a number of points in order for the bet to pay off; index contracts pay out like grades where

each percentage point accrues value to the owner with the maximum payout being made for the perfect outcome; winner-take-all contracts pay out like lottery tickets where the owners ticket either has the correct number and they get paid or it doesn't and they receive nothing. In each case the stock, implementing the contract type, will pay out between \$0 and \$100.

In researching spread contracts we found them to be unsuited for gathering forecasts in a corporate environment (Schreiber). Spread contracts are designed to capture the distance between two values; in forecasting, this would be mean absolute percent error (MAPE) between actual and forecast. Figure 3.1.1 illustrates how spread contracts capture information.

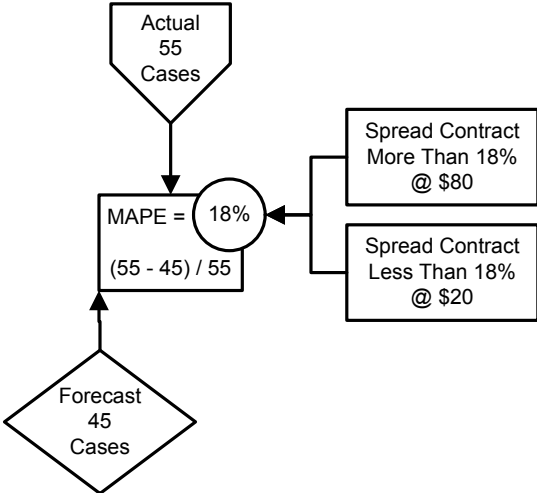


Figure 3.1.1 - Spread Contract

Figure 1 shows that 55 cases were sold and 45 cases were forecast yielding a spread of 10 cases or a MAPE of 18%; e.g. $(55 \text{ cases sold} - 45 \text{ cases forecast}) / 55 \text{ cases sold} = \text{an } 18\% \text{ MAPE}$. The spread contract asks participants to buy shares in whether or not they believe that the MAPE will be more or less than 18%; as illustrated in Figure 3.1.1, participants have the choice of two stocks one that costs \$80 showing the MAPE will be greater than 18% and another that costs \$20 showing the MAPE will be less than 18%. Because the prices for the stocks in the contract must add to \$100, we know that the participants believe that there is an 80% chance that

the MAPE will be greater than 18% and only 20% believe that the MAPE will be less than 18%. This illustrates why spread contracts are not used in forecasting; most companies are not interested in knowing whether the MAPE is greater or less than some percentage without having a base forecast to reference the MAPE against. While spread contracts have application in other areas, such as sports betting, the general consensus is that they are inappropriate when forecasting in a corporate environment (Shreiber).

Index contracts closely mirror current forecasting practice because they deliver point forecasts based on price (McCabe) using a single stock to implement the contract. When configuring an index contract, a scaling factor is used to link the stock price and forecast value, between \$0 and \$100 dollars, and the quantity being forecast; an \$18 stock price could equal 18 cases or it could equal 1,800 cases or it could equal 3,600 depending on the scaling factor used.

Figure 3.1.2 provides a visual representation of an index contract.

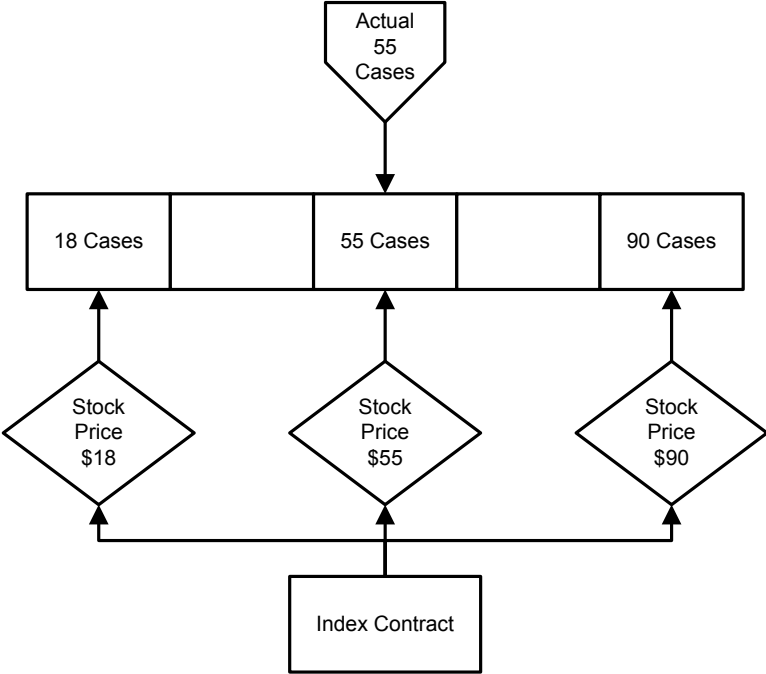


Figure 3.1.2 - Index Contract

In this example, the scaling factor of stock price to actual cases is one to one; e.g. an \$18 stock price equates to 18 cases, a \$55 stock price equates to 55 cases and a \$90 stock price equates to 90 cases; it is important to note that an index contract may only take on one value at a time. The price of the stock within the Prediction Market is determined by the number of shares that participants buy and sell; if participants buy shares, the price and hence the forecast goes up; if participants sell shares, the price and hence the forecast goes down. If participants agree with the current forecast, then no shares would be bought or sold; because index based contracts link the price with the forecast they can reduce trading activity once the forecast matches what participants expect it to be. Index contracts provide a method for capturing forecasts that closely mirror current forecasting practice by delivering a point forecast using a single stock. Because they create point forecasts, we expected to run all of our Prediction Markets using index contracts.

Winner-take-all contracts are the most prevalent contract type used to capture information from Prediction Markets (Wolfers & Zitzewitz 2004). In most instances winner-take-all contracts are associated with ranges of forecasts (Hopman). Each forecast range is implemented using a single stock and the price of that stock is determined by how many shares participants buy and sell for that stock; buying and selling of shares within a range may change the price of the stock, but will not change the value of the forecast represented by that stock. Thus, a winner-take-all contract will generally be implemented with more than one stock. Figure 3.1.3 illustrates how forecast ranges, each representing a stock, are used to implement winner-take-all contracts.

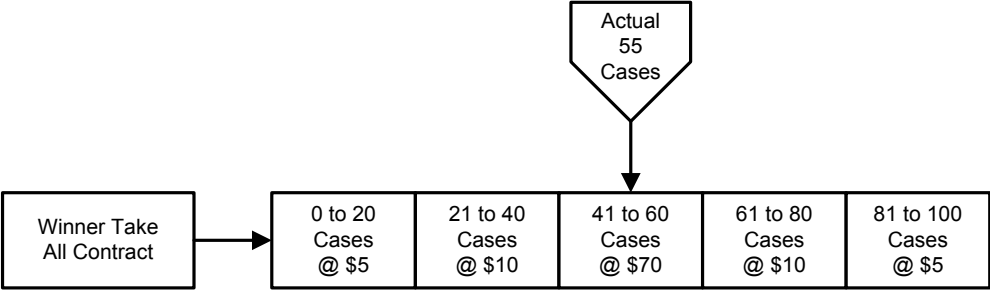


Figure 3.1.3 - Winner-take-all Contracts

Figure 3.1.3 must be interpreted using a two step process. In the first step, forecast ranges for each of the contracts are identified; e.g. 0 to 20 cases, 21 to 40 cases up to 81 to 100 cases. The second step, for winner-take-all contracts, is interpreting the price; all stock prices must add to \$100. Since all of the prices add to \$100, the contract price becomes the probability that the actual value will fall into the stocks forecast range. In Figure 3.1.3, the price shows that there is a 70% chance that the actual sales will fall in the range covered by the 41 to 60 cases stock. Selecting the right contract type was important to our study, to ensure we made the right selection, we tested both index and winner-take-all contracts in our pilot Prediction Market.

3.1.2 Selection of Prediction Market Pilot Participants

Prediction Markets require informed and uniformed participants in order to operate properly (Wolfers & Zitzewitz 2004). As with many corporations, employees tend to operate in their own organizational silos; at General Mills, these silos are organized along brand lines. At General Mills, a demand planning manager for cereals would not have detailed knowledge about the yogurt forecast. Thus, the nature of General Mills’ organizational structure ensures that a Prediction Market focused on cereals would have one informed demand manager and the other brand’s demand planning managers would be uninformed because cereal was not their primary focus though they would have some sense of what a reasonable cereal forecast would be.

It would seem that whether or not uninformed participants purchased stocks in a market would be of little importance, but they are required to provide liquidity to the market (Wolfers & Zitzewitz 2006); Intel disagrees with this conclusion and uses its Prediction Markets solely with informed participants (Hopman). Since there was not a consensus in the literature we decided that participation was an item we would need to measure in our pilot. We decided to measure the degree to which uninformed participants joined a market by counting the number of participants who bought or sold stocks in each of the Prediction Markets; to measure the degree of information we followed up with a debriefing session where we asked the participants how much information they had for each of the Prediction Markets that they bought or sold stocks in. By comparing the two numbers we were able to gauge number of uninformed and informed participants for each of the Prediction Markets.

To test for uninformed participants we configured three Prediction Markets for our pilot. The first Prediction Market asked, “Who will win the Minnesota senate race, Coleman or Franken?” This question was selected because the campaign had been very contentious and many conversations were taking place at the General Mills water cooler; as a result we expected most of the participants would purchase stocks in this Prediction Market. The second question we selected was, “What will General Mills second quarter shipments be?” This question was designed to test the knowledge of demand managers and sales people regarding aggregate corporate numbers; we expected a majority of the participants to participate in this market. The third question was, “What will second quarter shipments of Product One be?” This question was designed to test brand specific knowledge; we expected very low participation in this market because only a few people had direct knowledge of this Prediction Market.

If the pilot metrics showed that uninformed participants participated heavily in the markets and overwhelmed informed participants, then we would decrease the number of participants to those individuals who General Mills judged to have knowledge about the markets in our study. If, on the other hand, the pilot showed that uninformed participants did not overwhelm the pilot Prediction Markets, we would expand the number of participants in our study to gather information from people who did not participate in the current forecasting process.

3.1.3 Testing for Prediction Market Manipulation

Hanson et al. found that Prediction Markets were immune to manipulation. Mangold et al., in contrast, found that Prediction Markets can be manipulated. Given these contradictory points of view, we needed to determine if participants would attempt to manipulate our pilot Prediction Markets.

With the specter of manipulation looming over our Prediction Markets, we decided to ensure that participants would have the ability to manipulate markets; to ensure this we gave each participant \$25,000 to invest in the three Prediction Markets we configured for the pilot. By making our own investments in the pilot Prediction Markets, we knew that every \$5,000 invested in a stock would yield a \$15 rise in a contract's price. Therefore, if a participant invested all of their money in a single stock, they could increase the price by \$75.

If we return to the generic index contract in Figure 3.1.2, we could envision a single participant being able to drive the price of the index stock to \$75 corresponding to a forecast of 75 cases. If this forecast was outside the realm of believability, then other participants may not participate in the market because the manipulator owned all of the shares between \$0 and \$75; in essence, there would be no way for any other participants to own shares because any purchases

would further increase the price. We have many of the same issues for the generic winner-take-all contract in Figure 3.1.3. We could envision a single participant being able to drive the price of one stock to \$75 and decrease all of the other prices because they must add to \$100. Thus, unless other participants thought that the forecast was in another range, they would not participate in the market because the manipulator had locked them out of the market by increasing the price beyond what most participants might be willing to pay.

Because of our initial testing, we knew that with \$25,000 participants would be able to manipulate the Prediction Markets if they chose to. Once we understood how the participants would manipulate the Prediction Markets we would then be able to implement safeguards to prevent manipulation during our study.

3.2 Prediction Market Configuration

The parameters used to configure a Prediction Market have a significant effect on the value of the forecasts that it generates. While many studies have been performed on the Iowa Election Markets and Betfair.com, there has been little research performed in corporate environments. Because we conducted a Prediction Market pilot, we had gained an understanding of how Prediction Markets functioned in General Mills' environment. This section lays out what we learned in our Prediction Market pilot and describes the parameters we used to set up the Prediction Markets we used to gather data for our research.

3.2.1 Participant Selection

The Prediction Market pilot provided us with information regarding how participants would interact with the Prediction Markets. Table 3.2.1 presents the number of participants that took part in the three Prediction Markets set up as part of the pilot.

Question	Participants	Percent of Total
Who will win the Minnesota Senate race, Coleman or Franken?	14 of 18	77%
What will General Mills second quarter shipment volume be?	10 of 18	56%
What will second quarter shipments of Product One be?	5 of 18	28%

Table 3.2.1 - Prediction Market Participation

As we expected, the number of participants dropped off dramatically as the Prediction Markets asked for more specific information. Table 3.2.1 shows that uninformed participants do not participate in Prediction Markets where they do not have information; this can be seen by the decrease in the percentage of the participants from 77% on the broadest question to 28% on the most specific question. The eleven demand managers confirmed that they did not participate in markets where they felt like they did not have any information; they went on to say that they believed that others within the General Mills organization would react in a similar manner and that it was unlikely that uninformed participants would overwhelm informed participants. Based on this feedback, we worked with General Mills to increase the number of participants who would take part in the main Prediction Market study from 18 to 168.

In addition to increasing the number of participants, we also increased the number of business functions involved in the study. Table 3.2.2 presents the number of participants by function plus the number of participants who took part in the initial pilot.

Departments	Participants in Research Study	Participants in Prediction Market Pilot	Part of Current Forecasting Process
Customer Service Center	24		No
Finance	8		Yes – through conversations
GroceryCo Product Sales Manager	11		Yes – through Conversations
GroceryCo Sales	15	3	Yes – through trade planner
Marketing	9		Yes
Product Sales Manager	3		Yes – through conversations
Corporate Sales Management	8		No
Demand Planning	39	11	Yes
BoxCo Product Sales Manager	19		Yes – through conversations
BoxCo Sales	33	4	No
Grand Total	169	18	

Table 3.2.2 - Summary of Participants by Job Function

As shown in Table 3.2.2, we were able to include participants from departments that did not have a direct voice in the planning process. By including the individuals from these departments we hoped to determine if better information existed that could be shared through the Prediction Market.

3.2.2 Prediction Markets Configured at General Mills

In addition to expanding the number of participants, we also expanded the number of Prediction Markets from three to twenty four to gather three categories of forecasts: general, GroceryCo specific and BoxCo specific. General Prediction Markets were set up to gather forecasts for General Mills in aggregate; e.g. What will Product One Q3 deliveries be? GroceryCo Prediction Markets were set up to gather GroceryCo specific forecasts for promotions and other product categories; e.g. What will GroceryCo deliveries be for January & February across all Special Events? BoxCo Prediction Markets were set up to gather specific forecasts for everyday low price product categories; e.g. What will Product One deliveries be for BoxCo? In addition to assigning a category to each of the Prediction Markets, we also assigned each market a unique ticker symbol for easy reference. A common attribute of a well defined market is having a specific question accompanied by the data required for participants to perform research and make intelligent decisions. Table 3.2.3 presents all of the Prediction Markets we configured for our study.

Prediction Market Market Type	Prediction Market Forecast Type	Prediction Market Ticker	Prediction Market Description
Corporate	Product Category	PONEQ3	What will Product One Q3 deliveries be?
Corporate	Volume	JAN	What will the total January deliveries be for US Retail?
Corporate	Volume	MAR	What will the total March deliveries be for US Retail?
Corporate	Volume	Q3	What will the total Q3 deliveries be for US Retail?
Corporate	Product Category	PTWOQ3	What will Product Two Q3 deliveries be?
GroceryCo	Promotional	GROCAR	What will GroceryCo deliveries be for January & February across all special event brands?
GroceryCo	Promotional	GROCARPSEVEN	What will GroceryCo deliveries be for January & February on the Product Seven special event brands?
GroceryCo	Promotional	GROCARPSIX	What will GroceryCo deliveries be for January & February on the Product Six special event brands?
GroceryCo	Promotional	GROCARPFIVE	What will GroceryCo deliveries be for January & February on the Product Five special event brands?
GroceryCo	Volume	GROMAR	What will the total March deliveries be for GroceryCo?
GroceryCo	Volume	GROQ3	What will the total Q3 deliveries be for GroceryCo?
GroceryCo	Promotional	GROPFOUR	What will GroceryCo deliveries be for Product Four brands during February & March?
GroceryCo	Promotional	GROPTWO	What will the GroceryCo total deliveries be for January & February for Product Two items?
GroceryCo	Promotional	GROPTHREE	What will GroceryCo Product Three deliveries be February/March?
BoxCo	Product Category	BOXPONEQ3	What will the Q3 Product One deliveries be for BoxCo?
BoxCo	Volume	BOXJAN	What will the total January deliveries be for BoxCo?
BoxCo	Volume	BOXMAR	What will the total March deliveries be for BoxCo?
BoxCo	Product Category	BOXPTWO	What will the Q3 Product Two deliveries be for BoxCo?
BoxCo	Volume	BOXQ3	What will the total Q3 deliveries be for BoxCo?
BoxCo	Product Category	BOXPTHREEQ3	What will the Q3 Product Three deliveries be for BoxCo?

Table 3.2.3 - Prediction Markets Configured at General Mills

The Prediction Markets were configured so that each group (see Table 3.2.2) involved in the study would have at least one area where they had expert knowledge that others did not share. The general markets were set up to capitalize on the overall picture that finance, Demand Planning and Marketing have for General Mills as a whole. The GroceryCo and BoxCo markets were set up to capitalize on the detailed picture that the GroceryCo and BoxCo sales organizations and customer service have by interacting directly with GroceryCo and BoxCo employees.

A key attribute of Prediction Markets is that at some point, they all come to an end. This end state can be based on a date or other objective attribute that can be used to judge all of the stocks that participants hold in that market to declare a winner. In each case, the winner of the Prediction Market is the person who has made the most money.

