Prediction Markets: How They Can Work in Foresight

A prediction market is a competitive betting game designed to tap into the collective intelligence of a large group of participants so as to predict the occurrence of specific events in the short-term future. This approach may generate dynamic predictions that evolve in real time until the issue has been resolved.

APPLICATIONS SCOPE

Time frame: short-term events (two years at most).

Domain: politics, geopolitics, technology, regulation, actor decision-making.

Number of participants: from a few dozen traders to several thousand.

Prediction markets are **useful when (i)** knowledge is decentralized and information is distributed among many people or difficult to gather; **(ii)** new information arrives continuously, requiring forecasts to be frequently updated; **(iii)** little relevant or reliable past data exist to make projections.

TECHNICAL REQUIREMENTS

Subscription to prediction market provider: direct use or through a provider for set-up/running.

Public prediction markets: Betfair, HSX, Hypermind, Predictlt, SciCast.

Service providers: Lumenogic, Hypermind, Cultivate Labs.

RELEVANCE AND USE IN FORESIGHT

This tool application is emerging in the foresight process so there is limited feedback available. Prediction markets could serve well as a complementary tool to build and share some specific prediction or forecast on short-term events. In classical foresight, a retrospective view would shed additional light on most topics and a study may exceed a time horizon of two years.

TIME FRAME

Survey Design: approx. 1 month including test.

Implementation: 6 to 24 months depending on the event horizon.

BASIC CHECKLIST

- Check for legal clearance to run the market (legal rewards, protection of the personal data).
- Ensure topics are relevant to the foresight process.
- Set out precisely defined observable outcomes.
- Do a trial run with a small group, especially if several questions are listed.
- Maintain the quality and interest of the participant crowd.
- Animate the community during the process.
- Make sure the project or community manager has experienced a prediction market at least as a participant.

TOOL IMPLEMENTATION COSTS

Material: subscription to a prediction market service or license from a software provider.

Cost: 1 workday spread over a 1-month period for project manager to build and test question.

Additional: rewards (gifts, gift certificates or equivalent).

Communication with the panel along the survey and with the final customer(s) after.

Prospective and Strategic Foresight Toolbox

January 2017

Prediction Markets

by Émile Servan-Schreiber*

*Managing director of both Hypermind International Ltd, operator of the Hypermind.com prediction market, and Lumenogic, LLC, Émile Servan-Schreiber consults with leading companies worldwide to design and implement prediction markets and other collective intelligence applications aiming at improving forecasting and accelerating innovation. Trained in cognitive psychology at Carnegie Mellon University (Ph.D., 1991), he is the author of several research publications about human learning, prediction markets and collective intelligence.

He may be reached at ejss@lumenogic.com

The Prospective and Strategic Foresight Toolbox is a Futuribles International and CAP Prospective project
Editorial assistance: Kathryn Radford

www.futuribles.com

Abstract

A prediction market is a competitive betting game designed to tap into the collective intelligence of a large group of participants to predict specific future events or quantities. An efficient way to crowdsource forecasting on a large scale, a prediction market is particularly useful in generating dynamic predictions that evolve in real time until the issue is resolved.

The first modern prediction market began in 1988 as an academic research project at the University of Iowa's Tippie College of Business and offered forecasts on that year's US presidential election. The World Wide Web soon enabled the launch of larger prediction markets targeted at the general public, sometimes bearing other names, like betting exchanges or idea futures. Over the years, this form of "crowd wisdom" has acquired an impressive track record of accurate forecasting in diverse fields ranging from sports and film to business, elections, geopolitics and even medicine.

In its most common form, a prediction market resembles a financial market for binary options or futures. People can buy and sell predictions about a particular event occurring, *e.g.*, John Smith's winning the election. A prediction market may also yield a variable's future value, *e.g.*, candidate John Smith's share of the vote. In each case, the expiration price of the prediction will be determined by the actual outcome. In the binary example given above that could mean $100 if John Smith wins the election and $0 if he loses. The second example might translate to $1 for each percentage point of vote share.

