USING PREDICTION MARKETS TO HARNESS COLLECTIVE WISDOM FOR FORECASTING

By Aleksandar Ivanov

Prediction Markets (PM) method aggregates a large amount of information from various individuals to generate a forecast … the method responds to changes in market conditions quickly … it can be used for forecasting sales of new and existing products.

Prediction Markets (PM) are a collaborative forecasting tool based on methods from Decision Theory, Collective Intelligence, and Crowdsourcing. Here, participants of PM use play-money to bet anonymously on a specific company outcome. Betting influences the outcome. The outcomes may include forecasts of existing and new products, deadlines, and new product ideas. The PM collects the opinions of all the participating employees, weighs them, and then computes a single-number forecast. In practice it is done with online PM software that is deployed in the company intranet. If the incentives of a PM are well placed, it can give excellent results.

Let's say we want to predict the sales of a certain product in the fourth quarter of 2010. In PM, the forecast is represented by the price of a virtual stock; let's call it stock "Q4, 2010." During the PM forecast, which typically takes a few weeks, participants will buy or sell that stock, which will increase or decrease its price, just like in a real stock exchange. The key point is that the final payout value of that stock at the end of a forecast will be set to the actual sales in Q4, 2010, and paid out to traders in play-money. Participants will therefore buy or sell shares of that stock depending on whether they think its current price is higher or lower than their expectations about the actual sales in Q4, 2010. Obviously, we are ultimately interested in a sales forecast well before actual sales are known. The current price of the virtual stock represents that forecast. If, for example, the price of stock for Q4, 2010 comes to $55 by the buy and sell orders of traders, then this corresponds to a PM sales forecast of $55 million for that period.

ALEKSANDAR IVANOV

As Head of Advanced Methods at Analyx US and Germany, Mr. Ivanov leads the development of new forecasting methods and tools. He has implemented advanced forecasting methods for global clients in retail, telecommunications, and high-tech industries. His expertise includes econometric modeling, agent-based modeling and simulation, and prediction market design. He graduated from Dresden University of Technology with an M.Sc. in Economics.

The above incentive mechanism is what drives PM, because no one wants to own shares, which will be worth less at the end, according to his individual expectations, or miss an opportunity to buy shares, which will be worth more at the end, again according to his individual expectations. Why does this give us an excellent sales forecast? Consider the following example: Say we forecast sales of Q4, 2010 using PM. The PM trading may last 2 weeks. After 2 weeks of trading the final stock price may have reached, say $55 per share. This tells management that the sales forecast from the PM for the Q4, 2010 period is $55 million, and this information can be used for decision making well before Q4 arrives. Suppose now that actual sales in Q4, 2010 happen to be $55.4 million. The PM forecast is quite close to that. Why? For participants the final payout would be set at $55.4 per share, according to the actual sales. Hence the incentive was to sell (or short-sell) shares when the stock price was above $55.4 and to buy otherwise. Purchases and sales of traders contribute to a good forecast and the PM gives traders the incentive to forecast as close as possible.

To arrive at a good PM forecast, we need participants who are knowledgeable about what we want to forecast. For a sales forecast, we can have people from production, sales, finance, marketing, and warehouses, as well as anyone else who has knowledge about the product and the market. To initiate the PM approach you have to have a starting stock price—in this case, the initial stock price of "Q4, 2010." The starting price can be the sales of that product in Q4 of the previous year or the best guess of a sales, marketing product

manager, or combination. You need to have an incentive to forecast correctly. The incentive here is how much play-money you wind up with in the end over and above what was given in the beginning of the program. The play-money earned can be turned into real prizes, which may be a one-day paid holiday, luncheon with the president, or company-wide recognition. The PM forecasting must have a starting and ending date. In our example, for a forecast of Q4, 2010 sales, we can start the program in Q1, 2010 and end it after four weeks. If you want to forecast multiple products at the same time, you will need enough participants. As a rule of thumb, at least 30 active participants per forecast (one person may participate in more than one forecast simultaneously) are needed.

