

Experience from Hosting a Corporate Prediction Market: Benefits beyond the Forecasts

Thomas A. Montgomery
Ford Motor Company
MD 2122 RIC, Box 2053
Dearborn, MI 48121
01-313-337-1817
tmontgo1@ford.com

Paul M. Stieg
Ford Motor Company
MD 2122 RIC, Box 2053
Dearborn, MI 48121
01-313-323-2098
pstieg@ford.com

Michael J. Cavaretta
Ford Motor Company
MD 2122 RIC, Box 2053
Dearborn, MI 48121
01-313-594-4733
mcavaret@ford.com

Paul E. Moraal
Ford Motor Company
Suesterfeldstr 200
Aachen, Germany
49-241-942-1218
pmoraal@ford.com

ABSTRACT

Prediction markets are virtual stock markets used to gain insight and forecast events by leveraging the wisdom of crowds. Popularly applied in the public to cultural questions (election results, box-office returns), they have recently been applied by corporations to leverage employee knowledge and forecast answers to business questions (sales volumes, products and features, release timing). Determining whether to run a prediction market requires practical experience that is rarely described.

Over the last few years, Ford Motor Company obtained practical experience by deploying one of the largest corporate prediction markets known. Business partners in the US, Europe, and South America provided questions on new vehicle features, sales volumes, take rates, pricing, and macroeconomic trends.

We describe our experience, including both the strong and weak correlations found between predictions and real world results. Evaluating this methodology goes beyond prediction accuracy, however, since there are many side benefits. In addition to the predictions, we discuss the value of comments, stock price changes over time, the ability to overcome bureaucratic limits, and flexibly filling holes in corporate knowledge, enabling better decision making. We conclude with advice on running prediction markets, including writing good questions, market duration, motivating traders and protecting confidential information.

Categories and Subject Descriptors

J.1 [Computer Applications]: Administrative Data Processing—*business*; Social and Behavioral Sciences — *economics*

General Terms

Algorithms; Management; Measurement; Performance.

Keywords

Artificial markets; prediction markets; forecasting; organizational knowledge; social media.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
KDD'13, August 11–14, 2013, Chicago, Illinois, USA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-2174-7/13/08...\$15.00.

1. INTRODUCTION

Prediction markets leverage the wisdom of crowds [21], the knowledge that is dispersed among the members of a group of people [8], through a virtual stock market mechanism. Participants buy and sell answers to questions such that the stock price (the current price of an answer to a question) is a prediction of the likelihood of that answer being the right answer. Compared to surveys, traders in a prediction market invest in answers to questions according to what they *think will happen*, instead of answering questions according to what they *want to happen*. For example, a trader who believes that Candidate A will win an election will invest in A's stock even if they would prefer to have Candidate B win [3] [10].

The result is a mechanism that is relatively insensitive to the demographics of the participants, is fun, and can be deployed anywhere at any time. Prediction markets can be used to fill holes in corporate knowledge on

- External factors such as macroeconomic trends or competitors' actions
- Corporate performance metrics such as sales volumes or market share
- Internal performance such as the timing of internal gates or which advanced projects to pursue.

In Section 2 we present highlights of the history and results of running the Ford Prediction Market (FPM) on a wide variety of business questions, comparing the forecasts made by “the crowd” to what ultimately happened in the real world. This demonstrates how a prediction market can be more accurate than other forecasting mechanisms, but also shows where Ford employees tend to see things differently. It even includes some surprises, like an unexpected bias against short selling that we found.

In Section 3 we describe the benefits of running a prediction market that go well beyond just the forecasts made. This includes tracking changes in stock prices over time, leveraging comments made, overcoming bureaucratic limits, educating employees, and deploying it whenever and wherever information is needed.

In Section 4 we present advice on running a prediction market including writing good questions (even those with unverifiable answers), market duration, motivating participation, protecting confidential information, addressing market manipulation, monitoring liquidity, and more.

The practical experience described here should prove valuable for any corporation considering running their own prediction market.

2. FORD PREDICTION MARKET DEVELOPMENT AND RESULTS

The Ford Prediction Market has gone from in-house software used experimentally for sales forecasting by a few traders in Research and Advanced Engineering (R&A), to commercial software used by thousands of employees globally (U.S., Europe and South America) to predict favorability of new vehicle features, price points, macroeconomic trends, sales and more.

2.1 The Ford Prediction Exchange

The Ford Prediction Exchange (FPEX) was the first prediction market at Ford, developed in 2006. Instead of buying and selling stock, it used a scored polling mechanism in which traders made forecasts by specifying ranges of values (a low and high value for each answer). The more confident traders were, the smaller the range they would use. Monthly sales volumes and feature take rates were forecast, and the error rates were lower than official forecasts in most cases. The internally developed tool provided a consensus density graph which characterized the voting strength across the range of possible answers. The success of this work gave us confidence to continue investigations.

2.2 The Ford Prediction Market – Test

In the next phase we moved away from the FPEX, partly because we wanted to be able to predict events that were not numbers, such as which new vehicle features held the most promise. We faced a make versus buy decision and considered extensions to FPEX, open source solutions, and commercial tools. We ultimately chose Inkling's [17] hosted commercial solution for its ease of use, low cost, and technical support. Inkling also allows traders to post comments, which proved to be an unexpected benefit we describe later in Section 3.2.