When the outcome is still unknown, the price of a prediction may evolve dynamically in response to supply and demand from the participants within the market (the traders), according to their individual forecasts for the expiration price. As in a financial market, the transaction price is negotiated directly among the traders themselves rather than being set by a bookmaker. The price of a prediction thus captures the collective forecast and stands as the point at which the traders collectively agree to disagree. For instance, a price of $65 might predict a 0.65 probability that John Smith will be elected or will garner 65% of the vote. The trader who thinks that prediction is too pessimistic would have incentives to buy and push the price up. The trader who thinks that prediction is too optimistic would have incentives to sell and push the price down.

As in any kind of market, prediction traders are rewarded in proportion to their ability to "buy low and sell high". Accuracy, precision, and timeliness all play a crucial role therein. Indeed, prediction markets provide incentives for both the timely and truthful revelation of opinions.

If designed properly, they may also exhibit learning. Over time and many predictions, the best forecasters will profit at the expense of those whose predictions are less often correct. They will have increasingly more money to invest, thus more and more influence on the market prices. In contrast, poor forecasters' ability to set prices will progressively diminish as their capital shrinks. This continuously-updated performance-based weighing of individual opinions helps sharpen the accuracy of the market's predictions. ■

Keywords

Prediction Market ❙ Information Market ❙ Predictive Market ❙ Crowd Forecasting
Wisdom of Crowds ❙ Collective Intelligence

Description

Application Domain Resources

Prediction markets are generally used to forecast short-to-medium-term observable events, rarely more than 24 months into the future. Public prediction markets have focused primarily on sports, elections, geopolitics, medicine, science and technology. They have also been used for some macro-economic questions not traded on existing financial markets. Private prediction markets are usually run within a company, *e.g.*, Google, or with an expert panel, *e.g.*, doctors. These markets focus on various predictions relevant to business, *e.g.*, product sales, project deadlines, and industry regulation.

The market's ability to tap into collective intelligence and produce real-time forecasts is particularly useful in the following instances:

- **Knowledge is decentralized:** Information or expertise is distributed among many people, hard to gather or difficult to verbalize (implicit knowledge).
- **The situation is fluid**: New information comes into the market continuously so forecasts must be updated frequently.
- **The past is irrelevant:** Little relevant or reliable data are available to make projections.

Requirements

The requirements to run a prediction market are as follows:

- **Legal clearance to run the market.** Note that real-money markets are illegal in most jurisdictions, and subject to heavy regulation where they are legal. The alternative is to run the market as a "play-money game" which is subject to lighter regulation.
- **Questions** with precisely **defined observable outcomes.**
- **A crowd of participants,** who are somewhat knowledgeable or curious about the questions asked.
- **Performance-based incentives** to participate.
- **Software to run the market online**. Managing, matching, plotting and reporting trades in real time requires a sophisticated software platform. Fortunately, several companies do sell or rent this capability to would-be market operators.
- **One or more community managers** to formulate and create questions, rule on outcomes, manage the community of traders, and distribute rewards.

How a Prediction Market Works

In a prediction market, various predictions/outcomes are listed, as in the stock market. Bets are made by buying and selling shares of possible outcomes. A pay-off rule specifies how these shares will be valued when the question is settled. Typically, when an outcome occurs, its shares are paid $100 each, but when it fails to occur, the shares are worth nothing ($0). In the context of an election, we could list shares for the outcome: "John Smith will win the election". Should he be elected, each share would be worth $100 and if he lost, $0.

Until the question is settled, share prices vary according to the supply and demand among traders. In a classic design, called a Continuous Double Auction (CDA), everyone can offer to

buy or sell shares at his or her chosen price. These offers are on public display in an "order book", and when a buyer and a seller agree on the price, a transaction takes place. In another popular design, called the "Logarithmic Market Scoring Rule", trades are conducted with an algorithmic market maker which progressively increases the price when a trader buys shares or decreases the price when a trader sells shares.