3.2.3 Prediction Market Contracts

As described in section 3.1.1, contract type selection is critical to a Prediction Markets' success. If the wrong contract type is selected, then two issues will emerge; first, the Prediction Markets will not be able to gather meaningful information and second, the participants will not understand how to interact with the Prediction Market. Based on our Prediction Market pilot, we believe that if either situation occurs Prediction Markets will fail.

As a result of the Prediction Market pilot, we realized that index contracts would not work in the General Mills' environment because the clearing price represents the forecast. Therefore, once the price equated to the forecast that participants believe to be accurate, then no one can buy or sell shares because the price and hence the forecast would change. We found that the price sensitivity of index contracts was so high that one participant could drive the price beyond a reasonable forecast range; e.g. the first two trades, in pilot Prediction Market, the "What will General Mills second quarter shipments be?" forecast had jumped to 394 million cases, 220 million more cases than any quarter in General Mills history. After interviewing the participants, we discovered that they wanted to buy shares in the index contract when it reached the market clearing price; "I saw that the price was equivalent to the forecast that I thought was right and so I decided to buy" was what one of the participants said. This behavior meant that participants would drive the price of index contracts outside the range of believability within

minutes of the market opening. It was for this reason that we abandoned index contracts for developing Prediction Market Forecasts at General Mills.

The participants commented that they felt that they could easily understand and interpret winner-take-all contracts. The general consensus in the debriefing was that winner-take-all contracts were the best method for gathering forecasts because:

1. Participants could buy as much of any stock as they wanted without changing the value of the forecast. See Figure 3.1.3 for a pictorial representation of winner-take-all contracts.
2. Participants could easily determine the current forecast by looking at the price of each stock within the Prediction Market. See Figure 3.1.3 for a pictorial representation of winner-take-all contracts.

As a result of this feedback, we selected winner-take-all contracts for all of the Prediction Markets within our study.

3.2.4 Setting Stock Ranges in a Prediction Market

When using winner-take-all contracts to gather forecasts using Prediction Markets the number of stocks used to do so becomes important. If there are too few stocks, then the participants will purchase a single stock and the Prediction Market Forecast range will be too large to be meaningful; e.g. if a Prediction Market contains two stocks, one with a range of 0 to 1,000,000 cases and another with a range of 1,000,001 to 2,000,000, the range for each of the stocks will be so broad that the results will not be meaningful. On the other hand, if there are too many stocks, then the participants will not know which stock to purchase and the Prediction Market will not be able to converge on a forecast; e.g. if a Prediction Market contains 2,000,000 stocks each representing a range of one, there will be too many choices for participants to select and the

market will not converge on a forecast. Therefore, selecting a reasonable number of stocks is important to gathering meaningful forecasts with Prediction Markets.

The first approach we evaluated involved setting an overall range and then slicing it into a number of divisions; we will call this the range selection approach. The second approach we evaluated involved selecting a central range and then expanding a set number of ranges above and below to construct the total range spanned by the market; we will call this the expansion selection approach. We chose the expansion selection approach because we believe it yields forecast ranges that are easier for participants to understand.

When setting up Prediction Markets, organizations either know the range that the forecast can exist in or the unit increments that participants think in. The unit increment is the level of detail that participants forecast in; e.g. 10's of cases, 100's of cases and so on. If the forecast range is known, then it is best to slice it into a number of equal divisions that get turned into stocks. If the unit increment is known then it is best to select a central point and expand the range in even unit increments until the range of forecast possibilities is covered. The following section provides a detailed description of the two range setting methods.

The range selection approach provides the simplest method for setting ranges within a Prediction Market. Figure 3.2.1 provides an illustration of this approach.

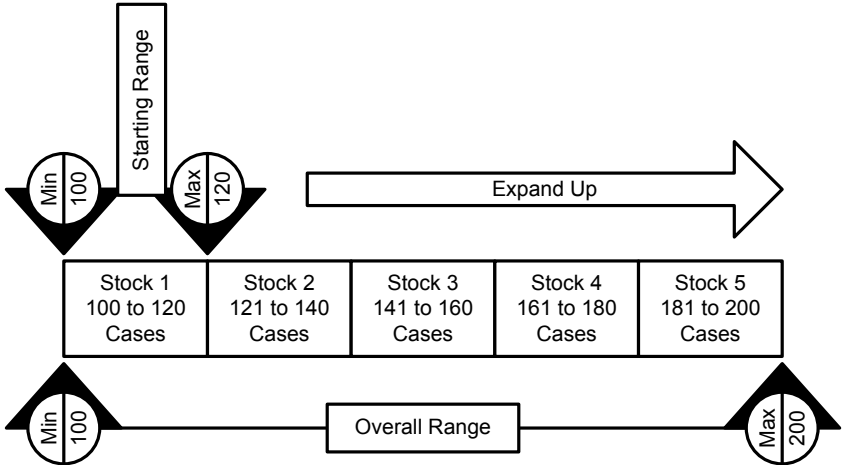


Figure 3.2.1 - Range Selection Approach

The first step in using the range selection approach is to choose the overall range covered by the Prediction Market; e.g. the overall forecast range for the prediction shown is 100 to 200 cases. The determination of the minimum and maximum value in the Prediction Market range is set by the group in charge of the Prediction Market; in our case, if we had used this method, we would have gotten this range from the Demand Planning group at General Mills. The second step is selecting the number of divisions for the Prediction Market; five divisions have been selected in this case. The third step sets the range for each stock; e.g. there are five stocks and a range of 100 cases yielding a range of 20 cases for each stock. The fourth step sets the range for the first stock; e.g. the first stocks range is set as 100 to 120 cases. Step four is repeated for the remaining stocks in the Prediction Market until the maximum value is reached.

The expansion selection approach provides another approach for calculating the ranges within a Prediction Market. Figure 3.2.2 provides an illustration of this approach.

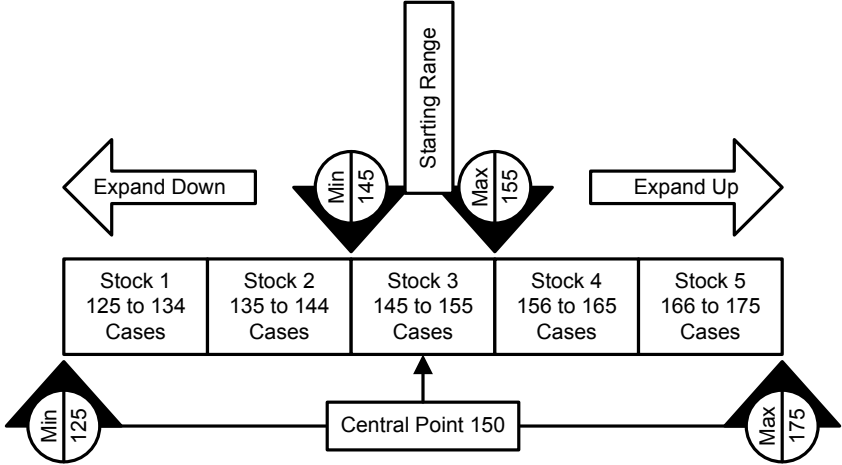


Figure 3.2.2 – Expansion Selection Approach

The first step in using the expansion selection approach is to choose the central point covered by the Prediction Market; e.g. in this case, the central point has been determined to be 150 cases. The determination of the central point for the Prediction Market range is usually set by the group in charge of the Prediction Market; in our case, we used the Operations Forecast computed by the planning group at General Mills. The second step is selecting the starting range for the Prediction Market; in this case, the starting range was set to ten cases yielding a stock range of 145 to 155 cases; e.g. 5 cases on either side of the central value of 150 cases. The third step adds divisions above and below the starting range. Step three is repeated until the desired range of values is spanned by the Prediction Market; in this case, the range covered by the five stocks in the Prediction Market is 125 cases to 175 cases.

Both the range (Figure 3.2.1) and the expansion (Figure 3.2.2) methods are effective methods for setting up Prediction Market stocks. We used the range selection method for configuring the pilot Prediction Market. The issue we found with the range selection approach was that the participants found the forecast ranges for each stock to be too large due to the number of stocks we selected; as a result all of the participants purchased shares in the same stock and were frustrated because they had more accurate information that they could not act on

due to the span of the forecast range. Another issue that surfaced was that the majority of the participants based their forecast, at least in part, on the current Operations Forecast; during the pilot Prediction Market debriefing the participants suggested that having forecast ranges branch out from the Operations Forecast would make the stocks much more meaningful. We decided to use the expansion approach for setting the stock ranges.

3.2.5 Setting the Number of Stocks in a Prediction Market

Because we used the expansion selection approach for setting our stock ranges, we opted to have an odd number of stocks in our Prediction Markets; we wanted to have an even number of ranges on either side of our central range. In the pilot Prediction Market, we had six stocks to capture the forecasts. Combining feedback from the pilot Prediction Markets (demand management group) with input from the software solution provider (ConsensusPoint), we set the number of stocks to nine so the overall range would cover a wide enough range (Perry and Kittlitz).

Figure 3.2.3 illustrates the configuration we used to set up the JAN Prediction Market; all of the Prediction Markets (see Table 3.2.3) were configured using the same approach.

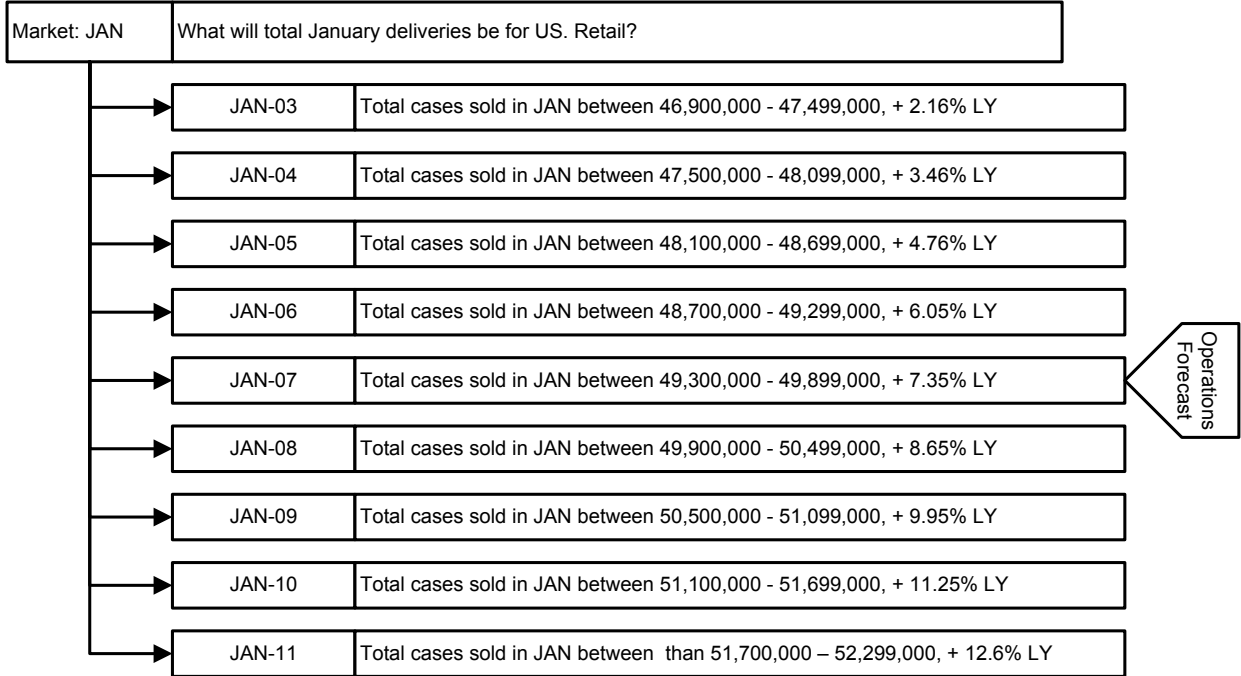


Figure 3.2.3 - JAN Prediction Market Configuration

As illustrated in Figure 3.2.3, there are nine stocks that comprise the JAN Prediction Market. Each stock has a unique ticker that is derived by adding a number to the end of the Prediction Market symbol; we started the numbering system at three because we did not know whether or not we would have to add additional stocks to the Prediction Market and therefore wanted to leave two numbers open at the bottom of the range: JAN-01 and JAN-02 to cover this possibility. The center stock JAN-07 was set to contain the Operations Forecast and then each stock radiated out from that point with a forecast range of 600,000 cases. Using the approach laid out in Figure 3.2.2, we see that the central point is 49,600,000 cases, the starting range is 49,300,000 to 49,899,000 cases, the minimum is 47,499,000 cases and the maximum is 51,700,000. Thus the Prediction Market JAN was configured to cover a forecast range around the Operations Forecast of plus or minus 5.04%;

$$\frac{52,299,000 - 49,600,000}{49,600,000} = 5.04\% \text{ to } \frac{46,900,000 - 49,600,000}{49,600,000} = -5.04\%.$$

This approach was used to set up the ranges for all of the Prediction Markets listed in Table 3.2.3.

3.2.6 Determining the Winner of a Prediction Market

Before exploring the remaining configuration options within a Prediction Market it is important to understand the method by which the winner of a Prediction Market is judged. *The answer is simple – the participant who has made the most money in the Prediction Market is the winner.*

There are three trading strategies that participants can use to make money in Prediction Markets: buy and hold, buy and sell and selling short. At the end of section 3.2.2 we discussed that all Prediction Markets end and the winner of it is determined; this section will examine how the judgment process determines the winner.

To execute a buy and hold strategy, a Prediction Market participant simply purchases shares in the stock representing the forecast range that they believe will contain the actual result. Because we used winner-take-all contracts, the winning stock will pay \$100 and all others will pay \$0. Figure 3.2.4 illustrates a buy and hold strategy for three participants investing \$3,000 into the COJAN Prediction Market.

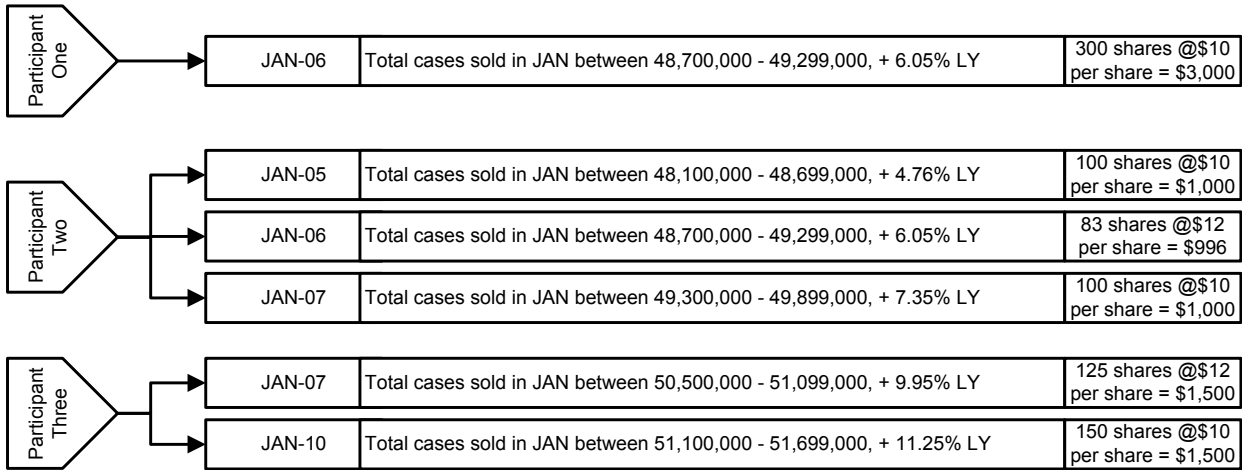


Figure 3.2.4 - Buy and Hold Strategy

Figure 3.2.7 illustrates three participants purchasing shares in the JAN Prediction Market. Participant One purchased 300 shares of JAN-06 for \$10 per share. Participant Two purchased 100 shares of JAN-05 and JAN-07 for \$10 per share and 83 shares of JAN-06 for \$12 per share. Participant Three purchased 125 shares of JAN-07 for \$12 per share and 150 shares of JAN-10 for \$10 per share. While each participant invested the \$3,000, the final value of their shares will be determined when the Prediction Market is judged and the winner is declared.

Given that participants in the JAN Prediction Market hold their shares (see Figure 3.2.4) until the market is judged, each has the chance of winning the Prediction Market. In order to determine the winner of the JAN Prediction Market we will examine the portfolios of three participants. The scenario presented assumes that the JAN Prediction Market event takes place and the actual deliveries for the Prediction Market are 48,933,000 cases; given this outcome, the JAN-06 stock would pay out \$100 per share and all other contracts would pay out \$0. Table 3.2.4 summarizes the results for each of the participants.

Participant	Share Value	Net worth
Participant One	Final Share Value $300 \text{ shares JAN} - 06 @ \$100 = \$30,000$ Less $300 \text{ shares JAN} - 06 @ \$10 = \$3,000$ Net Gain $\$30,000 - \$3,000 = \$27,000$	$\$27,000 + \$3,000 = \$30,000$
Participant Two	Final Share Value $100 \text{ shares JAN} - 05 @ \$0 = \$0$ $83 \text{ shares JAN} - 06 @ \$100 = \$8,300$ $100 \text{ shares JAN} - 07 @ \$0 = \$0$ Less $100 \text{ shares JAN} - 05 @ \$10 = \$1,000$ $83 \text{ shares JAN} - 06 @ \$12 = \$996$ $100 \text{ shares JAN} - 07 @ \$10 = \$1,000$ Net Gain $\$8,300 - \$2,996 = \$5,304$	$\$5,304 + \$3,000 = \$8,304$
Participant Three	Final Share Value $125 \text{ shares JAN} - 07 @ \$0 = \$0$ $150 \text{ shares JAN} - 10 @ \$0 = \$0$ Less $125 \text{ shares JAN} - 07 @ \$12 = \$1,500$ $150 \text{ shares JAN} - 07 @ \$10 = \$1,500$ Net Gain $\$0 - \$3,000 = -\$3,000$	$-\$3,000 + \$3,000 = \$0$

Table 3.2.4 - Buy and Hold Strategy Outcomes

Since Participant One made \$30,000 in the Prediction Market, they would be designated the winner of the JAN Prediction Market. Participant Two invested in two markets that did not span the actual, but did own shares in JAN-06 and so was able to make some money; Participant Two ends the market with \$8,304 because he invested in two stocks that went to zero. Finally, since none of the markets that Participant Three invested in spanned the actual deliveries, the value of Participant Three’s portfolio drops to \$0 causing Participant Three to lose the entire value of his investment. Table 3.2.4 clearly illustrates that when using a buy and hold strategy it is critical to purchase the maximum number of shares in the winning market at the lowest price.