WHY PM FORECASTING?

A number of studies have shown that PM outperforms significantly other models of forecasting including surveys. Following are the four key reasons for that:

1. **Wisdom of Crowds:** Depending on the openness of your organization and the nature of the forecast, PM brings in participants ranging from a few well-chosen experts and analysts who have a specialized knowledge in the area to a large group of people from various departments including sales, marketing, production, finance, and warehouse. As such, the PM forecasting process collects a large amount of distributed knowledge as opposed to only the opinion of a few forecasters or experts (which, of course, may participate too, and even emerge as powerful traders).
2. **Incentives:** In PM those who make good predictions increase their play-money, which can later on be converted into prizes; e.g., iPods, lunch with the CEO, or recognition for being a "Trader of the Month." This encourages participants to forecast as well as they can. Also, traders who feel strongly that the price will go up will buy more shares than those who are not so confident. This way the PM not only reveals traders' true opinion but also the level of confidence they have in it.
3. **Performance-Based Weighting:** Those who make poor forecasts or deliberately try to manipulate the PM forecasts will lose their play-money as described above. Such participants, left with less play-money, can place only small bets with the result that they will have little impact on the overall PM forecast. In other words, with this mechanism, good forecasters will have more impact on a forecast than the poorer ones.
4. **Real-Time Dynamics:** Unlike surveys, which are one-time snapshots, the PM forecasts are open 24/7. If somebody knows something first that person will have a strong incentive to trade on the basis of his or her information before other traders, thereby removing the profitable opportunity from others by affecting the price. With the result, the PM forecast is updated quickly and automatically.

AREAS OF APPLICATION

PM can be applied in a wide range of areas. However, below are a few areas where it is most suited:

- **Sales Forecasting:** PM is particularly well suited for sales forecasting because it allows hundreds and thousands of people from all areas to participate in generating a forecast. Hewlett-Packard has been using it successfully since 2002 in forecasting printer sales.
- **New Product Development:** For new product developments, companies often use suggestion boxes, brainstorming, and other creative techniques. The problem with suggestion boxes is that management has to devote time and resources to assess the value of each idea. PM, on the other hand, can quickly and easily generate forecasts of key performance indicators of new products including estimated development cost and time as well as lifetime sales potential. General Electric has been using PM in several of its business units to filter out the best ideas from a large amount of suggestions.
- **Project Management:** Siemens, Microsoft, and others have been using PM to forecast a product launch date and project completion date. With PM one can predict delays in launches early on while project managers still think everything is "going fine." Nobody has an incentive to report the bad news, but the PM-based forecast can give early warning signals due to its incentive system. In PM, it is the "crowd who says so," and nobody has to take the sole responsibility of being the "messenger of bad news."
- **Market and Economic Indicators:** Forecasting market growth, market shares, and economic growth have always been difficult, and more so in an increasingly inter-connected and volatile economy. However, simple averaging of forecasts from various analysts has been shown to outperform the individual forecasts of an analyst. Since PM is based on performance-based weighted average, it gives even better forecasts.
- **Cross-Company Collaboration:** Some of the hardest forecasting problems are those that involve multiple participants both within and outside the organization. For example, a manufacturer of household appliances or consumer goods would like to know how their wholesalers and retail partners think about future demand. At times politics also stand in the way. But PM is capable of capturing all the knowledge distributed over various functions and business entities without any involvement of politics. (Again, remember those who

forecast poorly are automatically weighted down).