The example plotted in Figure 1 demonstrates how the price of John Smith shares has evolved and is currently set at $80. This means a consensus prediction that John Smith is 80% likely to win the election.

If a trader buys a John Smith share at $80 and the candidate goes on to win the election, that share will be paid out at $100, yielding a profit of $20. On the other hand, if the candidate loses the election, the same share will be worthless and the trader will have lost $80.

FIGURE 1. EVOLUTION OF THE SHARE PRICE OF A FICTIONAL OUTCOME: "JOHN SMITH TO WIN THE ELECTION"

If the candidate wins, the shares will be worth $100. If he loses they will be worthless ($0). The latest price of $80 captures the most recent consensus forecast. It indicates a probability of 0.8 that John Smith will win.

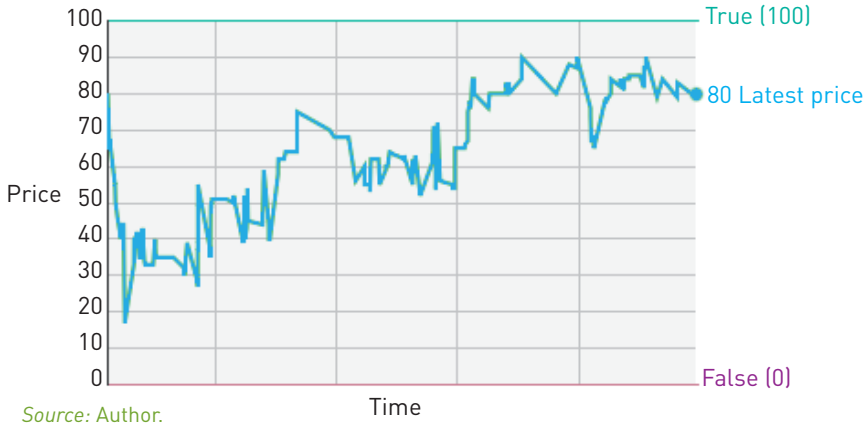

Source: Author.

Unlike a typical betting platform, the prediction market allows traders to settle their wagers even while the outcome has yet to be determined. How? Traders simply sell back the previously purchased shares. Those who buy low and sell high register a profit. The others will take a loss, but selling sooner rather than later may help them limit the damage.

For example, if a trader had bought a John Smith share at $60 before it went up to $80 a few days later, perhaps after the candidate's excellent TV-debate performance, he or she could resell it now at $80 and immediately make a $20 profit, no matter who wins the election. On the other hand, that trader would be forfeiting an opportunity to win $20 more, should John Smith eventually win and his shares be worth $100. The trader's decision will depend on a personal assessment of the chances that the candidate will win the election: higher than 80% (hold), or lower than 80% (sell)?

However, requiring traders to buy shares (which puts upward pressure on the price) before they can sell them (which puts downward pressure on the price), although highly intuitive, runs the danger of biasing the market towards higher prices and dreaded bubbles. It would be like requiring people to believe in a prediction before they are allowed to find it too optimistic. So in addition to buying and selling, the market will typically offer a third way to deal in shares: selling shares that you don't own yet. In financial markets, this is called "shorting". The market operator lends you shares against a deposit, so you can then sell them and collect the cash, but you'll have to buy the shares back later, at a moment of your choosing, to return them to the market operator against the deposit money. If you succeed in buying the shares back at a lower price than what you sold them for, you'll make a profit. The winning strategy in that case is "sell high, buy low". Shorting allows traders to bet directly against over-optimistic predictions without having invested in them first.

For example, if the outcome "John Smith to win the election" is priced at $80 and a trader thinks it is too high because John Smith is significantly less than 80% likely to win, she can borrow 1 share against a deposit of $100. When she sells the share on the market at $80, she collects $80 in cash. If the share price later falls to, say, $65, she can decide to buy the share back. Her deposit money is used for the buy-back, so the market operator gives her back the unspent portion : $100-$65 = $35. The trader sold high at $80 and bought low at $65, so she made of profit of $15.