To execute a buy and sell strategy, a Prediction Market participant simply buys and sells shares in a stock based on the current price. If the price of the stock increases enough, it is possible that the Prediction Market winner may not even own shares in the winning market. In the ideal case, the buy and sell strategy is used to move from a stock that does not match the

actual into one that does. Figure 3.2.5 illustrates a buy and sell strategy for two participants investing \$3,000 into the JAN Prediction Market.

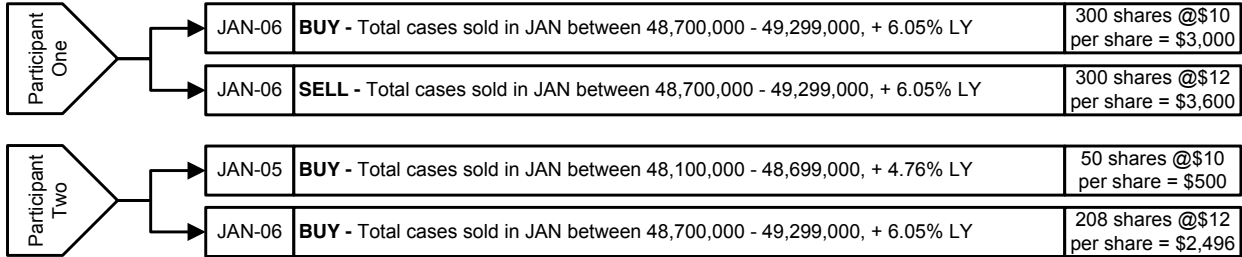


Figure 3.2.5 - Buy and Sell Strategy

Participant One purchased 300 shares of JAN-06 for \$10 per share and then re-sells them for \$12 per share. Participant Two purchased 50 shares of JAN-05 for \$10 per share and 208 shares of JAN-06 for \$12 per share. Because Participant One sold his shares of JAN-06, he will not be affected by the actual results of the market because he has locked in a profit of \$600 by selling the shares at a \$2 profit. Participant Two, on the other hand, will win or lose the market depending on what the actual deliveries for January are. If the deliveries are less than 48,100,000 or greater than 49,299,000, then Participant One will win because Participant Two’s shares will be worth nothing because the actual delivery value was not spanned by the contracts he owned; Participant One will win because he made \$600 in the market by selling his shares. If, on the other hand, the actual delivery value falls between 48,100,000 and 49,299,000, Participant Two will win because either JAN-05 or JAN-06 will pay off at \$100 per share and he will have made the most money in the market. Table 3.2.5 illustrates the outcome if actual deliveries are 47,000,000 cases.

Participant	Share Value	Net worth
Participant One	Buy and Sell Share Value	
	<i>Buy 300 shares JAN – 06 @ \$10 = \$3,000</i>	
	<i>Sell 300 shares JAN – 06 @ \$12 = \$3,600</i>	
	Final Share Value	
	0 shares JAN – 06 @ \$0 = \$0	\$3,000 + \$600 = \$3,600
Participant Two	Less	
	0 shares JAN – 06 @ \$0 = \$0	
	Net Gain	\$600 – \$0 = \$600
	Final Share Value	
	50 shares JAN – 05 @ \$0 = \$0	
Participant Two	208 shares JAN – 06 @ \$0 = \$0	
	Less	
	50 shares JAN – 05 @ \$10 = \$500	
	208 shares JAN – 06 @ \$12 = \$2,496	
	Net Gain	-\$2,996 + \$3,000 = \$4
	\$0 – \$2,996 = -\$2,996	

Table 3.2.5 - Buy and Sell Strategy

Because Participant One bought and sold shares, he leaves the JAN Prediction Market with a \$600 profit. As a result of this when the market is judged, Participant One has locked in his winnings while Participant Two ends up with a balance of \$4. Table 3.2.5 clearly shows that selling shares in a market can lead to victory even if it does not improve the accuracy of the Prediction Markets forecast.

The short selling strategy is an indirect strategy for winning Prediction Markets. To execute this strategy, a Prediction Market participant simply sells shares in a stock that he does not own. Figure 3.2.6 illustrates a short selling strategy for two participants investing \$3,000 into the JAN Prediction Market.




Figure 3.2.6 - Short Selling Strategy

Figure 3.2.6 illustrates two participants, one selling short and the other purchasing shares in the COJAN Prediction Market. Participant One sold 150 shares of JAN-05 and JAN-06 for \$10 per share. Participant Two purchased 125 shares of JAN-05 and JAN-06 for \$12 per share. Because Participant One sold his shares short he will receive a payout if the actual deliveries are not spanned by JAN-05 or JAN-06. Participant Two, on the other hand, will win or lose depending on what the actual deliveries for January are. In this case, Participant One is betting that the actual deliveries will fall outside of the range from 48,700,000 and 49,299,000 and Participant Two is betting that they will fall within the range. Table 3.2.6 illustrates the outcome if the actual deliveries fall outside of the 48,700,000 and 49,299,000 range.

Participant	Share Value	Net worth
Participant One	Final Share Value	
	150 shares JAN – 05 @ \$100 = \$15,000 150 shares JAN – 06 @ \$100 = \$15,000	
	Less	
	150 shares JAN – 05 @ \$90 = \$13,500 150 shares JAN – 06 @ \$90 = \$13,500	\$3,000 + \$3,000 = \$6,000
Net Gain	\$30,000 – \$27,000 = \$3,000	
Participant Two	Final Share Value	
	150 shares JAN – 05 @ \$0 = \$0 150 shares JAN – 06 @ \$0 = \$0	
	Less	
	125 shares JAN – 05 @ \$12 = \$1,500 125 shares JAN – 06 @ \$12 = \$1,500	-\$3,000 + \$3,000 = \$0
	Net Gain	\$0 – \$3,000 = -\$3,000

Table 3.2.6 - Outcomes for Values Outside the 48,100,000 and 49,200,000 Delivery Range

Since Participant One sold short, they will make \$30,000 on JAN-05 and JAN-06 less the short selling price of \$27,000 for both the stocks. Because he started with \$3,000, Participant One ends up with a final asset value of \$6,000. Participant Two, on the other hand ends up with a final asset value of \$0 because neither of his stocks spanned the actual deliveries.

Regardless of the strategy used, buy and hold, buy and sell or selling short, the winner of a Prediction Market is the participant who generated the most money in that market. As shown in the buy and sell and selling short examples, the winner may not have owned shares in the stock

that spanned the actual deliveries; however, because the participant had made the most money in the Prediction Market, they will be declared the winner when the event occurs. It is this flexibility that helps Prediction Markets derive accurate forecasts from a large number of participants with little or no supervision (Schreiber).

3.2.7 Setting Stock Prices in a Prediction Market

When using winner-take-all contracts to implement Prediction Markets, the price of each stock represents the probability of the actual result falling in that stock's forecast range. The probability value for each of the stocks is maintained because the stock prices add to \$100; e.g. for the JAN Prediction Market (see Figure 3.2.3) the prices for the nine stocks in the prediction add to \$100. Just as with stock ranges, see section 3.2.4, we considered two approaches for setting prices: average pricing and normalized pricing. In reviewing the literature and talking to experts we discovered that average pricing is the most prevalent pricing mechanism used to set the prices. The following section describes the process we used to evaluate average and normalized pricing for our Prediction Market study.

We believe that average pricing is used to implement most Prediction Markets because it is easy to calculate. To calculate the price of each stock, using an average pricing approach, simply divide \$100 by the number of stocks in the Prediction Market; using this approach, each stock in the JAN Prediction Market would have received a price of $\$11.11 \left(\frac{\$100 \text{ dollars}}{9 \text{ stocks}} = \$11.11 \text{ per stock} \right)$. Based on the results of the Prediction Market pilot, we discovered that average pricing would not work in the General Mills environment because participants could easily exploit the pricing to win a market without having any knowledge.

As discussed in section 3.1.3, we deliberately configured the Prediction Market pilot to allow participants to manipulate the Prediction Market for personal gain. The issue that we

encountered with average pricing was that participants were able to win Prediction Markets without any market knowledge by selling the ends of the market short. To execute this strategy, a participant would sell shares in the outermost ranges (see Figure 3.2.3), JAN-03 and JAN-11, knowing that it was unlikely that these ranges would span the actual delivery value. When the market event occurred, he made the most money (see Table 3.2.5). When we debriefed with General Mills', demand planning managers told us that they had exploited this weakness in the average pricing to make money. Because it was easy for smart traders to short low probability events and win Prediction Markets, we elected not to use the average pricing approach for our Prediction Market study.

Normalized pricing is a more complex and controversial method for setting prices in a Prediction Market (Thomas). When using this method to set initial prices, the average pricing method is modified by multiplying each average price by a factor and then normalizing the resulting values to ensure they add to \$100. Figure 3.2.7 illustrates this pricing approach.

Market: JAN		What will total January deliveries be for US. Retail?		
		Average Price (\$100/9)	Normalization Factor	Initial Market Price
→	JAN-03	\$11.11	.07	\$4.42
→	JAN-04	\$11.11	.12	\$7.53
→	JAN-05	\$11.11	.20	\$12.57
→	JAN-06	\$11.11	.25	\$15.74
→	JAN-07	\$11.11	.31	\$19.48
→	JAN-08	\$11.11	.25	\$15.74
→	JAN-09	\$11.11	.20	\$12.57
→	JAN-10	\$11.11	.12	\$7.53
→	JAN-11	\$11.11	.07	\$4.42

Figure 3.2.7 - Normalized Pricing Approach

The first step, in Figure 3.2.7, sets the average price by dividing \$100 by 9, the number of stocks, to calculate an average price for each stock in the Prediction Market. The second step is to create a normal distribution using Excel’s NORMDIST function centered on the JAN-07 stock and multiplying it by the average price calculated in the first step of the process. Step two ensures that the price of the central stock, JAN-07, will be the highest and the outside stocks, JAN-03 and JAN-11, will be the lowest. The third step normalizes the prices to make sure that they meet the Prediction Market requirement of adding to \$100; this is achieved by taking each value, calculated in the second step, dividing it by the sum of the values and then multiplying the result by 100. The result of this process is in Figure 3.2.7 shown as Initial Price Profile.

We selected the normalized pricing strategy for the Prediction Markets based on additional feedback that we received from the Prediction Market pilot. A simple winning strategy was to purchase shares in the forecast range that matched the Operations Forecast; if a participant was the first person to purchase shares in the stock, they would have the lowest price shares in the market and, when the market event occurred, they would make the most money in the market. Normalized prices limited this ability by making the forecast corresponding to the Operations Forecast the most expensive forecast in the market; see sections 3.2.4 and 3.2.5 for a detailed description of the approach used to set forecast ranges for the Prediction Markets.

Another effect of the normalized pricing strategy is that it reduces the value of short selling the end stocks in the Prediction Market because the price of those shares is significantly reduced from the average pricing approach; thus a participant selling the end markets short only stands to gain \$4.42 with normalized pricing rather than the \$11.11 from the average pricing approach. In both cases, we expected the normalized pricing to incent participants to invest in the stocks that they believed would end up spanning the actual deliveries for the Prediction Market.

3.2.8 Setting Investment Caps in a Prediction Market

In a perfect world, Prediction Markets would operate without regulation using Adam Smith's invisible hand. Unfortunately, in the real world, it turns out that Mangold was right; participants will manipulate Prediction Markets in order to win. A powerful tool in preventing manipulation is setting an investment cap for each Prediction Market.

By preventing participants from investing all of their money in a single market we suspected that investment caps would push participants to invest in Prediction Markets where they were not experts. Having uninformed participants (Wolfers & Zitzewitz 2004) should improve the performance of the Prediction Markets by adding transaction volume as described in

section 3.1.2. During our Prediction Market pilot, we discovered that, if given a chance, participants will manipulate Prediction Markets. One of the participants in the pilot was able to manipulate the market to such an extent that he was able to double his money by investing all of it in a single market and then convincing others to buy shares in other stocks within the same market. In order to combat this, we needed to follow the advice provided by Hanson et al; prevent one participant or small group of participants from having sufficient funds to manipulate the Prediction Market.

Armed with this learning, we set a \$15,000 investment cap for each Prediction Market. With the increase from six to nine stocks in each market, a participant investing \$15,000 in a single market would only be able to move the price of a stock by two dollars. Because each participant could only move the market by two dollars, it would take a group of thirty-eight participants to create the same \$75 stock price increase we observed in the pilot Prediction Markets. Given that the only department participating in the study that had thirty-eight participants was the demand planning department (see Table 3.2.2), and they were the sponsors of the Prediction Market study, we felt it highly unlikely that all of the participants would collude to manipulate stock prices in order to win.

3.2.9 Setting Trading Hours in a Prediction Market

There is significant debate in the Prediction Market community regarding trading hours; some researchers feel that having open trading hours encourages participation while others believe that having limited trading hours encourages participation. For our Prediction Markets, we implemented limited trading hours.

There were two reasons that we implemented trading hours for our Prediction Market study: input from Intel and the results of our Prediction Market pilot. During our interview with

Intel, Hopman described how it used trading hours to limit the amount of time that participants would spend “fiddling” with Prediction Markets rather than working; Intel felt that having trading hours encouraged participants to spend a focused and consistent time in the markets rather than nervously checking the value of their portfolios. In addition to Intel’s input, we also observed that open trading hours promoted price manipulation in our pilot Prediction Markets.

By having the Prediction Markets open all of the time, sophisticated participants were able to take advantage of participants who only accessed the market on an infrequent basis; the sophisticated participants would note price changes and then place trades at the end of the day to benefit by driving prices up further or selling short to drive prices down. During our debriefing session many of the participants who accessed the market on an infrequent basis felt that they were at an unfair disadvantage because the active traders were able to profit from their lack of activity; the sophisticated traders mirrored this sentiment stating that they were able to profit by jumping in to the market and taking advantage of their less active counterparts.

As a result of our interview with Intel and the results of the Prediction Market pilot, we determined that we would have our Prediction Markets open for one twenty-four hour period each week during our study; we validated that this approach would provide sufficient time with the pilot participants. Based on their input we were confident that having weekly trading hours would provide sufficient time for both active and inactive traders to participate in the Prediction Markets without one group having an unfair advantage due to time.

3.2.10 Providing Training for Prediction Market Participants

Perhaps the most interesting revelation from the Prediction Market pilot was that several participants did not take part in the pilot because they did not understand how Prediction Markets operated and did not feel comfortable buying and selling shares of stock. As a result of this

feedback, we developed training materials to provide participants with the background they would need to effectively take part in our study. The training we prepared concentrated on three capabilities required for effective participation: how to access the Prediction Market, how to buy and sell shares and how to implement trading strategies.

To illustrate how to access Prediction Markets and how to buy and sell stocks, we provided participants with a simple five step process. By providing participants with an overview of how the Prediction Market process worked, they were able to understand how they should take part in the study. In addition to the process, we also provided screen shots coupled with hands on training to make sure that every participant knew how to execute trades and manage their portfolios. At the end of the training process we conducted a survey to ensure that the participants felt comfortable accessing the Prediction Markets; based on the results 75% of them did.

4 Methods for Analyzing Prediction Market Data

Studies of Prediction Markets in controlled academic settings were focused on simulating and learning about attributes of Prediction Markets, for example information aggregation was studied by Hanson, et al (2006). In Prediction Markets where the public participates, the studies were focused on accuracy of predictions and the speed of information discovery. A few studies have been done using Prediction Markets for business decision making. In their study of Prediction Markets at HP for sales forecasting, Chen and Plott (2004) used forecast accuracy and bias to assess the results. Since we had access to the transaction history from the Prediction Market and the identity of all the traders, we performed quantitative analysis on the Prediction Market data and qualitative analysis on data from surveys and interviews of market participants. These

analyses were done to find out if Prediction Markets can develop more accurate forecasts than General Mills current forecasting process.

4.1 Mapping Current Planning Process

We mapped the planning process at General Mills to help understand how they forecast and to choose the appropriate participants for the study. Our understanding of the planning process was developed through interviews with the Demand Planning group and the sales regions, and review of internal documentation.

4.1.1 Operations Forecasting Process

The Demand Planning team at General Mills is responsible for generating the Operations Forecast. The Operations Forecast from the demand planning department is used for execution by manufacturing and supply chain departments. The demand planning process at General Mills is comprehensive and takes into account three factors; customer forecasts from the customer sales regions; consumer insights from the marketing department using Nielsen market research data; and, long term and short term trends from historical data. The Demand Planning team is organized along product categories and product groups. The demand planning function is centralized across all customers and is located in General Mills headquarters. Key customers have dedicated demand planners who work with the customer key account managers at the sales region offices. The demand planning process starts with a bottom-up forecast by product stock keeping unit ('SKU') and week using statistical models. This forecast is aggregated to product group and month levels, and combined with promotional up-lifts determined by the Sales teams. This process also accounts for consumer insights from the Nielsen market research data provided by the Marketing team. The demand planning, marketing, sales and customer service departments participate in the consensus process and agree on a point forecast at the product

group and month level. This forecast is referred to as the Operations Forecast and is disaggregated down to product SKU, week and distribution center for execution. The following figure presents the high level forecasting process at General Mills.