CASE STUDY 1
CONSUMER GOODS RETAIL

One of our own PM implementation in sales forecasting took place in a large German retail chain with over 1,000 outlets across the country. The goal was to reduce costs of sales forecasts while maintaining current level of forecasting accuracy that they had achieved (87% R^2 with actual sales outcomes). The objective here was to forecast the sales of each outlet. The cost of doing it ran several million Euros a year. The PM system, on the other hand, was much less expensive for similar or better forecasts, and cut down forecasting cost by as much as 65 percent. The process used to establish the PM process was as follows:

1. **Identify what is to be forecasted and who are the PM participants:** These two elements are inseparable, as the group of people selected depends on what to be forecasted (forecast target). The PM process won't work if nobody has relevant information about what we want to forecast. In our case, we wanted to forecast product sales. As such, we selected an excellent mix of participants from various departments, who were knowledgeable about the products of the client. The PM participants included 20 people from headquarters (HQ), mainly from the S&OP department, who had the knowledge about the market and competitors. We also picked up 70 outlet managers across the nation, who had close contact with customers regarding their tastes and preferences. In addition, we chose 10 logistics and warehouse staff members who had no direct customer knowledge but were aware of the ordering dynamics from many different outlets. As such, we selected a total of 100 people.

2. **Setup PM project team:** As with any large-scale project, the team on this PM implementation was diverse and consisted of staff from the client and our experts.
   a. Champion: The most important person for any project is the champion who drives the process. PM is no different. In our case, the head of S&OP process personally supported the project and put together a team.
   b. Project Manager: Beyond the standard operational tasks the project manager has to make sure that everybody understands what PM is and how it works. He explains the idea of PM in simple terms to different stakeholders, most importantly the future participants of the PM process.
   c. Domain Expert: A person from the S&OP staff ensures that the PM team would design a process that fits into their existing set up.
   d. Outlet Manager: The outlet managers have several representatives in the HQ, one of which has to be on the team. The participation of outlet managers is one of the most vital factors in the success of a PM.
   e. IT Staff: One IT person is needed to install the PM software in the company intranet. For the participants, the PM should be just a click away in their web browsers.
   f. PM Expert: This is usually a person with an advanced degree in economics. He or she helps to design the incentives and set up the parameters for the PM program; e.g., the initial amount of play-money to be given to each participant.
   g. PM Operator: When PM is on, it has to be monitored by someone. The operator has to watch whether the system is working as intended, participants have no problem, and incentives are working well. This role does not require much additional bandwidth. It can also be performed by one of the above team members.
3. **Integrate the PM software platform into your organization's infrastructure:** For PM participants, the system should be just a click away. The typical approach is to make it available over the intranet as a standard web application. This saves a lot of time and technical hassles, and assures that a participant, no matter where he or she is, can participate.
4. **Design the PM incentive system:** The PM expert is responsible for calibrating the PM system to determine how much play-money each participant should initially get, and how sensitive the PM price should be to buy and sell orders. The play-money is given out free. Too much free money can make people careless with their prediction, and too little makes it difficult to participate in the market. The other important PM parameter is how sensitive the stock price should be to each buy and sell order. In other words, by how much the price should change with each share bought or sold. Finding the right value for a change is important because if the market is not quite responsive, participants have to place huge orders before any significant change in the virtual stock price occurs. People might run out of money before reaching the right stock price level, which is bad for forecast accuracy. On the other hand, if the price is too responsive to orders, a small buy or sell will cause a drastic change in the virtual stock prices. With the result, the forecast would become highly volatile. Finding the right value for the sensitivity parameter is a difficult task. It depends on many unknown inputs, which only become apparent after the first weeks of trading. These inputs include the average amount of play-money spent per order, the average number of orders per trader, and the average number of traders for a forecast.

In our case, the PM calibration led to an initial $13,000 of play-money per player that could be spent on 20 sales forecasts of 20 different products, which are to be forecasted every month. Besides this, the PM expert and the representative of the outlet managers have to come up with suitable incentives and prizes that are attractive enough to participants. This is crucial because incentives are an important driver of PM performance.

5. **Train your participants:** For participants, the system should have a browser application where they can buy and sell stocks in a highly simplified trade interface. In this case, we trained all participants in two-hour workshops on how to use the system and most importantly how to make profits. Again, a thorough understanding of the PM incentive system is one of the key conditions for a PM success. Since PM is an online system, participants at any location can be trained online.