The discussion thus far has focused on the binary question, the simplest kind possible in a prediction market. This type of question yields answers in the form of a probability, and as a basic unit of prediction may be combined to ask more complicated questions with more possible outcomes and richer answers. The most common types of questions are listed in Table 1.

TABLE 1. THE MOST COMMON TYPES OF QUESTIONS POSED TO A PREDICTION MARKET

A comparison of the last two rows demonstrates how the same question can be asked in various ways, depending on the type of answer sought: probability distribution or single-point forecast

Question Type	Market Answer	Example	Outcome(s)	Payoff Rule
Binary	Probability of the specified outcome	Will John Smith win the upcoming election?	Yes	If YES, shares are worth $100 If NO, $0
Discrete Winner-Take-All	Probability distribution over 3 or more mutually exclusive discrete alternatives	Who will win the upcoming election?	- John Smith - Jane Doe - Someone else	Shares of the winning candidate are worth $100; the other shares are worthless
Ordered Winner-Take-All	Probability distribution over 3 or more mutually exclusive alternatives on a continuum	What will John Smith's share of the vote in the upcoming election be?	- More than 55% - [50%, 55%] - [45%, 50%] - 45% or less	Shares of the correct interval are worth $100; the others are worthless
Linear	Single-point forecast for a continuous quantity	What will John Smith's share of the vote in the upcoming election be?	Vote share	Shares are worth $1 per 1% of vote share earned

Tips and Best Practices

- **Focus on short- or medium-term prediction:** Do not aim further than 24 months into the future. Beyond that point, the incentives offered to participate might prove too weak or might need to be dramatically increased to maintain traders' interest.

- **Define questions with utmost precision:** If a question or its resolution criteria are ambiguous, individual traders might forecast different things and render the collective forecast meaningless. They will also complain about the settlement, whatever it is, and lose faith in the fairness

of the market. For example, when asking who will win a US presidential race, make sure to specify whether the question is about the vote in the electoral college or about the popular vote. The more details and fine print provided in the description of the question, the better the answers.

- **Provide relevant information:** Contextual information may increase traders' interest in a question by conveying why it is important. If relevant historical information exists, sharing it helps bring all traders up to speed on the issue. If a regular flow of relevant news is available, feeding it to the traders helps keep all participants up to date in real time. In the same vein, allowing traders to interact in discussion forums increases the volume of shared information and insights.

- **Gather a diverse crowd:** The crowd of traders should include diverse points of view so that individual biases are more likely to be cancelled. Professional, geographic, ideological, and even ethnic diversity are all beneficial to market accuracy.

- **Collect a sufficiently large crowd:** The number of traders should be large enough to power a liquid and reactive market. Two constraints need to be considered: (i) each question in the market should be traded by at least a few dozen traders, and (ii) no trader can be expected to pay attention to more than 20 questions simultaneously. The greater the number of participants involved, the better the market performance, although there are diminishing returns.

- **Offer 'pseudonymity':** In most contexts, participants will be more willing to trade in unpopular or politically incorrect forecasts if their identities are protected. Full anonymity is detrimental to building a community, but letting everyone hide behind a self-chosen pseudonym is a good solution that provides cover while enabling social features such as leaderboards, forum discussions, and personal profiles.

- **Use a mix of incentives:** The incentives offered will differ according to whether the market is run with real money or play money. In real-money markets, as in financial markets, traders risk their own money in order to win other traders' cash. In that case, the profit motive alone usually incites participation from those who think they possess relevant information. In a play-money market, traders risk nothing but perhaps some of their precious time or reputation. Nonetheless they should be offered prizes or cash rewards in proportion to their play-money profits. The smaller the material rewards are, the more leaderboards, performance awards and discussion forums are needed to offer alternative incentives based on recognition and relationships.