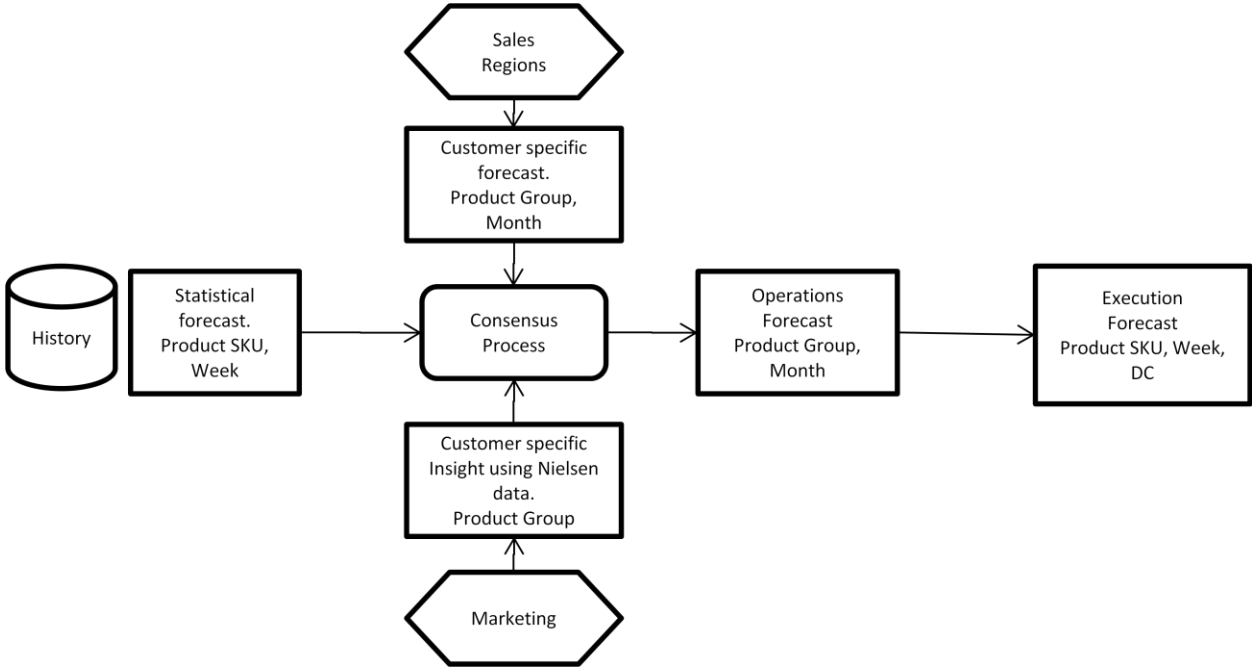


Figure 4.1.1 - Forecasting Process at General Mills

4.1.2 Sales Forecasting Process

General Mills sales regions are dedicated to working with individual customers. The sales regions establish relationships with the customer, plan store promotions in collaboration with the customer and provide detailed forecasts to other departments in General Mills. Sales regions prepare bi-annual sales plans based on guidance from the corporate marketing department. The sales regions plans include detail of promotional design and the costs of promotions. After the plans are approved by Marketing, the Sales teams collaborate with the customers to adapt the plan to meet customer’s needs. This process is completed four months prior to execution in the store. During the next four months, the sales regions update the promotion plans and the

expected forecast on a regular basis in discussions with the customer. The sales region offices are usually located near the customer’s headquarters and are organized into teams by product categories. Each team is responsible for one or more product categories. Each customer key account manager, aided by business process associates is responsible for detailed promotional plans, customer interaction and sales forecasts for one product category. Other associates in this team are responsible for category management and execution functions. The customer key account managers provide the forecasts to the demand planning department and participate in the consensus forecasting process on an as needed basis. The following figure represents the sales organization in a typical region.

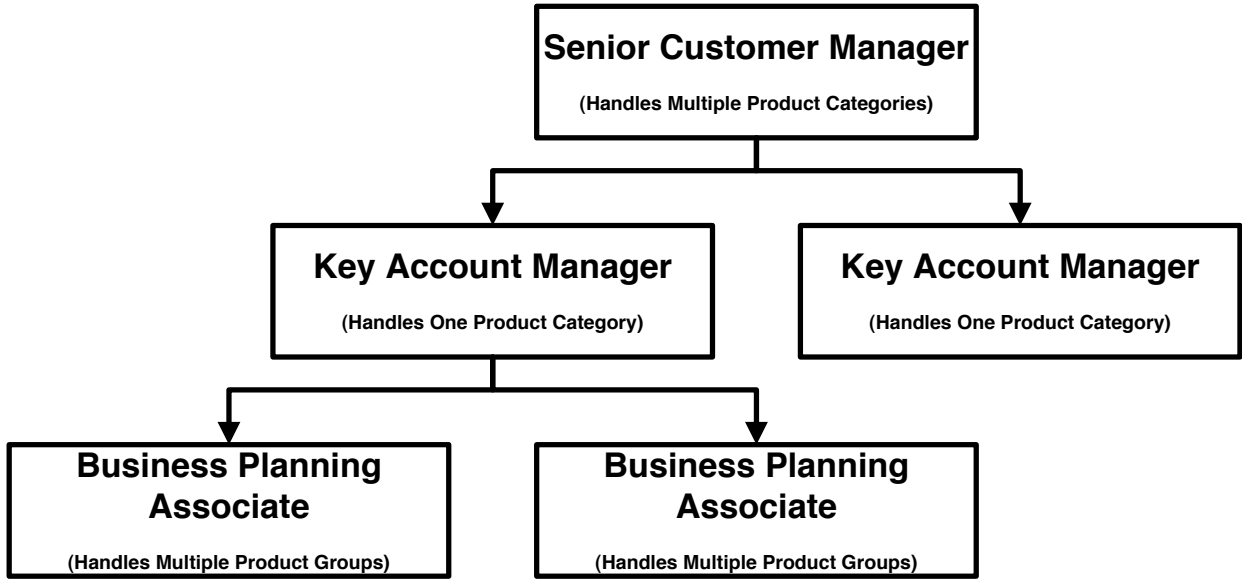


Figure 4.1.2 - Organizational Structure for General Mills Sales Regions

4.2 Quantitative Analysis of Data

Analysis of the prices of the stocks along with the quantity that was bought and sold was central to determining if Prediction Markets can generate better results than the current General Mills forecasting process. General Mills was interested in exploring Prediction Markets for business decision making in the context of Operations Forecasts. The applicability of Prediction Market

Forecasts to operational planning largely depends on forecast accuracy; equally important are aspects like, speed of information revelation as measured by Prediction Market shifts, probability expressed by the Prediction Market in each of the forecast ranges, as measured by the prices on a weekly basis, and the private information that winners possessed, as expressed by the stocks they invest in. We studied the Prediction Markets at both the individual Prediction Market level and at aggregate levels. Aggregate levels were chosen based on the type of Prediction Market question. Corporate questions were broad Prediction Market questions that enlisted participation from across General Mills; GroceryCo questions focused on promotions at GroceryCo; and BoxCo questions focused on a mix of general and category specific forecast at BoxCo. The data sources for this analysis came from Prediction Markets, the Operations Forecast, Sales Forecast and a naïve forecast using prior year actuals. The following table summarizes the categories for data analysis.

Data Category	Department	Frequency of update
Sales Forecast	Customer Sales Regions	As available
Operations Forecast	Demand Planning	Weekly
Year Ago Actual	Demand Planning	Once
Prediction Market Forecast	Market Consensus	Weekly

Table 4.2.1 - Table of Data Categories Used in Analysis

4.2.1 Prediction Market Data Description

The Prediction Market data was gathered using internet-based Prediction Market software provided by ConsensusPoint. The software vendor was chosen because of their prior relationship with General Mills and for their ability to:

1. Operate the Prediction Market at specified times. Business requirements suggested that the Prediction Market be open for one day each week.
2. Set ceilings on investments. To encourage participation and prevent manipulation, the investment that was possible in a single Prediction Market needed regulation.

3. Extract detailed transaction data. For our analysis we needed detailed transaction data from the trading activity.
4. Manage the right level of anonymity using ‘Leader Boards’. Showing ranking of participants in leader boards encouraged competition and maintained anonymity.

Selected employees from sales, customer support, finance, marketing and demand planning departments were setup as participants. A total of 20 Prediction Markets were setup as part of this study. Each Prediction Market asked a specific question; for example, “What will the total January deliveries be for US Retail?” The answers to this question were broken into 9 forecast ranges, each corresponding to an individual stock, see Figure 3.2.2 for details. The forecast ranges were centered on the Operations Forecast.

The price of a stock reflected the probability of the actual being in the forecast range as spanned by the stock. The price was influenced by the number of traders interested in the stocks, the quantity of stocks purchased and most importantly on the belief in the other stocks in the Prediction Market, see Figure 3.2.7 for details. The Prediction Market software ensured that the total price of all the stocks in a Prediction Market always equaled one hundred dollars. As explained earlier in the methodology section, the initial price of the stocks in a Prediction Market was normally distributed with the highest price for the stock corresponding to the forecast range which included the Operations Forecast. Once trading began the price of the individual stocks fluctuated and revealed the Prediction Market’s forecast. The software recorded four states of the stock price in each Prediction Market.

1. The initial prices at which the stocks of a Prediction Market opened for trading.
2. The price that a buyer or seller paid for the stock for each trade.
3. The closing price of stock at the end of each day.

- 4. The final price of the stock when the Prediction Market was closed and a winner was judged.

In addition to the four states mentioned above, the historical trading data captured the traders involved and the number of shares traded in each transaction. This historical record was captured from the moment the Prediction Market opened until it was finally judged and closed. This historical record allowed us to study a participant’s trade in detail. The following is a logical description of the main data elements that were used in our analysis.

The trading history presents a historical time line of trading activity for individual traders and the Prediction Market as a whole. This allowed us to study the individual traders and identify shifts in sentiments and contrast the behavior of winners versus the others in any Prediction Market.

The trading history recorded all trading activity: selling, buying and short selling. In addition to the stock and its price, the history included the trader who initiated the transaction, the date and time of the transaction and the number of shares purchased. The following table shows a sample trading history record.

Column Name	Description	Example
History Transaction ID	Unique number for each transaction	1278
Trader ID	Unique identifier corresponding to the participant	123 refers to John Doe
Stock ID	Unique identifier corresponding to the stock being traded	ABC-06
Quantity	The number of stocks being traded. Positive number for buy and negative number for sell	25 or -54
Price	The price at which the trade occurred	\$32.66
Other Owners	Number of other participants who owned the same stock	22

Table 4.2.2 - Description and Sample of Trading History

To study the overall market forecast we needed a snapshot of the price of all the stocks in a Prediction Market each week. The daily stock summary compiled the final closing price of every stock for each day. This summary price was used to study the confidence of the Prediction

Markets and determine the forecast accuracy on a weekly basis. The following table shows a sample daily stock summary record.

Column Name	Description	Example
Stock	Unique ID of the stock	ABC-06
Calculation Date	The date and time at which the snapshot was taken (one per day)	2009-02-19 00:00:00
Ending Price	The price of the stock at the end of the day	\$35.64

Table 4.2.3 - Description and Sample of Daily Stock Summary

Having access to a combination of detail and summary information enabled us to understand participant and market behavior.

4.2.2 Prediction Market Price Interpretation

As previously described, each Prediction Market answers one forecasting question. The price of the stocks in the Prediction Market expresses the confidence of the Prediction Market corresponding to the forecast range. Interpreting the implied forecast based on the price of the stock is illustrated below.

Consider a Prediction Market with a ticker symbol ABC corresponding to the question “How many cases of Product A will be sold in March 2009”. For ease of understanding let us assume that there were five stocks in this Prediction Market. At the end of each day, the daily stock summary table would record the price of each stock. The stock price represents the probability expressed by the market in each forecast range. The average forecast from the forecast range is multiplied by the implied probability to get the implied forecast from each range. The sum of all implied forecasts from all the ranges gives the Prediction Market Forecast. The equation for deriving the forecast from the prices in the Prediction Markets is shown below.

$$E(\text{Market Forecast}) = \sum_{i=1}^n P(\text{Forecast Range}) \times \text{Avg}(\text{Forecast Range})$$

Equation 4.2.1 - Equation for Deriving Prediction Market Forecasts

The following table illustrates this with an example.

Stock Symbol	Forecast Range of cases sold in March 2009	Stock Price in dollars	Implied Probability	Average forecast from the range in cases	Implied forecast in cases
ABC-01	1 – 10	\$10	10%	5.5	0.55
ABC-02	11 – 20	\$10	10%	15.5	1.55
ABC-03	21 – 30	\$20	20%	25.5	5.10
ABC-04	31 – 40	\$50	50%	35.5	17.75
ABC-05	41 – 50	\$10	10%	45.5	4.55
Prediction Market Forecast					29.5

Table 4.2.4 - Illustration to Derive Prediction Market Forecast from Prices

As seen above, a Prediction Market Forecast takes into account the confidence expressed by participants in each of the forecast ranges. By attributing a probability, the Prediction Market Forecast effectively aggregates opinions from diverse participants. Stock ABC-04 with a \$50 price shows that there is a 50% probability that the forecast will fall in the middle of the range 31 and 40 (35.5), yielding 17.75 cases (50% * 35.5). Repeating this process for all the stocks in the Prediction Market and adding the results gives an expected forecast of 29.5 cases.

4.2.3 Definition of MAPE for Measuring Error

The accuracy of a Prediction Market Forecast is an indicator of the effectiveness of the Prediction Market for forecasting purposes. The accuracy of a forecast, referred as forecast error measured the difference between forecast and actual. The lower the forecast error, the more accurate the forecast. While there are different ways to measure forecast error, like, Mean Percent Error (MPE), Root Mean Square Error (RMSE), Mean Absolute Percent Error (MAPE) etc, MAPE was chosen as the metric for measuring forecast error for the following reasons.

1. General Mills uses MAPE as its error measure for forecast accuracy.
2. MAPE is very easy to calculate.
3. MAPE can be measured at aggregate levels.
4. MAPE can be used to compare the accuracy of different datasets.

MAPE was calculated using the following formula.

$$MAPE = \frac{\sum_{i=1}^n Absolute(Forecast_i - Actual_i)}{\sum_{i=1}^n (Actual_i)}$$

Equation 4.2.2 - MAPE Calculation Equation

As shown above in Table 4.2.4 the Prediction Market Forecast is 29.5 cases, let us assume that the actual is 30 cases and that the Operations Forecast is 27. The following table illustrates the MAPE calculation for both the Prediction Market and Operations Forecast.

Prediction Market Forecast in cases	Operations Forecast	Actual	Prediction Market Forecast MAPE	Operations Forecast MAPE
29.5	29	30	= ABS(29.5 – 30)/30 = 1.6%	= ABS (29 – 30)/30 = 3.3%

Table 4.2.5 - MAPE Calculation Illustration

4.2.4 Prediction Market Groups for Analysis

Analyzing groups of markets and comparing the results is an important step to understanding if one Prediction Market is better than another. To infer the organizational, business strategy and process realities that differentiate one group from the other, we compared the aggregate forecast accuracy across Prediction Markets and participant groups. Working with General Mills we determined the need for three types of groups:

1. The first grouping of Prediction Markets were based on question type. For example if the Prediction Market question was an overall corporate question it was assigned to the corporate category.
2. The second grouping of Prediction Markets was classified based on type of forecast: volume forecast, promotional forecast and category forecast. For example if the question focused on promotional activity it was assigned to the promotional grouping.
3. The third grouping of Prediction Markets grouped all Prediction Markets that closed in March.

4.2.5 Analyzing Overall Results

Our study of the Prediction Markets compared the Prediction Market MAPE with the MAPE from the Operations Forecast. The analysis also included gathering attributes from the Prediction Markets such as the number of participants, duration, Prediction Market type and department of the winner. The following table shows a description of the metrics and an example of this study.

Column Name	Description	Example
Market Ticker	Symbol for the market	Q3, PONEQ3 etc.
Market Type	Type of market	Corporate, BoxCo etc.
Forecast Type	Type of forecast	Volume, Promotional etc.
Duration	Duration in number of weeks	1, 5, 10 etc.
Number of participants	The number of participants who traded in the Prediction Market	29
Highest Stock Match Actual Shipments (Yes/No)?	If the forecast range corresponding to the highest priced stock included the actual when the market was judged and closed	Yes
Top 3 Included Actual Shipments (Yes/No)?	If the forecast ranges corresponding to the top three highest priced stocks included the actual when the market was judged and closed	Yes
Winning Department	The department that the winner belonged to	Sales, customer support etc.
Operations Forecast MAPE	The MAPE between the Operations Forecast and the actual	3%
Prediction Market MAPE	The MAPE between the Prediction Market Forecast and actual	2%

Table 4.2.6 - Description and Example of Overall Results from Prediction Markets

4.2.6 Analyzing Forecast Accuracy

The in-depth analysis compared the MAPE from the Prediction Market Forecast, Operations Forecast, Sales Forecast and actual from last year. The comparison was done using a plot based on the data captured each week. While the Prediction Market Forecast and Operations Forecast changed each week based on new information, the Sales Forecast and prior year actual remained the same. By using MAPE we were able to examine the accuracy at any aggregate level.

Comparisons at aggregate levels allowed us to identify the groups that performed better than others. The weekly comparison allowed us to identify if the Prediction Market provided any timing advantages by revealing information swiftly and identified any lead or lag between the forecasts.

Consider the following table showing the MAPE on a weekly basis for the Prediction Market question “How many cases of Product A will be sold in March 2009”.

Forecast Generation Date	20-Jan	28-Jan	5-Feb	12-Feb	19-Feb	Average	Standard deviation across time
Prediction Market MAPE	3.00%	4.50%	3.00%	2.00%	1.60%	2.82%	1.12%
Operations Forecast MAPE	2.00%	2.50%	3.50%	4.50%	3.30%	3.16%	0.96%
Sales MAPE	6.00%	6.00%	6.00%	6.00%	6.00%	6.00%	0.00%
LY Actual MAPE	7.50%	7.50%	7.50%	7.50%	7.50%	7.50%	0.00%

Table 4.2.7 - Weekly MAPE data for Prediction Market “How many cases of Product A will be sold in March 2009”

The Table 4.2.7 allows us to answer the following questions

- Was the Prediction Market Forecast better than the Operations Forecast?
- Did the Prediction Market lead the Operations Forecast?
- Did the Prediction Market identify market shifts sooner?