Having completed these steps the PM was ready to be launched. The PM launch process takes about two to six months, and is divided into three stages to allow for sufficient fine-tuning and feedback from participants.

1. In the first stage, which lasts six weeks, the PM system was introduced to participants. It was run in parallel with the traditional forecasting approach to compare their results. After this phase, feedback from participants was collected and incorporated in the system. The first few weeks were required for participants to get used to the logic of PM forecasting. After having observed a few payouts, participants had a full understanding of the concept and the consequences of their actions, thereby enabling us to achieve the full potential of PM.

2. In stage two, the group of participants was expanded to 300 outlet managers, 25 HQ employees, and 20 logistics staff, covering a wider range of stores for stability in the results. The new participants were trained just as initial participants, and then added to the existing group.

3. The final stage was the full-scale rollout of the PM system. The PM system at this point was not changed any further. Only the PM parameters for the sensitivity of price were adjusted for a larger crowd. Now all 1,000+ store outlets were allowed to participate in the PM forecasting process using their intranet and/or Internet browsers. This does not mean that 1,000+ people were forecasting each product. The incentive scheme made sure that people would only forecast those products in which they felt they had the knowledge.

Once the PM was established, the S&OP could use it in a straight forward way. Whenever the sales of a new product had to be forecasted, the PM operator loaded the product into the PM system. All participants would see the new product. Participants who felt knowledgeable about the product would start making bets. This collaborative interaction lasted a week or two for a given product. When the forecast was over, the stocks of each product went into a pending state until the actual sales became known. Meanwhile, the PM forecast for each product went into the S&OP process and purchase decisions were made based on the PM forecasts.

While this completes the forecasting process from a business perspective, the pending stocks in the PM still had to be monetized. Let's look at one product: Once the actual sales of that product were known, e.g., 55,000 units, the PM operator fixed the value of each stock at $55 per share. Using that price, their outstanding stocks were converted into play-money. The best forecasters made the highest profits. As such, they could bet more in future forecast rounds, and thus they had more weight in future forecasts. In our case, the people from logistics and warehouses had the best performance. While they did not have direct contact to customers they could see incoming orders from outlets. The PM enticed the logistics staff to reveal their private information for their own benefit (individual profits), as well as for the betterment of the company as a whole.

The forecasting accuracy measured by R^2 correlation of the PM forecast with the actual outcome in the six weeks of stage one, excluding one week with abnormal trading behavior, was 81 percent compared to 87 percent with the traditional forecasting model during the same time period. Management was impressed that such a "simple" crowd sourcing approach was able to generate forecasts with an accuracy comparable to their long-standing and complex method of forecasting. With this result, it was easy to convince the steering committee of PM to advance this approach to stage two to see if its performance could be further enhanced when used for more than six weeks and with a larger crowd.

Furthermore, feedback from participants showed that the PM incentives worked very well. Some participants told us that they started doing their own research to improve their forecasts. The use of this methodology also resulted in increased employee satisfaction. One participant made this statement: "Finally, those guys from product planning have started asking me what can be sold and how much." Above all, with the use of PM, one gets recognition solely based on one's performance. As such, PM was not only perceived as a useful forecasting method but also a valuable tool for increasing employees' satisfaction.

CASE STUDY 2 A GLOBAL AGRI-BUSINESS FIRM

The objective here was to see whether or not PM is also suitable for a global company with sophisticated white-collar participants, located all over the globe.

In a global firm, a great deal of localized knowledge is needed to improve the quality of forecasts. PM helps to tap into that knowledge. With the experience of a global agribusiness client, we were able to prove that the PM forecasted successfully even critical market indicators by using the knowledge of an internal group of executives, strategists, key account managers and analysts.