Our primary business partner in 2009 was Feature Planning which spans Marketing and Product Development (PD). We designed the market around new vehicle features: which features would perform best in market research, how much would consumers say they would pay for a feature, which implementation would be best received, and what other vehicle features Ford should consider.

We invited about 1,000 people to participate in the test including everyone in R&A, and select employees in a few other divisions (IT, Marketing, and HR). 350 people registered to participate, and 250 ultimately traded in the market. The market ran for three weeks and generated about 7,000 trades and 350 comments.

Examining the results revealed some of the biases of the Ford traders in the prediction market when compared to traditional market research. The employee-based prediction market valued certain advanced features higher, and some other features dramatically lower. The ones that ranked much higher had similarities to recent notable Ford successes, while those that ranked much lower had been criticized by the traders for a poor cost/benefit ratio from an engineering perspective.

Important lessons learned while running the test market included how much effort is required to

- Obtain clearance to run the market
- Refine questions
- Monitor the market in real-time
- Analyze and report on the results

As we rolled out the prediction market to other regions of the company, we stressed the importance of socializing the concept early to avoid delays due to questions from key stakeholders such as senior management and HR. A key insight here is that the primary cost of running a prediction market is not the software, but the time spent by employees trading in the market.

We also allowed traders to suggest new vehicle features to add to this test market, and learned that contrary to some proponents, it is not easy to just “let the market run.” There was quite a bit of effort needed to vet the proposed ideas, decide which to add to the market, develop descriptions and images for them, and post them. Doing this real-time as the market runs requires a dedicated team.

2.3 The Ford Prediction Market – Pilot

In 2010, we expanded the Ford Prediction Market to all PD and Marketing employees in the U.S. A total of 10,000 employees were invited in an email signed by the Vice Presidents of PD and Marketing. 1,800 employees registered based on that single email, and 900 traded in a 2 week market in May, generating over 13,200 trades and 2,700 comments. We reopened the market in August and ran for another 3 months, pushing the total number of traders over 1,000.

We developed business partners from six different divisions of the company, and had questions on

- Rating and pricing potential new vehicle features
- Macroeconomic trends such as future average gas price
- Global feature strategies
- Adoption rates of advanced technologies
- Employee incentives

2.3.1 Europe

Also in August 2010, we transferred the prediction market technology to Europe where our European research arm, Ford Forschungszentrum Aachen, led a parallel market. They invited 5,000 PD and Marketing employees, 600 of whom registered, and about 400 ultimately engaged in trading. 60% of traders were based in Germany, and most of the rest were in the UK. They asked comparable questions as in the U.S., with appropriate adjustments made such as kilometers for miles, slightly different vehicle categories, and more diesel engines.

2.3.2 Comparison of Employee Prediction Markets with Traditional Market Research

The vehicle features tested using the FPM in the U.S. and Europe were the same features investigated by traditional market research in the U.S. and elsewhere. Coordination between traditional methods and prediction market was deliberate. Traders in the U.S. market were asked to forecast the U.S. research results.

A careful examination of the results reveals some insights. Similar to our earlier results, Ford employees rated certain high-tech features higher than predicted by the research. Some features were rated lower based on the engineering judgment of Ford employees and differing regional needs.

Since the results from traditional methods and prediction markets have so far proven to be different, Ford cannot just replace market research with prediction markets. However, each mechanism has its own strengths and weaknesses. Clearly neither method is right, because only the marketplace can measure the true value of a product or feature and that value is greatly influenced by its implementation and marketing. For these reasons, we see

prediction market results as complementary to traditional methods for evaluating products and their features.

Both methods provide valuable insight. Employees may bring corporate biases to their deliberation, but they also bring industry knowledge and experience. At least one successful corporation, Apple, has embraced the ideal that the company must determine what the customer will want. Or as our own Henry Ford is quoted (apocryphally) as saying, "If I'd asked my customers what they wanted, they'd have said a faster horse."

2.3.3 Comparison of Employee Prediction Markets and Sales Forecasts

During the US market, we experimented with a repeated, ongoing question. Over three-months, we asked traders to forecast the US sales volumes for selected Ford and Lincoln models. Traders predicted the future weekly and monthly volumes for each. The questions closed in advance of the actual projected sales period. The objective was to determine the accuracy of prediction markets, and whether they would improve over time as traders gained more experience as found by Google [4]. We also wanted to see if traders would continue to participate over a longer period.

Overall the Ford Prediction Market did very well. This does not mean that we can stop producing an official forecast and just rely on the prediction market. The heaviest traders in these questions turned out to be individuals involved in the official forecast. They brought specialized knowledge to the market. If that knowledge is removed, we might only have poor performers in the market (whose primary role is to provide liquidity by losing their money). This kind of feedback loop in which a forecasting effort informs prediction market traders has also been described between the FiveThirtyEight forecasts and the Intrade prediction market [19].