4.2.7 Analyzing Market Activity

Market activity as indicated by the volume of trading is an indicator of information flow.

Prediction Markets, much like their real world counterparts can get caught in bubbles. This happened when participants did not have information or had mis-leading information and traded on the wrong forecast ranges. We conducted a quantitative study using the measures such as flow (Lee) and confidence to determine if General Mills Prediction Markets exhibited bubble behavior.

1. Flow: Research showed that Prediction Markets revealed information sooner than traditional methods. This implied that participants will be motivated to trade on information as soon as they had it. Studying the volume of stocks traded on a weekly basis allowed us to map the flow of information in the Prediction Markets. By definition, Prediction Markets with consistent volume of trades during all weeks of trading implied

that new information was available to the participants on a regular basis. On the contrary if the trading only happened in a few spikes, then the information was only available a few times to the participants. Flow was measured as the cumulative percentage of stocks since the beginning of the market until the week being measured to the total number of stocks traded in the market. Flow was split into buy and sell trades. The equation for computing flow was as shown below:

n represents the number of weeks and m represents current week

$$Flow = \frac{\sum_{i=1}^m \text{Number of Stocks}}{\sum_{i=1}^n \text{Number of Stocks}} \%$$

Equation 4.2.3 - Equation for Computing Flow

The following table presents the flow for the Prediction Market “How many cases of Product A will be sold in March 2009”

Forecast Generation Date	20-Jan	28-Jan	5-Feb	12-Feb	19-Feb
Cumulative Buy	50%	60%	90%	95%	100%
Cumulative Sell	10%	10%	60%	100%	100%
Ratio of Buy to Sell	5:1	5:1	4:1	4:1	3:1

Table 4.2.8 – Transaction flow for Prediction Markets

The interpretation of this table helped answer the following questions:

- Did the market trading stay relatively constant during the weeks of trading?
- Was the spike in trading related to new information available to participants?
- Did the ratio of Buy-Trades to Sell-Trades decrease, indicating a shift in market perception?

2. Market Confidence: Just like in the real world stock markets, Prediction Markets were equally prone to “irrational exuberance”, when the participants traded vigorously on the wrong stocks because of wrong information or mistaken beliefs.

The confidence plot shows the probability expressed by the Prediction Market for each forecast range. The actual forecast range plus one range on either side was chosen to measure the accuracy of the market with a wider forecast range. The confidence in a forecast range is measured by the sum of the prices of the corresponding stocks in these three ranges.

Using the example of stock ABC for the Prediction Market question “How many cases of Product A will be sold in March 2009”, the market confidence is expressed in the following graph.




Figure 4.2.1 – Prediction Market Confidence

By examining Figure 4.2.1 we can see that the market was initially less confident in the forecast as indicated by the decreasing height of the lightly shaded rectangle in weeks 01-21, 01-30 and 02-05. The plot also indicates that the market changed course between 02-12 and 02-19.

To paint a complete picture we examined if the Prediction Market converged on a single price. To perform this analysis, the following questions were examined:

- a. Was the Prediction Market Forecast becoming accurate with time? This was measured using MAPE between the expected Prediction Market Forecast each week and the actual.
- b. Was the Prediction Market Forecast converging? The convergence implied that more people in the market agreed on a narrow forecast range.

The Prediction Market convergence was also referred as price convergence and was computed as a coefficient of variation (COV) for each week with respect to the expected market forecast.

$$COV = \frac{\sigma}{E(\text{Market Forecast})}$$

Where $E(\text{Market Forecast})$ is the forecast from the Prediction Market as laid out in section 4.2.2 and standard deviation σ is the weekly Prediction Market Forecast and is measured as follows

$$\sigma = \sqrt{\sum_{i=1}^{fcst\ Range} P(\text{Fcst Range}) \times (\text{Avg}(\text{Fcst Range}) - E(\text{Market Fcst}))^2}$$

Equation 4.2.4 - Equation of COV for Prediction Market Forecast

The following table shows the standard deviation calculated against the actual and the Prediction Market Forecast on each week.

Forecast Generation Date	COV	MAPE
20-Jan	73%	3%
28-Jan	80%	5%
5-Feb	78%	3%
12-Feb	79%	2%
19-Feb	69%	2%

Table 4.2.9 - Table of Coefficient of Variation and MAPE

The plot for convergence for the above table is as shown below.




Figure 4.2.2 - Graph of convergence Coefficient of Variation and MAPE

In this case we see that the Prediction Market is becoming more accurate each week as shown by the decreasing MAPE. The market however does not converge until the last week as shown by the high coefficient of variation before 02-19.

The interpretation of the plot allows us to conclude:

- If Prediction Market converged on the correct forecast?
- If the Prediction Market did not converge on the right forecast but became increasingly confident in its prediction each week?

- If the Prediction Market did not converge on a forecast at all?

4.2.8 Analyzing Trader Behavior

While being interviewed for our project, Prof. Doug Thomas, associate professor of Supply Chain at Penn State, stated that the biggest advantage of Prediction Markets was the honest revelation of information. We studied the Prediction Market for how different participants revealed information by their trades in the Prediction Markets. For the purposes of this study the participants in each Prediction Markets were categorized into:

- Winner – one per Prediction Market who has won the most money in that Prediction Market.
- Losers – Participants who made no money participating the Prediction Market since they did not own stocks in the correct forecast range (Note: Since the number of short sell trades in the Prediction Market was very low, participants who made money only by shorting stocks were still included in the losers category).
- Others – Participants who owned stocks in the winning Prediction Markets but did not own enough to become the winner.

This study compared the proportion of stocks invested in winning and non-winning claims by the three categories of users. This analysis showed if the winners bought into the winning stock from the very first week of trading and how the proportion of their portfolio changed during the weeks of trading. Based on this analysis we were able to conclude if the winners had private information that guided them to the stock choices.

Another area that we examined was the performance of the different job functions and departments that the traders belonged in order to understand if any group had information that would yield a better forecast. We used the following definitions for our analysis.

1. Job Function represents the role held by the participant within General Mills.
2. Number of Prediction Markets Participated In represents the number of Prediction Markets that the functional group owned shares in.
3. Forecast Right Trader Count represents the count of the number of traders who held a position in the right stock, the stock that spanned the actual shipments, when the Prediction Market closed.
4. Forecast Wrong Trader Count represents the count of the number of traders who held a position in the wrong stocks, the stocks that did not span the actual shipments, when the Prediction Market closed.
5. Total Traders adds the right and wrong trader counts together. It is important to note that this total count can overstate the number of participants in the Prediction Market because a single participant may own shares in both the right and the wrong stock.
6. Percent of Traders Right represents the percent of total traders who owned shares in the stock that spanned the actual shipments for the Prediction Market.
7. Average Right MAPE represents the accuracy of the Prediction Market Forecast for the stock that spanned the actual shipments (see section 4.2.3 and 4.2.5 for details on the calculations).
8. Average Wrong MAPE represents the average forecast error for stocks that did not span the actual shipments (see section 4.2.3 and 4.2.5 for details on the calculations).
9. Volume in Right Stock represents the net shares held in the right stock that spanned the actual shipments.
10. Volume in Wrong Stocks represents the net shares held in the wrong stocks that did not span the actual shipments.

11. Percent of Volume in Right Stock represents the percentage of shares that were in the stock that spanned the actual shipments.

The groups (see Table 3.2.2) were then compared against each other to determine the accuracy of their forecasting results.

Job Function	Total Traders	Forecast Right Trader Count	Forecast Wrong Trader Count	Percent Of Traders Right	Average Right MAPE	Average Wrong MAPE	Volume Traded In Winning Stock	Volume Traded In 8 Losing Stocks	Percent Of Volume Traded In Winning Stock
Customer Service Center	15	7	8	47%	0.05%	16%	13,392	30,403	31%

Table 4.2.10 - Job Function Forecast Summary

This table provided the ability to compare the forecast accuracy of multiple groups to determine the one that had the most accurate forecast.

4.3 Qualitative Analysis

Qualitative analysis was done to study the process and people dimensions of the Prediction Market. The analysis focused on the participants’ motivation to trade, sources of information and strategy for trading. Understanding participation was important to gauge the role of incentives in motivating trading. This was achieved through two surveys and in-depth interviews with selected participants.

4.3.1 Analyzing Participants Using Surveys

Surveys were used to gather feedback from the participants in the Prediction Markets. The surveys focused on participation and helped in understanding if any self selection was involved in the participants’ choice of Prediction Markets. Surveys also revealed the motivation behind participation, knowledge about the three types of Prediction Markets and their sources of information. The survey was taken online using third party software and was sent to the email of all the participants.

The first survey was sent to all the participants in the Prediction Market. The first survey was done after the first week of trading. The purpose of the first survey was to get an early read for the how the participants felt and more importantly un-cover any issues with the Prediction Markets in a timely manner. The initial survey concentrated on questions surrounding the ease of use of software, the training that was provided and participation.

The second survey was conducted after the Prediction Markets were closed for trading in early April. The second survey was used as a post mortem to understand certain noticeable traits, such as, the larger than expected number of uninformed traders in some Prediction Markets and the apparent lack of accuracy in many Prediction Markets. The second survey was more detailed than the first survey and was aimed at understanding the time spent by participants and their motivation for trading. The second survey also asked participants about their self selection criteria, the sources of their information and finally recommendations for appropriate use of the Prediction Market.

4.3.2 Analyzing Participants Using Interviews

Individual phone interviews were conducted with twelve selected participants who were categorized based on their transaction history in the Prediction Markets. These interviews were done to gain in-depth perspective on their trading behavior, their current role in the forecasting process, their sources of information, the process they applied to evaluating the Prediction Markets they wanted to participate in and the strategies they employed for trading on the stocks.

These interviews also helped us understand the motivation of the individuals to participate in the Prediction Markets, the time constraints that the trading imposed and the incentive structure that excited them. These interviews also provided a rare glimpse into why

participants chosen to join in markets where they had no insight on and how they sought information and processed the information from the Prediction Markets.

The participants formed a representative set chosen based on their participation history in the Prediction Market. For the purpose of choosing the persons to interview, we categorized the participants into eight categories as listed below:

Participant Category	Number of Interviewees
Participants who won the most with the least participation.	1
Participants who won the most.	6
Participants, who picked the right forecast the most number of times, yet did not win the Prediction Markets.	2
Participants who participated the most but neither won nor got the forecast right.	2
Participants who participated moderately but never won	2
Participants with the least participation	3
Participant who changed their positions in the Prediction Markets the most.	1

Table 4.3.1 - Categories of Participants for Interviews

5 Findings

Answering the question “are Prediction Markets appropriate for business forecasting?” is a complex affair. Prediction Market success or failure is determined by the interplay of people, the process they manage and the data they work with. The following sections present in-depth findings of forecast accuracy, trader behavior and market activity coupled with other quantitative and qualitative metrics. To aid in our analysis we classified the markets, first by customer type into Corporate, BoxCo and GroceryCo, second by the type of forecast into promotional, category and volume and a third category of only those Prediction Markets that ended in March. These findings answer many questions about the applicability of Prediction Markets and helped form our conclusions discussed in the next section.

5.1 Analyzing Overall Results

The following Table 5.1.1 summarizes the results from all the Prediction Markets.

Market Ticker	Market Type	Forecast Type	Duration (Weeks)	Participants	Highest Stock Matched Actual Shipments (Yes/No)	Top 3 Stocks Included Actual Shipments (Yes/No)	Winning Department	Prediction Market MAPE	Operations Forecast MAPE
PONEQ3	Corporate	Prod. Cat.	5	38	Yes	Yes	BoxCo Sales	0.58%	0.85%
JAN	Corporate	Volume	1	28	No	Yes	CSC-GroceryCo	0.77%	1.03%
MAR	Corporate	Volume	10	28	No	No	BoxCo Sales	1.22%	1.08%
Q3	Corporate	Volume	5	43	No	Yes	BoxCo Sales	0.46%	0.30%
PTWQO3	Corporate	Prod. Cat.	5	29	No	No	BoxCo Sales	6.21%	3.46%
GROCAR	Grocery Co	Promo	5	12	No	No	No Winner	17.27%	20.88%
GROCARPSEVEN	Grocery Co	Promo	5	14	No	No	GroceryCo Sales	3.71%	1.86%
GROCARPSIX	Grocery Co	Promo	5	10	Yes	Yes	GroceryCo Sales	9.06%	8.36%
GROCARPFIVE	Grocery Co	Promo	5	9	No	Yes	GroceryCo Sales	30.04%	11.43%
GROMAR	Grocery Co	Volume	10	17	No	Yes	GroceryCo CSC	6.15%	13.00%
GROQ3	Grocery Co	Volume	5	24	Yes	Yes	GroceryCo Sales	3.22%	1.73%
GROPFOUR	Grocery Co	Promo	5	15	Yes	Yes	GroceryCo Sales	10.83%	10.59%
GROPTWO	Grocery Co	Promo	10	21	No	No	GroceryCo Sales	35.06%	23.95%
GROPTHREE	Grocery Co	Promo	10	17	No	Yes	GroceryCo Sales	6.02%	7.05%
BOXPONEQ3	BoxCo	Prod. Cat.	5	26	Yes	Yes	BoxCo Sales	2.08%	1.90%
BOXJAN	BoxCo	Volume	1	18	Yes	Yes	BoxCo Sales	1.90%	2.21%
BOXMAR	BoxCo	Volume	10	23	No	Yes	BoxCo Sales	7.75%	9.52%
MAXPTWO	BoxCo	Prod. Cat.	5	23	Yes	yes	BoxCo Sales	4.44%	8.01%
BOXQ3	BoxCo	Volume	5	34	No	Yes	BoxCo Sales	1.90%	1.76%
BOXPTHREE	BoxCo	Prod. Cat.	5	20	No	No	Supply Chain	6.26%	6.38%
							Average	7.75%	6.77%

Table 5.1.1 – Overall results from Prediction Markets

We performed a correlation and regression analysis between each of the Prediction Market MAPE's and the Operations Forecast MAPE in table 5.1.1. The Prediction Market MAPE and the Operations Forecast MAPE share an 83% correlation with an R² of 0.7. The high correlation combined with high R² indicates that the Prediction Market Forecast and the Operations Forecast virtually move in unison. This is further confirmed by the average MAPE across all markets as shown above in Table 5.1.1.

Markets in which the public is allowed to participate are based on aggregating the opinion of participants to a broad question, this is also true of the questions used by BestBuy in their Prediction Markets (Jaedike). However the questions in our experiment were aimed at revealing and using information hidden in the organization to arrive at a better forecast. The dichotomy between opinion and information based markets is proven by regression with the duration of the Prediction Markets and number of participants as independent variables, and

Prediction Market MAPE as dependent variable. The low R^2 value of 0.32 shows that the number of participants cannot explain the variance in Prediction Market MAPE adequately.

A quick glance through the accuracy of the Prediction Market by the market type indicates that GroceryCo has the highest MAPE ranging from 1.73% to 23.95%. The promotional nature of GroceryCo's business contributes to the variability in the forecast. In contrast the more predictable business at BoxCo and the high aggregation level of the corporate Prediction Market questions contribute to their high accuracy. This is explained in detail in Section 5.2.

A powerful contradiction that we encountered was that Sales teams won the most number of Prediction Markets (refer to Table 5.1.1), yet as a group, Sales did not perform as well as the Demand Planning group (refer to 5.2.7). The implication of this is that in the absence of a definite way to identify the sales persons who have insight into the forecast, it is better to have a focused Demand Planning team.

5.2 Analyzing Forecast Accuracy

The ultimate measure of any forecasting solution is the accuracy of the forecasts it provides. To measure forecast accuracy we used MAPE (see sections 4.2.3 and 4.2.5) as our primary accuracy measure. In working through our findings, we discovered that there is a 69% correlation between the Operations and Prediction Market Forecasts based on MAPE.

5.2.1 Market Type Forecast Accuracy

The first grouping we used to analyze the data was along customer lines. Table 5.2.1 presents the market type definitions.

Corporate Prediction Markets		BoxCo Prediction Markets		GroceryCo Prediction Markets	
Forecast Type	Ticker	Forecast Type	Ticker	Forecast Type	Ticker
Volume	Q3	Volume	BOXPONEQ3	Promotional	GROCAR
Volume	MAR	Volume	BOXJAN	Promotional	GROWCARPSEVEN
Volume	JAN	Product Category	BOXMAR	Promotional	GROCARPSIX
Product Category	PONEQ3	Volume	BOXPTWO	Promotional	GROCARPFIVE
Product Category	PTWOQ3	Product Category	BOXQ3	Promotional	GROMAR
		Product Category	BOXPTHREEQ3	Promotional	GROQ3
				Promotional	GROPTFOUR
				Volume	GROPTWO
				Volume	GROPTHREE

Table 5.2.1 – Market Type Definitions

Table 5.2.2 summarizes the Prediction Market, Operations, Sales and Last Year (LY) Actual Forecasts using MAPE.