The objective here was to forecast the demand for seeds that would be used for each crop. Growing these seeds has a lead-time of up to one year, which is why decisions about demand have to be made up to one year ahead. There were 123 participants including executives, strategists, key account managers, and analysts. Half of these participants were from the United States and the other half from Asia, Europe, and Latin America. One of the participants was the CEO of the company. The PM was open 24/7 to allow people from all time zones to trade whenever they wanted to by just opening the secured PM website in their web browsers.

For eight forecasts, the PM forecasting process took five weeks. To validate the forecasting accuracy of this method, we first made some forecasts of the current (2008) year's U.S. acreage of some crops. The actual planted acreage of each of those crops was to be published a few weeks after the end of the PM forecast. This provided a fast feedback to top level management about PM accuracy for their business. We also showed to the management that we could not only get highly accurate forecasts but also update them in real-time as situation warrants.

Figure 1 shows how the PM forecast for the U.S. acreage for soybeans evolved from May 30 to June 15, 2008. The PM forecast was closed two weeks before the actual planted soybean acreage numbers were reported. The accuracy of the forecast was quite remarkable. It did an excellent job in aggregating the group's knowledge into a numeric forecast. But there is more to it than just accuracy of the PM forecast. As

FIGURE 1
ACCURACY OF U.S. SOYBEAN ACREAGE IN 2008 WITH PREDICTION MARKET MODEL

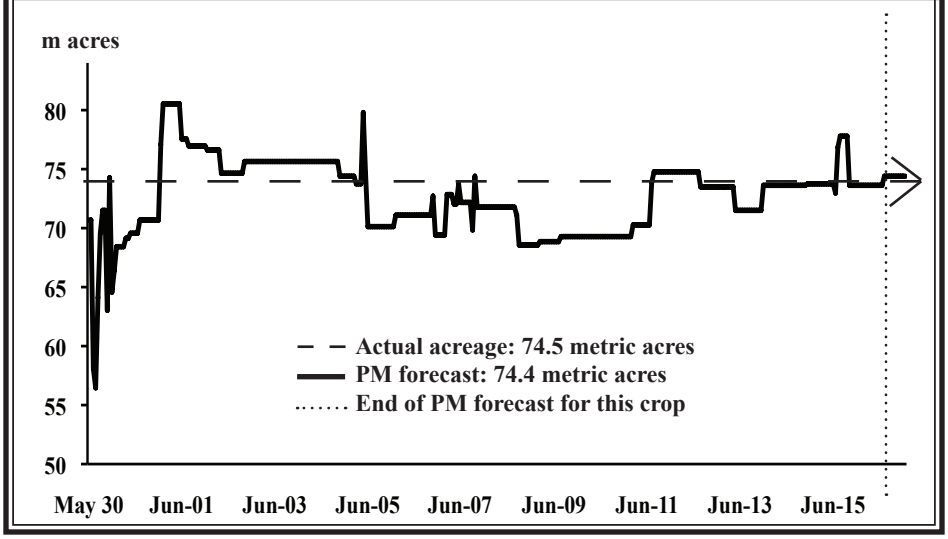

shown in Figure 1, the forecast had up and down swings over time, which reflected uncertainty during that time. The PM forecast started with around 50 million acres—the acreage of the previous year (2007). The participants knew that this was too low because farmers had planted much more in 2006 and much less in 2007. Hence, the PM forecast quickly went up to 75 million acres. It would have stayed around there, if there were no floods. But floods hit the Midwest on June 4, 2008. These floods might have destroyed some of the planted areas, making forecasting the planted acreage much more difficult. The PM forecast responded immediately with a drop from 75 to 70 million acres. The following days were characterized by up and down swings triggered by the uncertainty about the actual acreage hit by the flood. On June 7, 2008, another strong drop in the PM forecast occurred after various newspapers reported that the flood damage was quite serious. The actual flood damage was not known until the end of the PM forecast. But the PM participants were very experienced and had close contacts with farmer organizations, so they knew well in advance that the damage caused by floods was not as bad as reported in the news. To come up with a good forecast, the street-smart knowledge had to be somehow aggregated and transformed into actionable information. The PM did just that. The PM forecast went up to a more realistic value after June 9. The final forecast was very close to the actual acreage (actual 75.5 acres vs. forecast 75.4 acres) while newspapers were still speculating about the impact of floods on the planted acreage. The actual planted acreage was published by the U.S. Department of Agriculture two weeks after the end of the PM forecast.