2.3.4 South America

In 2011, the FPM moved to Ford South America (FSA) to provide local insight on vehicle features for the feature planning process. FSA employees participated at unprecedented rates, both in the number of people who traded, and the amount of activity by each trader. Of the 3,000 employees invited in Brazil, 1,300 registered to participate (43% of those invited registered, compared to 18% in a similar U.S. prediction market). Of those who registered, 65% of the Brazilians made a trade or posted a comment, compared to 50% in the U.S. The average Brazilian trader also posted nearly twice as many comments (average of 5.7 in Brazil versus 3.0 in the U.S.), and made nearly three times as many trades (41.4 versus 14.7).

The biggest challenge faced in extending to South America was the need to accommodate the languages of Portuguese in Brazil and Spanish in Argentina. Since this was the first time InKling ran in a language other than English, Ford South America became the beta-testers of this new feature and helped with the translations.

As of this writing, Ford has had a total of 2,300 traders in their prediction markets around the world, in the U.S., Germany, U.K., Belgium, Brazil and Argentina. This makes the Ford Prediction Market one of the largest known corporate prediction markets.

2.4 Bias against Short Selling

We conclude this section with a surprising observation in which traders exhibited a bias against short selling. We expected traders to be just as likely to buy or sell since the software interface looks the same regardless of whether a trader wants to drive the price up

(and they will be buying) or down (and they will be selling short). This was not the behavior we observed.

Figure 1 shows the closing stock prices of vehicle features tested in one FPM phase. Traders knew that at most 25 features would cash out at \$100, the rest at \$0. As can be seen, there are more stocks close to \$100 than close to \$0. In fact, 28 stocks closed above \$75, while only two closed below \$25. The opportunity to make money by selling short a "certain loser" was much greater than that afforded to buying a "certain winner." The lowest rated stock would have earned a trader over \$13/share if sold short, whereas buying the highest rated would earn less than \$4/share.




Figure 1: Closing Stock Prices of Vehicle Features Illustrating Short Sell Bias

There are a few explanations. There could be a positive bias since the features tested were ones that had been through a vetting process, so few ideas were poor. The InKling software also has an inherent bias against short selling for multiple choice questions, which we illustrate with an example from Ford of Europe.

The share price for two answers shot up just before closing in a question on the sales rate of future technology. Contrary to suspicion, this was not malicious manipulation by a single trader; several traders appeared to share the overly optimistic view on the popularity of one technology in 2020. Other traders reacted by shorting the stock, but due to the particular market mechanisms, they could not significantly reverse the upward trend.

Since the market had obviously not come to a consensus, we extended the trading period for this question. Over this extended period, the share price dropped again, though not back to the average level it had been trading at; it closed at 5.4%. In order to ascertain the extent the share price might be influenced by the mechanism of share price calculations and reserves required when short-selling stock, we introduced a new question specifically on just one of the technology options. This stock closed at 0.5%, only a fraction of the closing value for the same question as described above, and arguably a more realistic value. With this, we showed that the particular mechanism of short-selling can substantially influence the results, especially for low-probability events.

The difference is largely due to the way that short-selling works in the software we used: when a trader sells shares that he doesn't own, he is promising to buy back that number of shares later – speculating the price will drop. To be sure the trader will have enough money left to buy these shares later, even if the price has gone up, money is locked up. In fact, in this type of multiple choice question, the amount of money held is the difference of the current stock price and the maximum possible price (100%). Therefore, the lower the stock price, the more "expensive" it is to sell short. This makes sense when exactly one of the answers can be correct – any could end up at 100%. However, when looking at things such as market share, it is less appropriate. For questions

predicting a single number value, the money kept in reserve to cover shorts is twice the current price. As a result, traders don't tie up a lot of money shorting stocks in the 0-5% range.

Some implementations avoid the issue altogether. General Electric did not offer short selling in its own corporate prediction markets (they used bid-ask style markets), due to concern over traders' difficulty in understanding how to apply the concept [12].

3. BENEFITS BEYOND FORECASTS

Running a prediction market offers more benefits than just forecasts made with the wisdom of crowds. How the crowd arrived at its consensus is also valuable, including the changing stock prices and the comments made by traders attempting to convince each other of their positions. Furthermore, prediction markets overcome bureaucratic limits, can be flexibly applied to fill holes in corporate knowledge, and offer side benefits such as educating employees. We explore each of these benefits in turn.

3.1 Stock Prices and Trends

The results presented so far have been based on the closing stock prices for each question, but that is not the whole picture. Figure 2 shows a situation where one answer (one stock) took an early lead and held it for the duration of the market. There is increased confidence in the result of the market when an early lead holds up.




Figure 2: Stock Prices with a Clear Winner

In some cases there is no clear winner, reducing confidence in the final stock price. Figure 3 shows a situation where the market did not come to a consensus. Instead, two answers traded the lead position four times over the market, and likely would have continued to trade positions had the market run longer.



Figure 3: Stock Prices without a Clear Winner

Sudden changes in stock prices (Figure 4) suggest a shift in sentiment. By watching for shifts instead of just waiting for the final result, sponsors of questions can determine whether they need to react. It is important to understand the reason for a sudden change in the forecast for a product timing or sales volume question. Were there internal changes such as new requirements or personnel changes? External changes such as competitive actions or economic disruptions? Given the change, the business may need to respond quickly. As we discuss next, the first place to look for reasons is the comments posted by traders.



Figure 4: What Caused this Stock Price to Drop?