- **Train participants:** Trading does not come naturally to most people hence any kind of training material helps boost participation, *e.g.*, one-on-one tutorials, webinars, videos or user manuals, in decreasing order of effectiveness.

Errors to Avoid

- Do not mix linear questions with other kinds of questions. The payoffs in linear questions are generally an order of magnitude lower than those in binary or winner-take-all questions. When different types of questions are featured side by side, traders usually prefer investing in the latter.

- Do not feature questions about which participants have only very limited knowledge. Prediction markets consolidate informed guesses into hard forecasts. If the guesses cannot possibly be informed, the result is garbage-in/garbage-out.

- Do not reward only the very best performers. Stock markets would not be sustainable if only the very few best investors were allowed to make a profit. For every Warren Buffet, there are

countless small investors who reap small-scale profits. The same applies to prediction markets. When only the big winners are rewarded, everyone else soon loses interest and the market dies. In real-money markets, this is a non-issue. In play-money markets, however, if only a few big prizes are featured, managers could try to sustain interest by offering raffle tickets in proportion to play-money profits.

Frequently Asked Questions (FAQs)

Are prediction markets vulnerable to bubbles, like stock markets?

No. In financial markets, a bubble occurs when a stock is significantly overpriced for no good reason. This often happens because of a fundamental asymmetry in stock markets. Basically, there is no upper limit on the price of a stock, so it is less risky to buy a stock than to sell it short. Thus the pressure on the price to rise is generally higher than the pressure for it to fall, causing bubbles to inflate sometimes dramatically. Whereas bubbles may be considered systemic in financial markets, it is not the case with prediction markets. Although mispricing can naturally occur in prediction markets, there is no asymmetry between buying and shorting. In fact, it is equally risky to buy as it is to short because there is an upper limit, *e.g.*, $100, on the value of shares. This means that overpricing and underpricing errors are equally likely. When they occur, they arise exclusively from collective mis-estimation (a good reason) rather than from mere market mechanics (no good reason).

Are prediction markets vulnerable to price manipulation?

No. Although any individual trader with enough capital (real or virtual) may temporarily move the price of a prediction dramatically higher or lower, such manipulation is unsustainable. Any significant price movement away from what seems generally reasonable to the rest of the trading population will immediately create large profit opportunities for those willing to bring the price back to where it belongs. Interestingly enough, manipulation attempts studied in controlled laboratory experiments as well as in live commercial settings concluded that the presence of manipulators tends to increase market accuracy by providing even more incentives for informed traders to participate.

Is speculation compatible with forecasting?

Yes. There are two basic ways to profit in a prediction market: either buy and hold the correct prediction until it occurs or repeatedly buy and sell to take advantage of price movements and shifts in the collective opinion. The former is closer to pure forecasting, while the latter is akin to pure speculation with little attempt at an honest forecast. Yet, both are essential to a well functioning market. First, speculators play the role of market makers, providing trading opportunities for informed buy-and-hold forecasters. Second, over the term of a question, as ground conditions evolve, few traders are comfortable being locked into a pure buy-and-hold strategy. In fact, most participants will become speculators themselves occasionally or temporarily. This flexibility contributes to the market's reactivity.

Are real-money markets more accurate than play-money markets?

No. There is no evidence to support the intuition that prediction-market accuracy depends on participants' putting their money where their mouth is. Empirical comparisons of real-money and play-money prediction markets have found equivalent levels of accuracy. There are pros and cons to each paradigm, though. In a real-money market, the possibility of losing money, rather than just some prestige or a shot at some prize, may focus the mind of those who choose to participate anyway, and deter those who do not know enough about the questions listed. On

the other hand, in a play-money market, the net worth of each participant and his or her ability to impact prices are closely correlated with their past success in the market. This is not the case in real-money betting when the richest may not necessarily be the wisest.