The usefulness of this real-time capability of PM becomes apparent when one imagines how a more traditional forecasting process would have dealt with the uncertainty created by unexpected events. A traditional forecast report using scenario analysis would have to be updated after the floods. The scenarios of the report would have to be redefined and the results and conclusions had to be processed again. Alternatively, management might have refrained from issuing a report update and would have attempted to make a gut-feeling decision. Both approaches are sub-optimal. The former requires time and additional resources; the latter is likely to be less accurate. The sudden floods turned out to be an unexpected opportunity for the PM to demonstrate its

real-time forecasting capability. With the results of this first use of PM, the client decided to deploy this tool for its U.S. operations. Branches in other countries are also considering experimenting with it.

KEY SUCCESS FACTORS

While company executives can easily get excited about the potential of PM, it is important to understand the process before starting it, which I discussed earlier. The main message is that PM is not shrink-wrap software, which you just set up and the forecasts start flowing. PM is social software, which completely relies on participants' willingness to participate. There are a number of things that should be looked into before considering it.

- Are the forecasts you want to generate relevant to PM participants. Otherwise, no body would bother to spend time on it and make rational bets.
- Do participants have the relevant knowledge? If not, PM won't have useful information to aggregate.
- Do you have enough participants? More information can be aggregated with a larger group than with a smaller group. As a rule of thumb, there should be at least 30 active participants per forecast (one person may participate in more than one forecast).
- Is your organization ready for PM? PM is a Crowdsourcing method and the three rules of Crowdsourcing are openness, willingness to share knowledge, and reward those who perform well. Executives must ask themselves whether their organization's culture would support such a collaborative initiative.

CONCLUSIONS

The PM forecasting process gathers a large amount of information from many different persons and automatically weighs and combines it into a single number forecast. PM draws on an incentive system and aggregation of knowledge. With this method, forecasts are automatically updated if the situation warrants. The goal of a PM system is not to replace existing methods or experts, but to provide additional input for making better decisions. To reap its full potential, organizations must have a collaborative corporate culture. ■

(http://www.ibf.org)

SAFETY STOCK...

(Continued from page 7)

nothing more than a good guess. The more collaborative input that goes into setting a safety stock target the better should be the result.

SAFETY STOCK EXCEPTION FLAGS

After the safety stock levels have been agreed upon and implemented, it is important to maintain them. One way to do this is through a report that shows the following:

- Some items have less than 60% of the inventory target, or coverage duration less than or equal to seven days. Items that show up here imply that their safety stocks are too low.
- Some items have even less than one week of inventory. This also implies that their safety stocks are too low.
- The "Goal Post" report provides a view of the inventory position on all items. It shows whether inventory is between the min (safety stock) and max (safety stock + standard order quantity) values (i.e., between the goal posts). As such, it flags items that have either too much or too little safety stock.
- Monthly inventory points out the variations from the plan, if any. It may flag items, families, or locations where safety stock needs to be adjusted. If you see variations from the plan, the best thing is to find out what causing them.

CONCLUSIONS

In summary, remember the following while determining safety stocks:

- Consider all of the factors that influence the safety stock levels including demand variability, supply variability, and strategic components.
- There is no right or wrong answer. The best process is to agree collectively on a safety stock that is likely to minimize inventory and ensure availability of products.
- Develop a system that is suitable to your needs. Every company is different, so use a model that makes sense for you. ■

(http://www.ibf.org)