3.2 Comments

When reviewing the results from the Ford Prediction Market with our business partners, every one of them emphasized the value of the comments generated. The market had ignited an online conversation – a debate – which we could watch and capture. As we have seen, closing stock prices can be good point estimates, and stock trends can illustrate changing sentiment, but comments get at why people believe one answer is better than another.

For example, in forecasting the average price of gasoline 5 years in the future, the forecast price changed in predictable ways with changes in the U.S. and global economy. However, the strategy team who sponsored the question recognized the value of the comments as a source of environmental factors to consider in their analyses. This sample shows typical depth of thought seen in comments from our traders:

I think the current situation in the Gulf will have a ripple effect in fuel pricing this year. We may very well approach \$4.00. With more fuel efficient powertrains I would anticipate a drop in demand in the US, but emerging markets such as China will more than offset...

Ford employees participating in the prediction market saw it as an outlet for expressing their opinions to management and demonstrated a strong desire to communicate. During our May 2010 market, Ford employees posted more comments than any of Inklings's other clients ever had over a similar period. Inklings list of clients includes a broad range of organizations, including Chevron, CSX, Dow Corning, Johnson & Johnson, DOE, Cisco, Harvard Business School, The World Bank, Capital One, and Wells Fargo. According to Inklings, none had exhibited the same passion by commenting as frequently as Ford employees.

Surprisingly, the length of comments also increased over time (in contrast to other trader activity metrics which dropped over time as shown later in Section 4.3). As the prediction market ran, traders posted longer and more detailed comments. Figure 5 shows how the length of comments (measured by the average number of characters per comment) grew by about 50% from the beginning to the end of the market.




Figure 5: Average Comment Length (in sets of 50 consecutive comments)

In some cases, the direct influence of comments on stock prices can be seen. Figure 6 shows the prediction market trend on monthly sales volumes of two vehicles. The highlighted line is for a vehicle that was new to the market and had received very positive reviews. Employees were naturally bullish about its sales prospects. However, some traders commented that the sales

forecast had exceeded the production volume (it was forecasting sales of more vehicles than had been produced), the market corrected itself, and the forecast lost its irrational exuberance.




Figure 6: Stock Price Dropped due to Comments Posted

The inclusion of comments comes at a price: the requirement for administrative oversight. In a corporation, there are some policy and legal implications. Although we labeled our site confidential, we did not want traders to post sensitive information. We were also concerned that "flame wars" not erupt among traders, especially when discussing macroeconomic and political issues. Consequently, we required traders to acknowledge when registering that they agreed to abide by the Ford corporate communications policies. We also monitored all comments made during the markets and used administrative abilities to edit or remove problematic comments. When such action was needed, we contacted the traders so they knew that we were modifying or removing the comment and the reasons. While we had a number of instances of editing or removal, none were serious, and traders were always understanding and cooperative.

3.3 Overcoming Bureaucratic Limits

In corporations, it is common for employees to tell their management what they think they want to hear. In contrast, prediction markets are egalitarian and encourage candor. All traders begin as equals, each starting with the same money to invest. Traders can express their opinions openly since the interface offers anonymity (participants trade under self-chosen user names). Because they are rewarded for making accurate predictions, they invest according to their beliefs. In fact, traders who invest wisely will have more money to invest in future questions, and therefore will have greater influence on the market over time. This results in a positive feedback loop based on merit, independent of organizational structure or corporate politics.

Proof that traders are not swayed by corporate rank was clearly demonstrated in one of the markets we ran. In it, a vehicle feature was promoted by a Ford executive who posted multiple favorable comments about the feature. In spite of the executive's rank (who had chosen not to be anonymous in the prediction market), the market as a whole was not convinced by his comments, and the feature closed with a relatively low stock price.

When prediction markets span organizational boundaries, the dialogue created can uncover synergies between the organizations. For example, a prediction market that spanned government contractors revealed an opportunity to solve a larger problem once it was understood that different aspects of it were already being engaged by each contractor.¹

3.4 Flexible Deployment (Time and Location)

Prediction markets can be deployed anywhere in the world, any time of the year. All that is needed is an internet connection,

¹ Researchers at MITRE described this during the conference, Using Prediction Markets in Government, held 9/22/10 in Mclean, VA. <http://governmentwisdom.eventbrite.com/>. Of course, this benefit is not unique to prediction markets as it can be found in other collaborative environments.

traders, and questions to answer. Ford has made use of this capability to fill holes in data generated by market research. We have also timed prediction market runs to match Ford's business cycles, providing data when it is needed by the corporation.

For example, an idea was proposed to combine two features normally not offered together. An engineering team debated the relative merits of the idea, knowing that they had to decide before planned market research could test the idea. To support their decision, they added the idea to the Ford Prediction Market. The idea performed extremely poorly, confirming their original opinion and giving them increased confidence to end the debate.

3.5 Education of Employees

Another unexpected benefit of running a prediction market is education. In a post-market survey, 93% of employees said they learned something by participating (Figure 7).

- The questions revealed strategies Ford is considering.
- Traders exchange knowledge as they attempted to convince each other of their opinions.
- Traders are inspired to research answers.