Market Type	Accuracy	01-01	01-21	01-30	02-05	02-12	02-19	02-26	Average
Corporate Prediction Markets	Market MAPE	0.72%	0.65%	0.69%	0.54%	0.93%	0.79%	0.79%	0.73%
	Operations MAPE	0.71%	0.71%	0.36%	0.45%	0.54%	0.67%	0.84%	0.61%
	Sales MAPE	1.16%	1.16%	1.33%	1.33%	1.33%	1.33%	1.33%	1.28%
	LY Actual MAPE	5.48%	5.48%	5.04%	5.04%	5.04%	5.04%	5.04%	5.17%
BoxCo Prediction Market	Market MAPE	4.75%	3.71%	3.95%	3.77%	3.06%	3.01%	3.01%	3.61%
	Operations MAPE	4.86%	4.86%	4.39%	4.10%	3.44%	2.85%	2.85%	3.91%
	Sales MAPE	5.45%	5.45%	6.05%	6.05%	6.05%	6.05%	6.05%	5.88%
	LY Actual MAPE	9.74%	9.74%	8.77%	8.77%	8.77%	8.77%	8.77%	9.05%
GroceryCo Prediction Markets	Market MAPE	7.27%	7.34%	7.26%	5.56%	4.17%	3.64%	3.64%	5.56%
	Operations MAPE	7.55%	7.55%	8.04%	7.03%	6.14%	3.95%	3.95%	6.32%
	Sales MAPE	6.44%	6.44%	6.37%	6.37%	6.37%	6.37%	6.37%	6.39%
	LY Actual MAPE	20.27%	20.27%	20.27%	20.27%	20.27%	20.27%	20.27%	20.27%

Table 5.2.2 – Market Type MAPE Comparison

With the exception of LY actual, the MAPE for all of the forecasts is less than 7%. It is interesting to note that corporate Prediction Markets have the most accurate forecasts with a MAPE of less than 1%; General Mills provides household staple foods that are consumed every day and so in aggregate they are very predictable. BoxCo is next with a MAPE of less than 4%; BoxCo uses an everyday low price strategy for the products it sells. GroceryCo has the worst performance with a MAPE of less than 6%; GroceryCo uses high/low pricing for merchandising

the products it sells. We can see from Table 5.2.2 that business strategy has a significant impact on forecast accuracy.

In order to understand whether or not the average forecasts in Table 5.2.2 were different from each other, we analyzed the standard deviation of the MAPE for each forecast generation week. While the Prediction Market MAPE presented in Table 5.2.3 is on average 0.33% better across all markets than the Operations Forecast, we find that the operations MAPE and Prediction Market MAPE are virtually the same based on the standard deviation data presented in Table 5.2.3.

Market Type	Accuracy	Average	Standard Deviation
Corporate Prediction Markets	Market MAPE	0.73%	0.12%
	Operations MAPE	0.61%	0.17%
	Sales MAPE	1.28%	0.09%
BoxCo Prediction Market	Market MAPE	3.61%	0.64%
	Operations MAPE	3.91%	0.87%
	Sales MAPE	5.88%	0.29%
GroceryCo Prediction Markets	Market MAPE	5.56%	1.74%
	Operations MAPE	6.32%	1.72%
	Sales MAPE	6.39%	0.04%

Table 5.2.3 - Standard Deviation of Market Type MAPE

From this table we are able to conclude that the Prediction Market Forecast and the operations MAPEs are almost identical because they are within one standard deviation of each other. This leads us to conclude that while the Prediction Market Forecast is on average 0.33% more accurate; this difference is not significant enough to allow us to describe the Prediction Market Forecast as more accurate than the Operations Forecast.

5.2.2 Forecast Type Forecast Accuracy

Another data slice of the Prediction Markets that we examined was to view the markets based on the type of forecast being generated: volume, product category and promotional. The tickers that make up each forecast type are listed in Table 5.2.4.

Volume Prediction Markets	
Market Type	Ticker
Corporate	Q3
Corporate	MAR
Corporate	JAN
GroceryCo	GROQ3
GroceryCo	GROMAR
BoxCo	BOXJAN
BoxCo	BOXQ3
BoxCo	BOXMAR

Product Category Prediction Markets	
Market Type	Ticker
Corporate	PONEQ3
Corporate	PTWOQ3
BoxCo	PTHREEQ3
BoxCo	BOXPONEQ3
BoxCo	BOXPTHREEQ3

Promotional Prediction Markets	
Market Type	Ticker
GroceryCo	GROCAR
GroceryCo	GROCARPSEVEN
GroceryCo	GROCARPSIX
GroceryCo	GROCARPFIVE
GroceryCo	GROPFOUR
GroceryCo	GROPTWO
GroceryCo	GROPTHREE

Table 5.2.4 – Forecast Type Prediction Markets

Table 5.2.4 shows that volume and category Prediction Market cut across multiple market types as defined in Table 3.2.3 while the promotional Prediction Markets are limited exclusively to GroceryCo. By grouping the Prediction Markets in this manner we were able to isolate the type of forecast being developed and measure the MAPE for each of the forecast type groups.

Through our conversations with General Mills, we discovered that volume Prediction Markets were forecasting underlying data that was very stable because there was little promotional activity. Product Category Prediction Markets were forecasting data that was less stable due to moderate national promotional activity. Promotional Prediction Markets were forecasting volatile data because General Mills initiated new promotional activity surrounding a major special event. By understanding the data type being forecast we were able to compile the results presented in Table 5.2.5.

Forecast Type	Accuracy	01-01	01-21	01-30	02-05	02-12	02-19	02-26	Average
Volume Prediction Markets	Market MAPE	1.51%	1.41%	1.46%	1.15%	1.39%	1.21%	1.21%	1.33%
	Operations MAPE	1.55%	1.55%	1.26%	1.34%	1.26%	1.29%	1.44%	1.38%
	Sales MAPE	1.76%	1.76%	2.01%	2.01%	2.01%	2.01%	2.01%	1.94%
	LY Actual MAPE	5.99%	5.99%	5.28%	5.28%	5.28%	5.28%	5.28%	5.48%
Prod. Cat. Prediction Markets	Market MAPE	3.12%	1.96%	2.21%	2.01%	1.51%	1.46%	1.46%	1.96%
	Operations MAPE	3.17%	3.17%	1.69%	1.30%	1.30%	0.94%	0.94%	1.79%
	Sales MAPE	4.38%	4.38%	4.38%	4.38%	4.38%	4.38%	4.38%	4.38%
	LY Actual MAPE	10.51%	10.51%	10.51%	10.51%	10.51%	10.51%	10.51%	10.51%
Promotional Prediction Markets	Market MAPE	7.65%	9.34%	9.99%	14.43%	11.55%	12.86%	12.86%	11.24%
	Operations MAPE	6.24%	6.24%	23.81%	19.48%	15.82%	5.80%	5.80%	11.89%
	Sales MAPE	10.14%	10.14%	9.73%	9.73%	9.73%	9.73%	9.73%	9.84%
	LY Actual MAPE	44.58%	44.58%	44.58%	44.58%	44.58%	44.58%	44.58%	44.58%

Table 5.2.5 – Forecast Type Prediction Market MAPE

From Table 5.2.5 we see that the majority of the forecast error comes from promotional activity; promotional Prediction Markets have an average MAPE of 11% while the other markets have an average MAPE of 2%. One item that does stand out is that the Sales group appears to have some insight based on the accuracy of their promotional forecasts; however, this may be coincidental because the Sales Forecast performed poorly in all other data segments. Section 5.2.3 provides detail for this claim by segmenting the Prediction Market Forecasts by job function.

Once again, we can see that the Prediction Market and Operations Forecast have similar levels of forecast accuracy. This can be seen by examining the standard deviations in Table 5.2.6.

Forecast Type	Accuracy	Average	Standard Deviation
Volume Prediction Markets	Market MAPE	1.33%	0.14%
	Operations MAPE	1.38%	0.13%
	Sales MAPE	1.94%	0.12%
Prod. Cat. Prediction Markets	Market MAPE	1.96%	0.59%
	Operations MAPE	1.79%	0.98%
	Sales MAPE	4.38%	0.00%
Promotional Prediction Markets	Market MAPE	11.24%	2.37%
	Operations MAPE	11.89%	7.67%
	Sales MAPE	9.84%	0.20%

Table 5.2.6 – Forecast Type Standard Deviation of MAPE

From this table we are able to conclude that the Prediction Market Forecast and the Operations Forecast have similar forecast accuracy because the MAPEs are within one standard deviation of each other. Thus with the results from Tables 5.2.5 and 5.2.6 we conclude that the Prediction Market Forecast and Operations Forecast have the same level of forecast accuracy.

5.2.3 Functional Forecast Accuracy

One question that is often asked, in any forecasting process, is which group has the most accurate forecasts. For a complete description of the table columns please refer to section 5.2.3. As part of the analysis we classified each participant in the forecasting process (see Table 3.2.2) by job function so that we could measure the MAPE of each group. Table 5.2.7 lays out our findings.

Job Function	Number of Prediction Markets Participated In	Total Traders	Forecast Right Trader Count	Forecast Wrong Trader Count	Percent of Traders Right	Average Right MAPE	Average Wrong MAPE	Volume in Right Stock	Volume in Wrong Stocks	Percent of Volume in Right Stock
Customer Service Center	20	15	7	8	47%	0.07%	23%	11,992	49,816	19%
Finance	14	5	2	3	40%	0.05%	26%	2,165	7,795	22%
GroceryCo Product Sales Manager	14	5	2	3	40%	0.03%	20%	557	10,612	5%
GoceryCo Sales	16	14	6	8	43%	0.04%	21%	8,769	35,728	20%
Marketing	6	5	2	3	40%	0.08%	8%	593	848	41%
Product Sales Manager	10	2	1	1	50%	0.05%	13%	1,219	5,000	20%
Corporate Sales Management	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Demand Planning	20	23	9	14	39%	0.06%	11%	3,183	39,719	7%
BoxCo Product Sales Manager	12	16	8	8	50%	0.07%	18%	17,995	38,191	32%
BoxCo Sales	12	14	7	7	50%	0.11%	16%	20,750	25,260	45%
Total	14	99	44	55	44%	0.06%	17%	67,223	212,969	24%

Table 5.2.7 – Job Function Prediction Market MAPE

The Number of Prediction Markets Participated In column illustrates how effectively participants self select Prediction Markets they feel they have information about (Berg et al. 2008). We compared Table 3.2.2 and Table 5.2.7 and saw that the number of actual participants varied from 50% to 80% of the potential participants which was in line with what we expected based on the Prediction Market pilot. Customer Service and Demand Planning were the only groups who participated in all of the Prediction Markets; this makes sense because both groups are organized to manage the entire General Mills portfolio of products across all customers.

Many forecasters question if Marketing has any information to add to the planning process. Marketing participated in only 30% (6 of 20 Prediction Markets) versus the other groups who participated in an average of 80% of the Prediction Markets (16 of 20 Prediction Markets). The number of shares owned by Marketing shows that it did not have a strong commitment to its forecast; the average group held 35,000 shares while Marketing held 1,441. The Prediction Markets exercise was aimed at understanding information flow between Demand Planning and

Sales teams, Marketing team was only invited to participate, this explains the lack of participation from Marketing

Another conclusion that can be drawn from Table 5.2.7 is that Demand Planning has the best overall results across the broadest range of Prediction Markets. This can be seen in Demand Planning's higher than average holdings of 43,000 shares versus the average of 34,000 shares across all Prediction Markets; the average forecast accuracy is also better at 11% versus an average of 17% across all markets. The data in Table 5.2.7 suggest that the forecasting process and organizational structure that General Mills uses (see section 4.1) to develop its Operations Forecast is funneling the right information into Demand Planning enabling them to develop more accurate forecasts.

5.2.4 Location Forecast Accuracy

We explored how location affected forecast accuracy to determine if information was being transferred to headquarters or if the data remained in the regions; in performing this analysis we hoped to determine whether the regions or headquarters developed more accurate forecasts. For a complete description of the table columns please refer to section 5.2.3. We excluded the plant from consideration because one person purchased 100 shares in 12 separate markets which skewed our analysis.

Location	Number Of Prediction Markets Participated In	Total Traders	Forecast Right Trader Count	Forecast Wrong Trader Count	Percent Of Traders Right	Right MAPE	Wrong MAPE	Volume Traded In Winning Stock	Volume Traded In Losing Stocks	Percent Of Volume Traded In Winning Stock
Basset Creek Office	23	11	5	6	45%	0.08%	17%	6,233	37,311	14%
Headquarters	23	27	10	17	37%	0.04%	8%	4,134	38,823	10%
GroceryCo Region	16	27	12	15	44%	0.05%	24%	17,564	70,159	20%
BoxCo Region	12	32	16	16	50%	0.08%	17%	39,660	65,128	38%
Total	19	97	43	54	44%	0.06%	17%	67,591	212,421	24%

Table 5.2.8 – Location Prediction Market MAPE

Just as in Table 5.2.7 we can see that General Mills forecasting process is yielding benefits because the participants located at Headquarters have the best overall forecasting results with 8% MAPE versus the 17% average. We conclude that headquarters personnel are able to synthesize data from the various groups and use it to develop more accurate forecasts.

5.2.5 Forecast Type Transaction Forecast Accuracy

In 2008 Berg et al. explored the information aggregation properties of Prediction Markets and discovered them to be significant. Before examining Table 5.2.9 please review section 4.2.1; our version of Berg et al.’s analysis is presented in Table 5.2.9. The difference between Table 5.2.5 and Table 5.2.9 is that Table 5.2.5 weights each forecast range by the probability (price) to get the expected forecast to aggregate all participant transactions while Table 5.2.9 uses the raw transactional data without the benefit of probability (price) weighting to aggregate the information.

Group	Average Wrong MAPE from Table 5.2.5	Average Wrong MAPE	Percent of Volume in Right Stock
Volume Prediction Markets	1.33%	10%	25%
Prod. Cat. Prediction Markets	1.96%	26%	30%
Promotional Prediction Markets	11.24%	1,073%	6%
Total	4.88%	370%	23%

Table 5.2.9 – Forecast Type Prediction Market MAPE

The ability of Prediction Markets to aggregate data is most obvious in the Promotional Prediction Markets where the Prediction Market is able to aggregate information in such a way that the Prediction Market has an 11.24% MAPE while the transaction level data has a 1,073% MAPE. From these findings we conclude that Prediction Markets have the ability to aggregate large amounts of transactional data to develop more accurate forecasts based on input from a large number of participants.

5.2.6 March Forecast Accuracy

The final area that we explored was the ability of Prediction Markets to forecast for the longer term. When we configured the Prediction Markets there were five markets that examined the ability of Prediction Markets to make longer term forecasts. Table 5.2.10 presents the results of the analysis looking at the Prediction Markets that were judged in March: MAR, GROMAR and BOXMAR, GROPF0VR and GROPTHREE.

Accuracy	01-01	01-21	01-30	02-05	02-12	02-19	02-26	03-05	03-12	03-19	Average	Standard Deviation
Market MAPE	3.15%	3.46%	2.75%	2.54%	2.68%	2.19%	2.19%	3.91%	3.77%	2.26%	2.89%	0.65%
Operations MAPE	3.37%	3.37%	3.01%	3.51%	3.52%	3.63%	4.16%	4.16%	4.16%	2.46%	3.53%	0.54%
Sales MAPE	2.75%	2.75%	2.75%	2.75%	2.75%	2.75%	2.75%	2.75%	2.75%	2.75%	2.75%	0.00%
LY Actual MAPE	5.14%	5.14%	5.14%	5.14%	5.14%	5.14%	5.14%	5.14%	5.14%	5.14%	5.14%	0.00%

Table 5.2.10 – March Prediction Market MAPE

On first blush, Table 5.2.10 suggests that we should use Sales’ estimates to address all long term forecasting because the Sales Forecast has the lowest MAPE; this would be a bad conclusion given the information presented in Table 5.2.7 and Table 5.2.8 where we saw, over a

broad range of Prediction Markets, the centralized planning team and not Sales had the most accurate forecasts.

Table 5.2.10, however, does present an interesting scenario for Prediction Markets; we have seen that the Prediction Market and Operations Forecast are nearly identical, though the Prediction Markets are generally more accurate than the Operations Forecasts, from a forecasting perspective. However, in this case, we see that the standard deviation does not span the average suggesting that the forecasts are in fact different. We believe further study is required to make a definitive statement on Prediction Markets ability to develop long-range forecasts.

5.3 Analyzing Market Activity

Analyzing the market activity from the detail transactions helped us understand the flow of information in different groups of markets. The flow of information influenced the trader behavior resulting in varying accuracies across the different Prediction Markets.

5.3.1 Flow for Forecast Type Prediction Market Groups

The timing of when participants buy and sell shares provided visibility into when information was available. Flow in Table 5.3.1 presents the cumulative percent of Buy-Trades and Sell-Trades on a weekly basis by forecast type. The flow of the Buy-Trades and Sell-Trades indicates the frequency with which new information is available to the participants. If the information reinforced current market trends, then the market consolidated by buying and when the information revealed new insights, the market shifted position by selling. Thus the ratio of number of Buy-Trades to Sell-Trades in each week helps illustrate the amount of turnover in the markets predictions.

Forecast Type	Transaction Type	01-21	01-30	02-05	02-12	02-19
Volume Prediction Markets	Buy Stocks	70%	76%	86%	89%	100%
	Sell Stocks	46%	58%	65%	73%	100%
	Ratio of Buy to Sell	6:1	5:1	5:1	5:1	4:1
Promotional Prediction Markets	Buy Stocks	54%	66%	79%	90%	100%
	Sell Stocks	44%	45%	60%	77%	100%
	Ratio of Buy to Sell	4:1	5:1	4:1	4:1	3:1
Category Markets	Buy Stocks	65%	74%	87%	94%	100%
	Sell Stocks	45%	52%	61%	76%	100%
	Ratio of Buy to Sell	6:1	6:1	6:1	5:1	4:1

Table 5.3.1 - Flow for Forecast Type Prediction Markets

The overall pattern of trading between the three groups is similar indicating that similar information was available across all markets. The comparison of buy flows between the three market groups on the first week of trading shows that participants had the most confidence in the information for the Volume Prediction Markets and bought 70% of total volume. Participants expressed least confidence in the information they had for the promotions market by buying only 54% of the total volume. An interesting variation is the difference in ratio of Buy-Trades to Sell-Trades, on the first week of trading, between the Promotional Prediction Markets at 4:1 and the other groups at 6:1. This indicates that lack of credible information drives market shifts and helps explain forecast accuracy discussed in section 5.2.