Case Study: Hypermind

Hypermind is a prediction market that sells forecasts to asset management firms, multinational corporations, government agencies, and media organizations. It was created by Lumenogic, the founders of the pioneering prediction market NewsFutures (2000-2010), following a long-running collaboration with the Good Judgment Project (GJP), winner of the Aggregative Contingent Estimation (ACE) geopolitical forecasting tournament sponsored by the American government's Intelligence Advanced Research Projects Activity (IARPA). This case study reviews how Hypermind integrates best practices and state-of-the-art technologies to provide reliable predictions.

- **Legal clearance:** Hypermind is set up on the Internet as a play-money prediction game. No purchase is necessary to participate. Although material rewards are distributed according to performance, the game is legal and very lightly regulated in the vast majority of countries and jurisdictions.

- **Questions:** The prediction market focuses on geopolitics, geo-economics, and client-defined questions related to business. Most questions are binary or winner-take-all, so predictions are expressed as probabilities. The time horizons of the questions vary from 1 to 24 months.

- **Traders:** The Good Judgment research revealed that forecasting is a stable cognitive skill that depends less on what participants know (prior domain expertise) than on how participants think (general intelligence, curiosity, and open-mindedness). Hypermind began by actively recruiting a panel of several thousand highly skilled forecasters from among the 20% top performers in various political, geopolitical, economic and scientific prediction-market contests run by NewsFutures and Lumenogic between 2000 and 2014. The company then enriched this core group by partnering with various domain-relevant blogs and media organizations in order to capture a continuous flow of promising new recruits. Participation is by invitation only, and candidates must score sufficiently high on a test of actively-open-minded thinking (AOT), a key psychological trait of good forecasters. Once admitted to the panel, traders are rewarded solely on actual performance. As a result, good forecasters thrive while poor ones quickly drop out. Hypermind traders are competitive, highly educated, cross-generational, professionally diverse, and embedded all around the world.

- **Incentives:** The market's currency is play money denoted as H for hyper money. Everyone begins with a one-time grant of 100,000 H and no possibility of refueling. In this zero-sum game, every H earned has been lost by someone else, so any trader's net worth closely tracks his or her relative forecasting ability. Various prediction contests on different topics are offered in parallel. Traders may choose the contests in which they wish to participate. They share cash prizes of several thousand Euros per contest *pro rata* with their virtual profits and level of participation. These material incentives are complemented with recognition and relationship rewards offered by leaderboards and discussion forums.

- **Software:** Hypermind uses Lumenogic's prediction-market software. The trading engine is a classic continuous double auction (CDA), but with liquidity enhanced by a sophisticated frictionless arbitrage algorithm that enables trades to be matched across any number of rival outcomes.

Community management: Another strong conclusion from the Good Judgment Project is that prediction accuracy improves when forecasters share information and feel part of a team. Hypermind traders are therefore encouraged to partake in forum discussions. Team spirit is further fostered by regular reminders that better collective accuracy translates into more paying sponsors and larger cash prizes. Hypermind also provides traders with a continuous feed of the latest relevant news articles automatically selected from the Google News service.

How Accurate Are the Predictions?

Hypermind's forecasts are probabilistic. They answer the question: "What are the chances that this or that outcome will come true?" The best way to assess accuracy in the absolute is over many events, by comparing the probability estimations against observed event frequencies.