Amount learned about Ford business through participation




Figure 7: Learning through Participation

4. ADVICE ON RUNNING A PREDICTION MARKET

Many important lessons are learned by doing. For example, if questions are not formed well, traders can get confused, and results can be ambiguous. Therefore, we present practical advice on a variety of topics including writing questions, running idea markets (questions without verifiable answers), market duration (episodic versus continuous markets), maintaining confidentiality, and motivating participants.

4.1 Writing Good Questions

Ideal prediction market questions are:

Useful: By covering topics that have clear business value, we avoid the critique that employees are "playing" on company time and reduce the overhead of reviewing trader-submitted questions.

Forward looking: Prediction markets are not surveys; traders are investing in answers that they believe will come true regardless of which answer they would prefer to see come true. So instead of asking which they would prefer (which feature they like best), ask which outcome will happen (which feature will sell best).

Knowable: Questions should be on topics within the realm of traders' personal knowledge. It is not necessary for all traders to be knowledgeable on all questions, but each question should have enough traders with access to relevant information to avoid "thin" markets in which only a few traders participate.

Unambiguous: Unambiguous questions avoid unnecessary debate and uncertain results. For example, in 2010 a public market asked the question, “Will Ford outsell GM this year?” Many details about how the question will be evaluated are needed: full year or a given month? fleet and retail? where? according to whom?

Complete: Questions should have a complete set of non-overlapping answers. Often one of the answers must be “other” or “none of the above” for completeness.

Verifiable: Ideally, at some point not too far in the future, the real world answer should be known. This allows stocks to be cashed out, rewarding those who made good investments and freeing funds for investing in new questions. The next section, however, shows that this rule can and must be violated at times.

4.2 Questions without Verifiable Answers

Traditionally, prediction markets questions are structured such that they predict a measurable outcome at a specific time, such as an election result. But there are many questions of business value that look too far into the future or are otherwise unverifiable.

With long-term questions such as “What will be the average price of gasoline in 2020?” it isn’t feasible to let the market run until the event occurs. But corporate planning usually involves a time horizon greater than traditional prediction markets; to be useful, corporate prediction markets must allow long-term questions.

An example of a non-verifiable question is “Which features should we develop?” Corporations do not have the resources to produce every possible product variant, so there is no way to determine which would have sold best. We found inspiration in General Electric’s “Imagination Markets” which is a variant of prediction markets, called an idea market, to solicit and rank ideas from employees for new technologies. Traders could submit their own ideas, and they could buy and sell shares in other ideas. GE awarded \$50,000 in research funding to the top rated idea [12].

To evaluate potential new vehicle features, we phrased questions to sound as if they were verifiable. Rather than asking “Which features should we develop?” we pretended market research was pending asking “Which features will do best in market research?”

But how to value a market without a verifiable answer? GE’s used the volume-weighted average price of the answer over the last 5 trading days of a question. We used the closing value of an answer at market end as the final value of the answer. While our solution worked well overall, it unfortunately creates a potential for market manipulation, which we discuss in Section 4.9.

4.3 How Long Should a Market Run?

Prediction markets can be continuous (with a regular cadence of new questions), very short term, or something in between. A continuously running market is necessary when forecasting regular business metrics such as sales, production, or warranty.

Running a market for weeks or months has the advantages of being long enough for trader interactions to mature, but short enough to maintain enthusiasm. Timing can be tied to a business cycle or a particular event, such as annual planning. This is the approach we have taken, as did GE [20].

In the Infosur Concept Exchange (iCE) prediction markets only run for hours or minute. Traders make trades only once, but can hedge their bets by buying more than one concept [11]. This approach has many of the advantages of longer prediction markets, but lacks the interplay between traders.

Maintaining portfolio values of traders over long term or repeated markets provides a feedback loop such that the best traders have more influence on the market. Those who have particular skill in prediction can be identified, and their expertise can be tapped by the corporation in other ways. Tagging questions by type can help identify those who are proficient, say, in financial forecasts.

Loss of interest over time is an issue affecting both overall markets and individual questions. For example, we conducted an extended market with weekly and monthly questions forecasting vehicle sales. As Figure 8 shows, the number of unique traders dropped significantly over four months.




Figure 8: Number of Unique Traders per Month

In order to counter the general loss of interest, various strategies may be used. We found that an effective strategy was to continue to introduce new questions over the life of the market, rather than asking all the questions at the beginning of the market. We also obtained bumps in activity by sending reminder emails to participants. We discuss motivating traders further in Section 4.6.

4.4 Monitoring Liquidity

One potential problem in ongoing markets is that traders may have invested most of their assets, such that they have no more money to invest in new questions. Reserves held to cover potential losses in selling short may tie up more assets than traders realize. Cash can also be locked up in questions that are closed but not yet cashed out. Consequently, the markets can lock up, compromising their accuracy as research has shown that there is a correlation between market liquidity and prediction accuracy [22].

Though traders can sell their positions in active questions to invest in new questions, this option was not always obvious to traders and should be promoted. If liquidity does get tight, additional money must be injected into all of the accounts to avoid lock-up.

4.5 Protecting Confidential and Sensitive Information

Caution is needed to avoid revealing proprietary information in a prediction market. Not only can photographs and technical descriptions provide intelligence to competitors, but simply asking a question can reveal sensitive strategies under consideration.