5.3.2 Flow for Market Type Prediction Market Groups

The flow of Prediction Markets by market type (Corporate, GroceryCo and BoxCo) is presented in Table 5.3.2. The variation in the ratio of Buy-Trades to Sell-Trades in the opening week is even more obvious between the groups. Corporate markets at 8:1 show the highest confidence in information while GroceryCo at 4:1 shows the least confidence in the information. This indicates

that GroceryCo markets do not have reliable information leading to market shifts and decreased forecast accuracy discussed in Section 5.2.

Market Type	Transaction Type	01-20	01-30	02-05	02-12	02-19
Corporate Markets	Buy Stocks	65%	74%	90%	93%	100%
	Sell Stocks	41%	51%	68%	84%	100%
	Ratio of Buy to Sell	8:1	7:1	7:1	5:1	5:1
GroceryCo Markets	Buy Stocks	64%	72%	84%	93%	100%
	Sell Stocks	52%	59%	69%	83%	100%
	Ratio of Buy to Sell	4:1	4:1	4:1	4:1	3:1
BoxCo Markets	Buy Stocks	68%	75%	81%	87%	100%
	Sell Stocks	38%	46%	49%	57%	100%
	Ratio of Buy to Sell	6:1	6:1	6:1	5:1	3:1

Table 5.3.2 - Flow for market type Prediction Markets

5.3.3 Flow for March Prediction Markets

The March Prediction Markets present surprising results regarding availability of timely information. The flow as shown in Table 5.3.3 does not reach 65% until 03-05, 7 weeks after the markets opened. Also, the ratio of Buy-Trades to Sell-Trades at 4:1 was low indicating lack of confidence in the information all through the trading period.

Market Type	Transaction Type	01-21	01-30	02-05	02-12	02-19	02-26	03-05	03-12	03-19
March Markets	Buy Stocks	24%	29%	33%	36%	42%	42%	65%	89%	100%
	Sell Stocks	18%	19%	20%	22%	36%	36%	47%	72%	100%
	Ratio of Buy to Sell	4:1	4:1	5:1	5:1	3:1	3:1	4:1	3:1	3:1

Table 5.3.3 - Table Flow for March Prediction Markets

5.3.4 Convergence and Confidence in Volume Prediction Markets

The ability to converge is important to the success of the Prediction Market, the question often asked is do they converge on the right number? The findings from confidence and convergence will highlight the underlying causes that help or hinder the convergence of Prediction Markets.

The Figure 5.3.1 shows the charts of confidence and convergence for the volume Prediction Markets. The confidence in the Prediction Market increases each week reaching 70%

by the closing week. The increase in confidence is accompanied by a decrease in the MAPE and coefficient of variation leading to better forecast accuracy and price convergence. Thus with each week of trading the Volume Prediction Markets are increasingly confident in predicting the actual.




Figure 5.3.1 Confidence and Convergence for Volume Prediction Markets

Thus we see confidence increases, MAPE decreases and price converges, which helps explain why category Prediction Markets are accurate.

5.3.5 Convergence and Confidence in Promotional Prediction Markets

The convergence and confidence chart for the promotional Prediction Markets in Figure 5.3.2 presents a different picture than that of the Volume Prediction Markets. It is interesting because the Prediction Market shows a moderately high 60% confidence on the actual + 2 range during the trading period; however the Prediction Markets are unable to converge within this forecast range. This leads to lower forecast accuracy as indicated in the convergence charts.

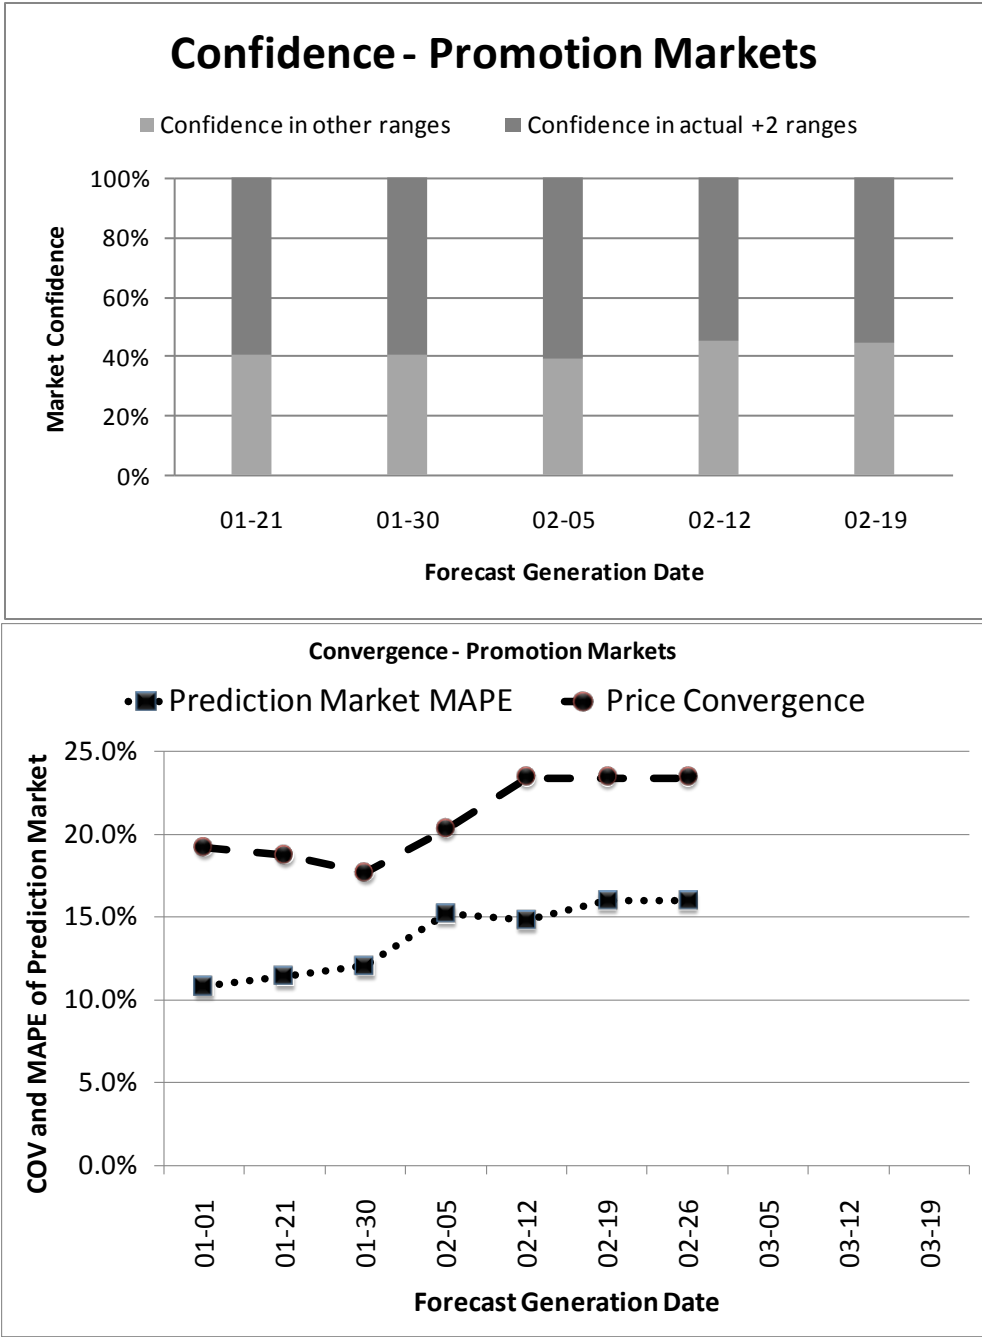


Figure 5.3.2 - Confidence and Convergence in Promotional Prediction Markets

Thus we see that confidence stays flat, MAPE increases and price does not converge which helps explain why promotional Prediction Markets are the least accurate.

5.3.6 Convergence and Confidence in Category Prediction Markets

The convergence and confidence charts for the category Prediction Markets shown in Figure 5.3.5 has the best behavior of all the market groups studied. The confidence shows increases with each week to almost 70% at market closing. The convergence chart shows that the MAPE is reduced from 3.5% to 1.2% during this time.

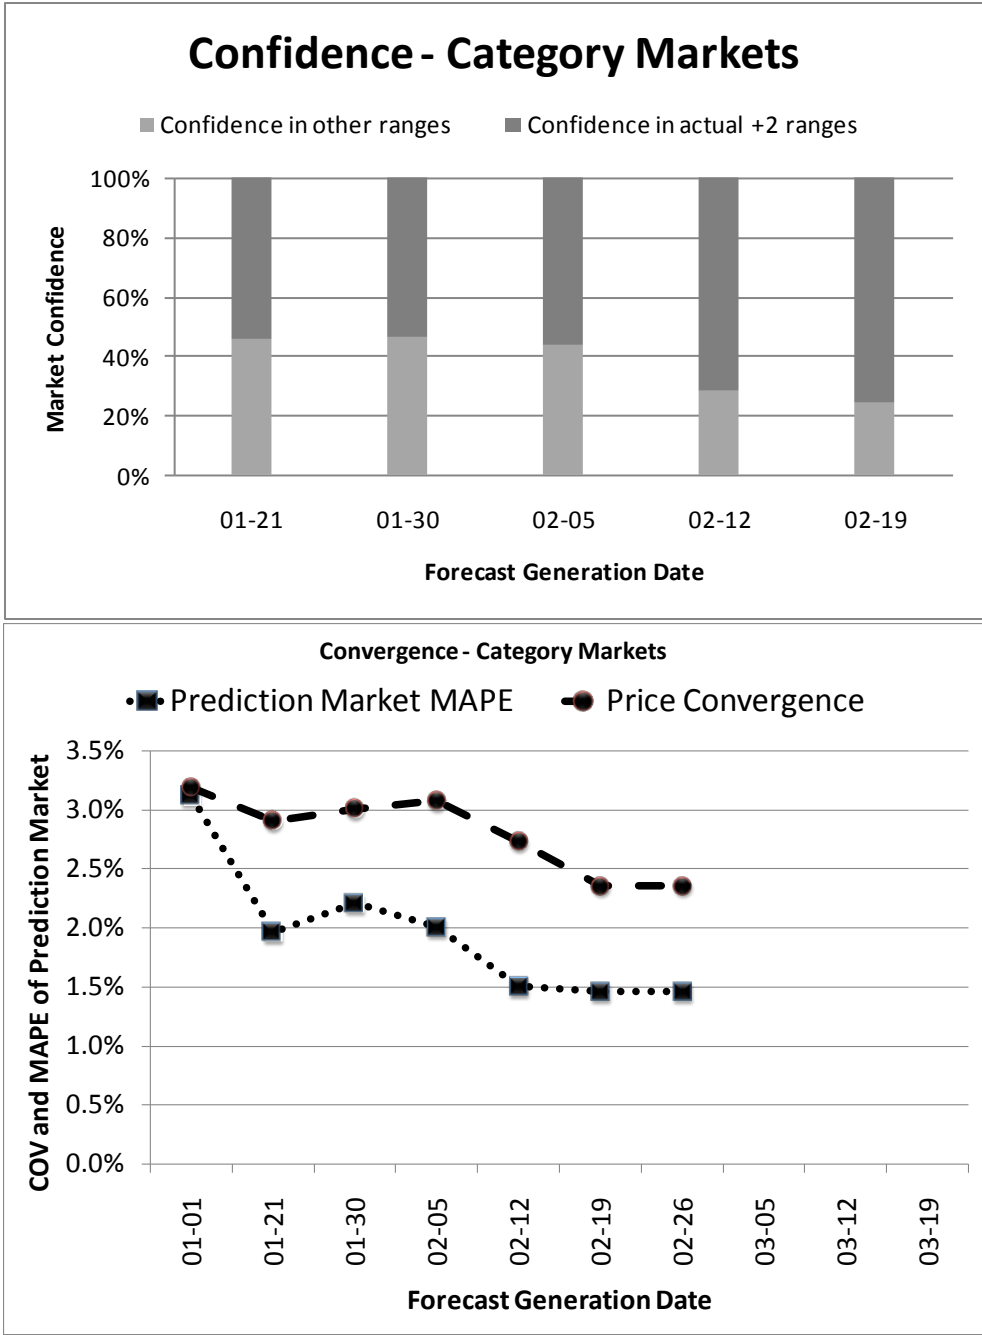


Figure 5.3.3 - Confidence and convergence in category Prediction Markets

Thus we see confidence increases, MAPE decreases and price converge which helps explains why category Prediction Markets are so accurate.

5.3.7 Confidence and Convergence of March Markets

The Figure 5.3.4 shows the confidence and convergence charts for the March Prediction Markets. The charts indicate that the confidence remained unchanged until active trading began in early March. The market shifted course in the flood of information available before getting back on track with reduced MAPE and better convergence in the closing week.




Table 5.3.4 - Confidence and Convergence in March Prediction Markets

Thus we see that confidence decreases and increases over time and MAPE and price convergence remain flat which explains why March markets maintain a stable level of accuracy over time.

5.4 Analyzing Trader Behavior

By joining a Prediction Market, traders reveal information and through their stocks reveal their belief in the forecast. Detailed analyses of the transactions show how the information about forecasts differs within participants of a Prediction Market. The following section analyzes the trader behavior with respect to the stocks they own, the information they possess and their confidence in that information.

5.4.1 Trader Behavior in Volume Prediction Markets

Figure 5.4.1 shows the chart for analyzing trader behavior for the volume Prediction Markets. Please refer to section 4.2.7 for more details about the definition of trader groups. The volume Prediction Markets included questions that were broad and thus enlisted large participation. The chart below shows that the “other” group of traders traded in a similar fashion to the winners.




Figure 5.4.1 - Trader Behavior in Volume Prediction Markets

The trading behavior of winners and “others” points out that both groups had similar information about the volume Prediction Markets.

5.4.2 Trader Behavior in Promotional Prediction Markets

The trader behavior chart for promotional Prediction Markets shows that the information is not dispersed through the organization; Figure 5.4.2 below shows that the proportion of winning and non-winning stocks held by the winner and the ‘other’ category differ significantly with the other category holding a smaller proportion of the winning stock.

One interesting observation is that winners hold the winning stock exclusively in the first and last week of trading but shifted their position in the intermediate weeks. This shows that the winners either had additional information that prompted them to trade multiple stocks or were swayed by the markets because of lack of confidence in the winning stocks.




Figure 5.4.2 - Trader Behavior in Promotional Prediction Markets

Thus we see that winners are able to recognize and respond to new information differently than other traders. What these charts do not reveal is whether the winners had misleading information in the weeks 01-30 through 02-12 or if they mis-interpreted the information. Further research is required to understand this topic.

5.4.3 Trader Behavior in Category Prediction Markets

Figure 5.4.3 below shows the trader behavior for the category Prediction Markets. It shows that winners exclusively held the winning stock throughout the trading session. The plot also indicated that the trading pattern differed between the winners and the ‘other’ group. The other group had significantly smaller proportion of the winning stock; this shows that the participants in the ‘other’ category were less informed about the forecasts.




Figure 5.4.3 - Trader Behavior in Category Prediction Markets

The overwhelming and complete confidence in the winning stock indicate that the winners are very sure about the Category Prediction Markets which also reflects in the very high accuracy of these markets.

5.4.4 Trader Behavior in March Markets

Figure 5.4.4, below, shows the trader behavior for the Prediction Markets that closed in March. The chart below shows that the winners held their stocks while the losers and ‘other’ category continued trading. The market as a whole became active in March suggesting that new

information became available to participants. The fact that winners changed their positions in March and that March markets were no more accurate than rest of the Prediction Markets indicates that new information did not improve forecast accuracy. However from a process perspective this indicates that information needed to get accurate forecasts do not become available until the last week of the Prediction Market.




Figure 5.4.4 - Trader Behavior in March Prediction Markets

The very high proportion of non-winning stocks held by winners in the March markets is a question that needs further research.

5.5 Analyzing Participants Using Surveys

In addition to our quantitative analysis we wanted to gain an understanding of trader behavior beyond the numbers tracked by the Prediction Market. We conducted two surveys to gauge trader sentiment towards Prediction Market forecasting.

5.5.1 Analysis of First Market Survey

The first survey helped us understand the motivation behind participation, the convenience of market timings and the usefulness of the training provided. The survey was sent to 176 participants with 36 respondents.

- To a question on how easy the trading software was on a scale on 1 to 5, 1 being hard and 5 being easy, 73% of the survey respondents said that it was very easy to understand and use. Our research and pilot had indicated that ease of use of software was important to encourage and sustain participation. Our choice of software provider was partly driven by the ease of use and the survey confirmed our choice.
- 95% of the respondents mentioned that information about the events was the main reason for participating.
- To a question on how well the training helped, over 95% of the survey respondents rated the training as 3 and above on a scale of 1 to 5. This confirmed that participants regarded training as very important in enabling them to use the Prediction Market software effectively.
- To a question on how much knowledge they had on each of the three groups of Prediction Markets, 33% of the respondents claimed that they had significant knowledge of BoxCo 23% of the respondents claimed that they had significant knowledge of GroceryCo and 39% of the respondents claimed that they had significant knowledge of Corporate Prediction Markets.
- A total of 13 survey respondents, who did not participate in the Prediction Markets, indicated that market timing was the key reason for not participating in the markets.

If General Mills wishes to implement Prediction Markets on an ongoing basis, training and simplicity will be crucial to its success.

5.5.2 Analysis of Final Market Survey

The second survey was sent to 170 participants with 48 respondents.