Limiting participants limits exposure. If the market is to discuss sensitive matters, one may restrict the market to employees and ask them to forecast public opinions. However, thousands of employee traders also create significant potential for leaks. Sometimes it is necessary to limit participants to members of a division, department, work team, or list of individuals.

We also limited exposure by asking more general questions. In one case, Ford was considering a new feature for a particular vehicle. However, that consideration in itself was too sensitive for dissemination through the employee base. We disguised the

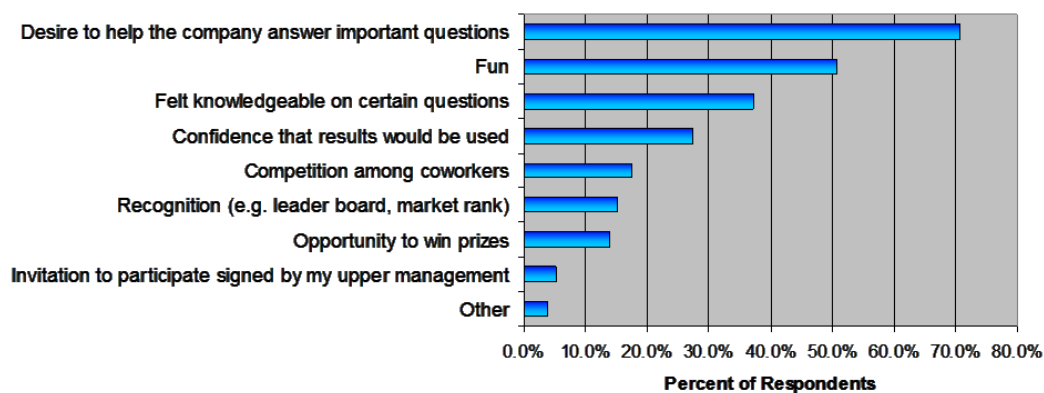


Figure 9: Reasons for Participating in the Ford Prediction Market

question by asking how well the feature would sell across all vehicle lines, with each vehicle being a separate answer. When the market was complete, we had our answer for the individual vehicle, along with comments for it. In addition, we had market data and comments for other vehicle lines to compare against.

Technology may also help in protecting intellectual property. One may disable the save feature on sensitive images, or watermark them so that if the image is leaked, the source can be identified.

There are also legal aspects to corporate prediction markets. U.S. SEC regulations prohibit insider trading. A corporate insider, using the results of an internal prediction market to time trading of the company’s securities, could be construed as violating the regulations [1]. Consequently, we deliberately avoided questions regarding Ford’s financial performance (e.g., predicting future Ford stock price).

4.6 Motivating Participation

Prediction markets with too few participants are known as “thin markets” and suffer from low trading rates, early lock-up of answers, undue influence from the few who do participate, and ultimately poor forecasts. Therefore it is important to motivate traders to participate.

In a corporate setting, this starts with getting HR and executive permission to run a prediction market – buy-in to allow employees to spend time trading in the market. Sending an invitation signed by senior management, asking employees to participate, tells them

- It is okay to spend time in the prediction market.
- Results will be used. Management is listening.

When we asked employees why they participated, the number one reason given by over 70% was the desire to help the company answer important questions (Figure 9). The business value of the questions we asked was obvious. For example, questions on vehicle features and future gasoline prices clearly influence product direction and the long-term viability of the company.

Over 50% of our participants said it was fun. Ford Motor Company frequently runs surveys of employees, dealers and customers for various reasons; fun would be a rare explanation any of these constituencies would give as to why they answered a survey. The opportunity to win prizes in the prediction market was not listed as a strong motivator by very many participants, whereas pecuniary rewards are common in surveys.

Employees also participated because they felt knowledgeable on certain questions, were confident the results would be used, and enjoyed the competition among their co-workers (bragging rights)

and the recognition of being ranked high on the leaderboard (personal pride).

4.6.1 Prizes

If prizes are to be given out, the prize structure must be carefully constructed to reward desired behavior. Prizes that are too large can motivate people to game the system, and have been demonstrated to reduce market accuracy by inciting risky behavior [5]. There are also potential legal pitfalls. U.S. regulations on gambling may come into play with large prizes [1]. We used prizes of nominal value, \$100 performance awards and weekend prototype vehicle drives. Similarly, Google found that employees really liked the T-shirts that they could win [4].

If prizes are only awarded to a few top performers, they will fail as a motivator for anyone who is not close to the top. We used a lottery system to award prizes, structuring it so that anyone could win. The odds of winning each participant had was based on

- Performance – the dollar value of their portfolio.
- Participation – one thousand additional “dollars” for every day that they participated by trading or commenting.

Surprisingly, even with this prize structure and relatively short prediction markets of a few weeks, no single trader participated every day in an attempt to increase their odds of winning.

Many of the early prediction markets used real money because economic theory suggested that traders had to “put their money where their mouth is.” However studies have shown similar results regardless of money or prizes [18]. In our U.S. and European prediction markets, we had similar participation rates regardless of whether prizes were offered (about 9% of invited employees participated in the U.S. where prizes were offered, and about 8% of invited employees participated in Europe where prizes were not offered).

We spent a lot of effort getting permission to award prizes, and many people believe that prizes are important. However, this experience shows that when traders are motivated by other factors (helping the company while having fun), prizes are not needed.

4.7 Trading Philosophy

A survey question we asked on trading philosophy supports our belief that Ford traders were overwhelmingly motivated to find the right answers to questions. Over 80% claimed that they traded on fundamentals: how they thought the real world would behave (Figure 10). In spite of the fact that some questions could not be judged according real world outcomes, traders still tried to make good forecasts. A much smaller fraction, 8%, invested according

Trading Philosophy

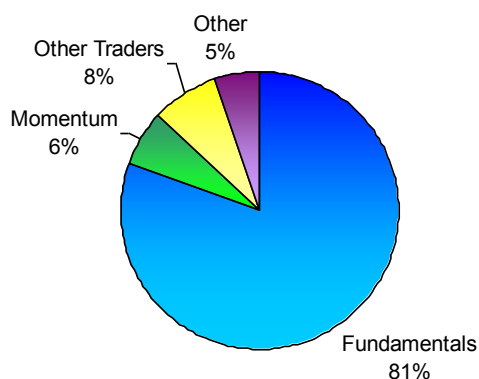


Figure 10: Most Traders Traded on Fundamentals

to how they thought other Ford employees would invest (they tried to predict Ford's internal biases), and 6% used a momentum strategy in which they watched for market trends and tried to buy stocks on the way up, or sell on the way down.

4.8 Barriers to Participation

When asked why they didn't participate more, the number one reason given was "too busy." While little can be done to provide people with more time, many other barriers can be addressed.

Many felt that if they did not get into the market early when trading opened, their opportunity to make money was limited. While we observed traders who did not invest early but were still very successful, this barrier to participation can be addressed by:

- Giving advance notice of when the market will open
- Introducing new questions while the market is running
- Setting initial stock prices at intelligent levels so that the first traders do not have obvious easy money plays

Some felt that investing was too complicated. This can be limited by providing clear instructions on how to trade, insights on trading strategies, and offering a simple software interface.

Finally, offering a variety of types of questions, each addressing different business needs, is likely to result in most people finding some questions of interest. Some people may be interested in new features, others in sales volumes, and others in economic trends.

4.9 Addressing Market Manipulation

One final concern is the possibility of market manipulation: whether traders can skew the market in their favor and reduce the validity of the prediction. Prediction market traders see a benefit from winning, which is one of the advantages and incentives offered over a traditional survey. But the question is whether the trader has the opportunity, and obtains a benefit, in cheating.

One way to reduce the temptation to cheat is to reduce the benefit. Prizes that are too large may inspire cheating and take the focus off the corporate benefits where it should be.

Another method is to watch for unusual activity, such as large purchases and sales of the same answer within a short period of time, or coordinated trades, wherein two traders consistently switch buying and selling in an answer. Two traders can collude through a scenario in which one trader buys a stock; the second drives up the price; then the first takes a profit by selling. This results in the transfer of funds from one trader to the other. While beneficial if prizes are awarded to the trader with the most money, since our prizes are lottery based, pooling money has no net effect

on the odds of winning. We did observe experimentation by traders in the FPM, but not collusion. When we did see unusual behavior, we contacted the trader, and were always satisfied with the response.

Research shows that the ability to manipulate prediction market prices is limited, especially when there is a sufficient number of traders in the market [14] [23]. Attempts by one trader to push a stock price are offset by other traders readjusting the price (however imperfectly). Thus, researchers conclude that the manipulation is only temporary.

The remaining risk is when there isn't enough time for the market to correct before trading closes. In one case, a trader exploited a shortcoming in our approach to cashing out questions without verifiable answers at final market value. This cash out approach makes it lucrative to buy large quantities of cheap stock just before closing. Doing so can generate a profit in excess of 1000%. (In normal trading, share prices drop with each share sold, hence buying many shares and subsequently selling them is financially neutral. However, cashing out at final price makes it better to be fully invested, and best to own a lot of cheap stock.)

One trader bought significant shares in a stock trading at around 5% in the final days of the market. His purchases pushed the price up to 25% at closing. This is good for the individual, but hurts the prediction. This mechanism can be limited by keeping the closing time of a question a surprise. GE tried to avoid this manipulation by not announcing the close date and time. If the traders didn't know when the market closed, they couldn't attempt a last-minute change. However, this was not an option in our software. GE's approach of using a last-5-days-average for cashing out also lessens the attractiveness of such trading strategies.

Comments also serve as an impetus for correction. When the market suspected that a given trader was trying to push an answer for his or her benefit, other traders would "call out" the trader as over-hyping an answer. The market then tended to sell the stock and correct the artificial inflation. This embodies the philosophy of the wisdom of crowds: sharing information improves results.

5. CONCLUSION

Ford Motor Company developed considerable experience with prediction markets, using both in-house and commercial software, and deploying them around the world to thousands of employees. This practical experience provides significant insights into the methodology. Forecasts of objective items such as sales volumes proved accurate, while forecasts of subjective results showed significant variances with predictions from market research.