- For a question on the time taken to participate, 60% of the survey respondents mentioned that the Prediction Markets took less than 10 minutes and another 13% responded that the Prediction Market took between 10 and 20 minutes. A low time commitment is needed for a successful rollout of the program on a larger and regular basis.
- For a question on the knowledge that the participants had on the different Prediction Market types, 60% of survey respondents ranked themselves 3 and above on a scale of five in terms of knowledge about the BoxCo events. To a similar question on GroceryCo markets, only 39% of respondents rated themselves 3 and above. This was similar to the response in the first survey and indicated the asymmetry in knowledge between the GroceryCo market participants and the BoxCo market participants.
- 47% of the survey respondents agreed that they participated in Prediction Markets where they had very little insight. Of these, 40% mentioned that excitement at being part of the experiment was the reason for participating in markets where they had no insight.
- While participating in markets they had no insight, 42 % guessed the forecast based on the prices in the markets and 39% used internal General Mills tools to make decisions. In addition participants mentioned that they spoke with other users who had information and even asked customers. A significant majority of people claimed to have intimate knowledge of the events.
- 47% of the respondents felt Prediction Markets was appropriate for promotional forecasting while 50% felt it was appropriate for new product introductions.

- While a small number of participants were not convinced that the Prediction Markets are useful, 40% of the participants said they would participate if this was a regular process. Also 50% of the respondents expressed the need to know how the results will be used.
- 3% of the respondents agreed that a survey would be a good substitute for Prediction Markets.

The second survey highlighted the sources of information that participants used and underscored the ability of the Prediction Market to arrive at an accurate forecast while not taking much time from each of its participants. While participants were ready to use the Prediction Markets as a process, they demanded to know about the use of information from the Prediction Markets.

5.6 Analyzing Participants Using Interviews

Interviews were done with selected participants to probe deeper into questions on participation, motivation, incentives, source of information etc. The interviews gave us an opportunity to interact and learn from the participants especially with respect to their strategy for trading, concerns etc. The interviews were conducted one on one over the phone. The following bullet points summarize the information we learned from the 12 interviews with the selected participants.

- **Trading Hours:** The participants liked the dedicated trading hours which allowed them to focus on the markets. Some participants missed participating because of other work and travel commitments and requested that the trading hours be extended to several days each week. All participants were thankful for the weekly reminder email sent by the project sponsor to announce the opening of trading markets.

- **Investment Cap:** All the participants we interviewed had participated in markets they had no insight on. The investment cap forced them to invest in markets they were not comfortable with and greatly increased the liquidity and accuracy of the markets. None of the participants complained of not having enough money.
- **Market Ranges in the stocks:** Most people were comfortable with the forecast range definitions on the stocks. A very few participants raised concerns about heavy Promotional Prediction Markets in which it was very difficult to predict which exact range the forecast would fall in.
- **Initial Normalized prices** did not cause concern among the participants. At least one participant deliberately chose to bet in the second highest priced stock so as to maximize his profits in case that stock included the actual.
- **Leader Boards:** The reaction to leader boards ranged from obsession to ignorance. Some participants were focused on staying in the top 10 leader board while many did not bother. At least one participant used the leader-board to gain insight into the markets based on the expertise of the persons on the leader-board. Most participants felt the leader board across all markets was good enough with no need for detailed leader boards per market or group of markets.
- **Trader Anonymity:** The software provided the right level of anonymity. While it was possible to know the participants and their net-holding it was not possible to know which markets they participated and what stocks they were holding. Participants mentioned that knowing who was winning increased their motivation and drove them to try and improve their results.

- Software trading interface: Excepting one interviewee who was not familiar with trading and did not attend the training, others were comfortable with the trading software.
- Incentive: The cash incentive currently was not a significant motivation to participate. Most participants would consider a non financial incentive like lunch with a senior executive favorably.
- Prediction Markets Buzz: Many participants indicated that they were involved in Prediction Market related chatter with their peers although no one mentioned discussing their strategy for the markets.
- Self selection. All participants recognized the need to possess unique information and participate in the market as early as possible to win the markets. No participant tried their luck in a completely unrelated market, for example a GroceryCo sales manager participating in BoxCo markets. Participants who had surplus money chose to invest in related markets where they could guess intelligently.
- Information Sources. While most participants relied on the sources they used during regular business, some went out of the way to learn from other sources. For these participants, the Prediction Market had an added benefit of broadening their horizon.
- Next Steps: Most participants were eager to participate in the experiment and mentioned they would participate regularly if they knew what the results were and how the results were going to be used.

6 Prediction Market Conclusions

In many ways this study of Prediction Markets has been eye opening; when we started the project we expected Prediction Markets to outperform the current General Mills forecasting process by a significant margin. As the findings section has shown, this is clearly not the case.

We think that proponents who claim Prediction Markets outperform experts and polls are overselling the benefits. However, Cain and Drakos, from Gartner, are premature in placing Prediction Markets, “in the trough of disillusionment.” Our findings clearly show that Prediction Markets are capable of developing very accurate forecasts, effectively aggregate information from multiple participants and may be able to provide improvement for long-range forecasting.

There are several factors that help Prediction Markets to succeed. Financial incentives bring participants to the Prediction Markets, but recognition and “bragging rights” are required to maintain participation. Finally, participants must have a very clear picture of how their input will be used for the benefit of the business; without this critical communication it is likely that Prediction Markets will lose participation once the initial novelty wears off. For a Prediction Market to be successful in the long-term, it needs to be championed with significant executive support. If these factors are taken into account Prediction Markets have the potential to improve forecast accuracy and increase communication throughout the organization.

6.1 Prediction Market Forecast Accuracy

We discovered that our Prediction Markets underperformed the Operations Forecast by 0.1% (7.75% MAPE vs. 6.77% MAPE); see Table 5.1.1. However, we do not believe the difference to be statistically significant because the variability of both Prediction Market and Operations Forecast MAPE span both forecasts in less than one standard deviation. Thus we would conclude that the MAPE of the Prediction Markets was equivalent to the Operations Forecast.

The forecasts generated by the Prediction Markets that closed in March show that Prediction Market Forecasts outperformed the Operations Forecast by 0.64% and that there was not overlap within one standard deviation; Table 5.2.10 clearly illustrates this point. Further

study is required to determine if General Mills can use Prediction Markets to improve the accuracy of its long-range forecasts.

6.2 Prediction Market Information Aggregation

In 2008 Berg et al.'s re-examined the Iowa Election Markets and found that Prediction Markets were able to outperform polls through information aggregation; these results are clearly corroborated by the comparison presented in Table 5.1.1. Understanding the difference between the transaction and daily summary data allows us to conclude that the Prediction Market is able to aggregate information. From our research, we recommend that General Mills use Prediction Market in any situation where a large number of participants must be brought together and a formal process for aggregating their opinions does not yet exist.

6.3 Participant Self Selection

Hanson et al. found that Prediction Market participants will select markets where they believe that they have information; our research confirms these findings. First, a comparison of Table 3.2.2 and Table 5.2.7 shows that 50% to 80% of the participants will buy or sell shares in a given market. Second, follow up surveys revealed that 70% of the participants who bought and sold shares did so because they felt that they had market knowledge. Finally, all of our detailed interviews revealed that a participant's primary decision for joining a market was based on the information they had. Thus we conclude that it is possible to open a Prediction Market to a large number of participants without adversely affecting the value of its forecasts. Before making this determination, however, General Mills should consult with its legal department to make sure that there are not any safe harbor issues (Hopman).

6.4 Prediction Market Manipulation

Mangold et al. found out in the Tech Buzz Game that Prediction Markets can be manipulated by participants. During our Prediction Market pilot we discovered that he was correct. As a result of the pilot markets manipulation (see sections 3.1.3, 3.2.7 and 3.2.8 for details) we took two steps to prevent manipulation:

1. Instituting per market investment caps is critical to preventing Prediction Market manipulation. During our detailed interviews, all of the participants mentioned that the investment caps prevented them from putting all of their money into the one or two markets where they were certain that they knew the forecast. Based on the Prediction Market pilot we learned that without per market caps participants would have been able to manipulate the Prediction Market.
2. Setting up normalized pricing prevented participants from shorting stocks that were outside the range of consideration and making easy money. Our analysis of trader behavior showed that Prediction Market participants were incented to buy and hold shares in stocks; in fact this is one of the crucial differences between winners and losers as illustrated in Figure 5.4.1, Figure 5.4.2, Figure 5.4.3 and Figure 5.4.4.

Our experience has shown us that participants will manipulate Prediction Markets. If General Mills chooses to add Prediction Markets to its forecasting process it will do well to continue to implement per market caps and normalized pricing to prevent manipulation.

6.5 Participant Behavior

Participant behavior is a key factor in the success of Prediction Markets. General Mills offered a wide variety of participants, from the active and informed participant to the gullible and curious participants. The analysis of transactional data followed by surveys and interviews allowed us to derive the following conclusions. Surveys showed that information about the event was the main

motivation to participate for 97% of the respondents. This implies that the choice of Prediction Markets should be such that participants from several departments are motivated to participate. 73% of the participants mentioned that they spent less than 20 minutes indicating that participants did not spend a lot of time in chasing down information. This is important for a larger rollout as it indicates that it does not take away too much time commitment from the individual participants. Non financial incentives, for example, recognition in a leader board, are powerful alternatives to financial incentives and can help sustain excitement and participation. Participants understand the value of a Prediction Market and are willing to participate if it is instituted as a process with defined goals.

6.6 Next Steps for General Mills

Overall, the Prediction Market study at General Mills generated significant interest from participants and showed that Prediction Market Forecasts perform as well as the current planning process. Based on our research we believe that there are three areas that General Mills should pursue using Prediction Markets.

First, General Mills should begin testing Prediction Markets for developing long-range forecasts; we would recommend setting a time horizon of six to twelve months. As part of our study we found that Prediction Markets performed better than the Operations Forecast for the Prediction Markets that closed in March. We believe that this trend will continue as the timeline is extended even further.

Second, Prediction Markets provide both a point forecast and distribution as they aggregate information. For products that have highly variable demand, General Mills could use the distribution from the Prediction Market to determine to move from a point based forecasting

process to a range based forecasting process. This would have the benefit of helping the organization plan for a range of outcomes rather than a single point forecast; the Prediction Market distribution would provide a systematic method for setting the forecast range that would automatically tighten or grow based on overall input.

Third, based on our research Prediction Markets are better at collecting opinions than gathering specific numerical information. Best Buy and Intel currently use their Prediction Markets to get “quick reads” on areas of uncertainty (Hopman) and visibility into employee sentiment (Jaedike) by quickly setting up a Prediction Market and asking participants to give input. Both felt that this process provided visibility that is not available through conventional means. General Mills could apply this same strategy for forecasting new product introductions and marketing campaigns. The results from our surveys and interviews suggest that participants would welcome the opportunity to give input into these areas; with the average participant spending less than twenty minutes per week to participate it would be very feasible to implement Prediction Markets in this way on a broad scale.

We would like to thank General Mills for the support that they provided on this project. The experience we had working with them was rewarding and filled with discovery. We would strongly encourage other researchers to reach out to General Mills as a research partner.

6.7 Additional Research

While Prediction Markets have a long history starting with the Iowa Election Markets, there is little research examining the use of Prediction Markets in a corporate setting. We would encourage others to continue to extend the work that we have begun. There are three questions that we were not able to address through our research efforts. First, we were not able to run Prediction Markets beyond ten weeks due to time limitations; based on our results we believe

this is an area where Prediction Markets have a clear advantage over conventional forecasting techniques. Additional research into this subject with other companies would be able to answer this question. Second, we used a normalized pricing strategy to incent participants to purchase shares in forecast ranges they believed to be accurate; based on our results our approach yielded the desired behavior we were looking for. Additional research into comparing average and normalized pricing strategies would answer the question of how pricing strategies affect Prediction Market performance. Third, we discovered an apparent contradiction in that, Sales and Customer Service won all of the Prediction Markets, but the Demand Planning group had the best overall performance; our results showed that participants with specific information were able to outperform participants with general information. Additional research comparing the performance of Prediction Markets comprised of experts compared to Prediction Markets comprised of generalists would shed light on the whether it is better to have more participants as Wolfers and Zitzewitz suggest or to have fewer participants as Hopman suggests.

Prediction Markets are a valuable tool for companies seeking to forecast future events. We hope that we have contributed to the overall understanding of the subject and look forward to seeing new research as it emerges.

Reference List

- Berg, J. (2008). Prediction Market accuracy in the long run. *International Journal of Forecasting*, (24), 28-300.
- Berger, R. (2003, April). 2003 GMA Logistics Study. Retrieved November 26, 2008, from <http://www.gmabrands.com/publications/docs/logistics2003.pdf>.
- Blank, E. (2009, January 2). MinnPost - Coleman-Franken recount: Some looming scenarios. News Blog, . Retrieved March 4, 2009, from http://www.minnpost.com/ericblack/2009/01/02/5549/coleman-franken_recount_some_looming_scenarios.
- Cain, M., & Drakos, N. (2008, August). Hype Cycle for Social Software 2008. Gartner.
- Chakravarty, S. (2005). ScienceDirect - Journal of Financial Economics : Stealth-trading: Which traders' trades move stock prices?*1. Retrieved December 21, 2008, from http://www.sciencedirect.com/science?_ob=ArticleURL&_udi=B6VBX-43DKSYF-5&_user=10&_rdoc=1&_fmt=&_orig=search&_sort=d&view=c&_acct=C000050221&_version=1&_urlVersion=0&_userid=10&md5=64727c26f52a1777567b8f0a1744de21.
- Demand Management Team. (2008, October). Interview With Demand Management Team - General Mills.
- Galton, F. (1884). Measurement of character. *Fortnightly Review*, 36, 179–185., (36), 179–185.
- Gruca, T., Berg, J., & Ci, M. (2005). Consensus and Differences of Opinion in Electronic Prediction Markets. *Electronic Markets*, 15(1), 13-22.
- Guo, Z., Fang, F., & Whinston, A. (2006). Supply chain information sharing in a macro Prediction Market. *Decision Support Systems*, 42(1), 194-1958.
- Hanson, R., Oprea, R., & Porter, D. (2006). Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior & Organization*, 60(4), 449-459.
- Hopman, J. (2008, November). Predictive Market - Lessons Learned At Intel.
- Hasbrouck, J. (1991, March). Journal Of Finance, Vol. XLVI, No. 1 - Measuring The Information Content Of Stock Trades. Retrieved December 21, 2008, from http://www.wiwi.uni-frankfurt.de/Professoren/empwifo/Lehre/Stat_2/Hasbrouck_Measuring.pdf.
- Hasbrouck, J., & Safianos, G. (n.d.). The Trades Of Market Makers: An Empirical Analysis Of NYSE Specialists JSTOR: The Journal of Finance, Vol. 48, No. 5 (Dec., 1993), pp. 1565-1593. Retrieved December 21, 2008, from <http://www.jstor.org/pss/2329060>.
- Ho, T., & Chen, K. (2007). New Product Blockbusters: The Magic and Science of Prediction Markets. *California Management Review*, 50(1), 144-158.
- IBM, R. S. G., & SAP, R. S. G. (2005, March). TRADE PROMOTION MANAGEMENT-REGAIN CONTROL OVER TRADE FUNDS SPENDING. Retrieved November 26, 2008, from file:///D:/Classes/Thesis/Papers/SPENDING_FACTS_GC10-4067-00.pdf.
- Jackson, K. (2008, October). Interview With Kathy Jackson - Demand Manager General Mills.
- Jaedike, B. (2009, March). Interview with Best Buy on their use of Prediction Markets.
- Justin, W., & Zitzewitz, E. (2006). FIVE OPEN QUESTIONS ABOUT PREDICTION MARKETS. *NBER*, JEL No. C9, D7, D8, G1, M2, 12060, 4, 5, 9.
- Kiviat, B. (2004, July 6). The End Of Management? *Time*, <http://www.time.com/time/insidebiz/printout/0,8816,1101040712-660965-1,00.html> (5 of 5)7/7/2004 8:11:57 PM.

- Leigh, A., & Wolfers, J. (2006, February). COMPETING APPROACHES TO FORECASTING ELECTIONS: ECONOMIC MODELS, OPINION POLLING AND PREDICTION MARKETS. NATIONAL BUREAU OF ECONOMIC RESEARCH. Retrieved from <http://www.nber.org/papers/w12053>.
- M C Lee, C. (1993). Earnings news and small traders - An intraday analysis. *Journal of Accounting and Economics* 15 (1992) 265-302. North-Holland. Retrieved December 21, 2008, from <http://deepblue.lib.umich.edu/bitstream/2027.42/30006/1/0000374.pdf>.
- Mangold, B., Dooley, M., Flake, G., Hoffman, H., & Dornfest, R. (2005). The Tech Buzz Game. *IEEE Xplore*, 38(7), 94-97.
- McCabe, J. (2004, April 1). *An Examination of the Predictive Abilities of Economic Derivative Markets*. Masters Thesis, New York University - Stern School Of Business.
- Perry, D., & Kittlitz, K. (2008, December). Prediction Market Configuration Session.
- Plott, C., & Chen, K. (2002). Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem. *SOCIAL SCIENCE WORKING PAPER* 1131, 2 to 31.
- Santiesteban, H. (n.d.). System and Method for Evaluating a ... - Google Patents. Retrieved December 21, 2008, from <http://www.google.com/patents?hl=en&lr=&vid=USPATAPP11535095&id=RsKqAAA AEBAJ&oi=fnd&dq=%22trader+metrics%22>.
- Schrieber, J. M. (2004, June). *The Application Of Prediction Markets In Business*. MIT. Retrieved from D:\Classes\Thesis\Papers.
- Schrieber, J. M. (2008, November). Predictive Markets - Reflections And Learnings.
- Surowiecki, J. (2004). *The Wisdom of Crowds*. N/A (1st ed., Vols. 1-1, Vol. 1). Doubleday.
- Thomas, D. (2009, March). Interview With Professor Doug Thomas, Associate Professor Supply Chain Management, Smeal College of Business (Pennsylvania State University).
- Thomas, G., Berg, J. E., & Cipriano, M. (2005). Consensus and Differences of Opinion in Electronic Prediction Markets, 15(1), 13-22.
- Wolfers, J., & Zitzewitz, E. (2004). Prediction Markets. *Journal of Economic Perspectives*, 18(2), 107-126.
- Wolfers, J., & Zitzewitz, E. (2006, March). PREDICTION MARKETS IN THEORY AND PRACTICE. NATIONAL BUREAU OF ECONOMIC RESEARCH.