This doesn't mean that forecasting objective items is a good idea, and subjective a bad idea. Regardless of whether questions were objective or subjective, it was clear that we could not simply replace an existing corporate function with a prediction market. In sales forecasting, our official forecasting team traded heavily and therefore played a significant role in the accuracy of the prediction market. The business question, then, becomes whether the potential increase in accuracy offered by the prediction market is worth the time and expense of running the market. In forecasting vehicle features, the results were different than traditional market research, but could be seen as complementary since the results tended to be tempered by engineering judgment.

Overall, our experience revealed that the value of prediction markets goes well beyond the closing stock prices. The price trends and comments provide additional value, and there is

significant benefit in a prediction market's ability to overcome bureaucratic limits and be deployed anytime, anywhere to fill holes in corporate knowledge. Prediction markets also engage and educate employees.

6. ACKNOWLEDGEMENTS

A project this large could not have been possible without the help of many. Thanks go to our executive sponsors; colleagues in Europe and South America who ran markets in their regions; Human Resources who navigated approvals; those who helped with registration, surveys, and monitoring; business partners from eight organizations who developed questions, worked with us to analyze results, and incorporated them in their decision processes; and Christina LaComb of GE who kindly shared her company's experience and approach. Finally, the Ford Prediction Market could not have succeeded without the employee traders who enthusiastically embraced this new technology, making thousands of trades and comments in their desire to help Ford Motor Company answer important questions. Thanks to all involved!

7. REFERENCES

- [1] Tom W. Bell, "Private prediction markets and the law," *The Journal of Prediction Markets*, vol. 3, no. 1, pp. 89-110, 2009.
- [2] Joyce E Berg, Robert Forsythe, Forrest Nelson, and Thomas Rietz, "Results from a dozen years of election futures markets research," in *Handbook of experimental economic results*, Charles Plott and Vernon Smith, Eds. Amsterdam: Elsevier, 2008, pp. 742-751.
- [3] Joyce Berg, Forrest Nelson, and Thomas Rietz, "Prediction market accuracy in the long run," *International Journal of Forecasting*, vol. 24, no. 2, pp. 285-300, 2008.
- [4] Bo Cowgill, Justin Wolfers, and Eric Zitzewitz. (2009, January) Using Prediction Markets to Track Information Flows: Evidence from Google. [Online]. <http://www.bocowgill.com/GooglePredictionMarketPaper.pdf>
- [5] James Duncan and R. Mark Isaac, "Asset markets: How they are affected by tournament incentives for individuals," *American Economic Review*, pp. 995-1004, 2000.
- [6] Foresight Exchange Partnership. (2012) The Foresight Exchange Prediction Market. [Online]. www.ideosphere.com
- [7] Alan Hall. (2011) Ford Prediction Market. [Online]. http://media.ford.com/images/10031/FTL_predmarkt.pdf
- [8] Friedrich Hayek, "The use of knowledge in society," *American Economic Review*, vol. XXXV, no. 4, pp. 519-530, September 1945.
- [9] Hollywood Stock Exchange, The Entertainment Market. [Online]. www.hsx.com
- [10] IEM - Iowa Electronic Markets. [Online]. tippie.uiowa.edu/iem
- [11] Infosurv Concept Exchange. [Online]. www.icepredict.com
- [12] C. A. Lacombe, J. A. Barnett, and Q. Pan, "The Imagination Market," *Information Systems Frontiers*, vol. 9, no. 2-3, pp. 245-256, 2007.
- [13] Steve Lohr, "Betting to improve odds," *The New York Times*, April 2008.
- [14] Marco Ottaviani and Peter Norman Sørensen, "Outcome manipulation in corporate prediction markets," *Journal of the European Economic Association*, vol. 5, no. 2-3, pp. 554-563, 2007.
- [15] L Page and R T Clemen. (2011, September) Do prediction markets produce well calibrated probability forecasts?. [Online]. <http://faculty.fuqua.duke.edu/~clemen/bio/PredictionMarkets.pdf>
- [16] David M Pennock, Steve Lawrence, C. Lee Giles, and Finn Arup Nielsen, "The real power of artificial markets," *Science*, vol. 291, no. 5506, pp. 987-988, 2001.
- [17] Prediction Markets - Inkling. [Online]. inklingmarkets.com
- [18] Emile Servan-Schreiber, Justin Wolfers, and David M Pennock, "Prediction markets: does money matter?," *Electronic Markets*, vol. 14, no. 3, pp. 243-251, 2004.
- [19] Nate Silver, *The Signal and the Noise: Why So Many Predictions Fail - but Some Don't*. New York: The Penguin Press, 2012, p. 335.
- [20] B Spears, C LaComb, J Interrante, J Barnett, and D Senturk-Dogonaksoy, "Examining trader behavior in idea markets: an implementation of GE's imagination markets," *The Journal of Prediction Markets*, vol. 3, no. 1, pp. 17-39, 2009.
- [21] James Surowiecki, *The Wisdom of Crowds: why the many are smarter than the few and how collective wisdom shapes business, economies, societies and nations*. New York: Doubleday, 2004.
- [22] P C Tetlock, "Liquidity and prediction market efficiency," Columbia Business School, SSRN Working Paper 929916, 2008.
- [23] Justin Wolfers and Eric Zitzewitz, Prediction Markets, April 19, 2004.