FIGURE 2. COMPARISON OF HYPERMIND'S ESTIMATED EVENT PROBABILITIES WITH ACTUAL FREQUENCIES OF OCCURRENCE

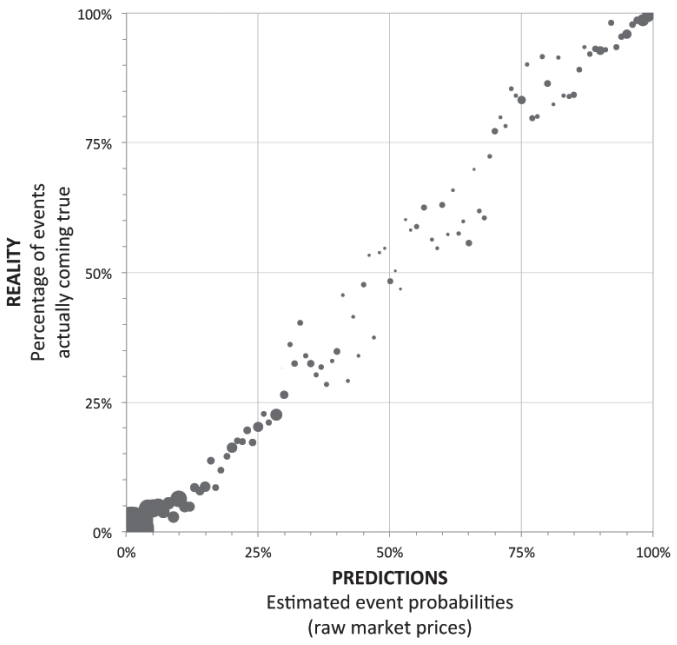

The probabilities were collected daily at a fixed time for all the outcomes traded on the market that day. When all the questions were resolved, it was possible to compute the proportion of outcomes that occurred at each level of estimated probability. Note that the size of the data points indicates the relative number of data collected at each level. This chart plots 50,463 probability estimations produced over 27 months, on 189 questions with 500 possible outcomes. See text for full discussion.

Source: Hypermind, August 29, 2016.

Figure 2 plots this comparison at every level of probability from 1% to 99%, over 27 months, on 189 questions with 500 possible outcomes. International topics ranged from elections, geopolitics, geo-economics, and business. The figure allows us to answer the following question: "Of all the events predicted by Hypermind to be *x*% likely, what percentage actually occurred?" The closer to *x*% the answer; *i.e.*, the closer the data points sit to the chart's diagonal (bottom-left to top-right), the more accurate the market's probabilities.

Another way to evaluate accuracy compares the market's results to other forecasting methods for the same events. Hypermind has generally outperformed polls and big-data statistical models in instances where comparison was possible. Recent examples include Hypermind's predicting the defeat of the Scottish referendum on independence, the results of the American 2014 Midterm elections (see Servan-Schreiber and Atanasov, 2015) and the 2015 electoral victories of Netanyahu in Israel and Cameron in the UK. The comparative results on Brexit were more ambiguous (see Servan-Schreiber, July 2016). Hypermind predicted Donald Trump's Republican Party nomination in 2016 earlier and more steadfastly than all the leading real-money prediction markets (see Servan-Schreiber, June 2016) and proved "less wrong" than most other methods in the general election (see Atanasov and Joseph, November 2016). In a head-to-head forecasting competition sponsored by IARPA on 36 geopolitical questions, Hypermind demonstrated the same accuracy as the Good Judgment Project's so-called "super forecasters", the current gold standard in geopolitical forecasting.

ACCURACY IN MARKET PREDICTIONS

Market predictions are typically expressed as probabilities, which means that accuracy can be measured only through many predictions. Accuracy may be defined as a product of both calibration and discrimination.

Calibration

Predictions are said to be well calibrated when the more probable events occur more often and the less probable events occur less often. For example, if we consider all the events to which the market ascribed 30% probability, we should observe that 30% of them actually did take place. Similarly, if we consider all the events to which the market ascribed 80% probability, we should observe that 80% of them indeed took place.

Discrimination

Discrimination refers to how extreme the predictions are. The closer the predictions are to 0% or 100%, the more discriminating the predictions are said to be.

Only an all-knowing god's predictions could be both perfectly calibrated and discriminating! Events would always be predicted to be 0% likely or 100% likely, and the prediction would always be correct. Baring such divine perfection, calibration is preferable to discrimination. In other words, a fuzzy but correct forecast is better than a categorical but incorrect forecast.

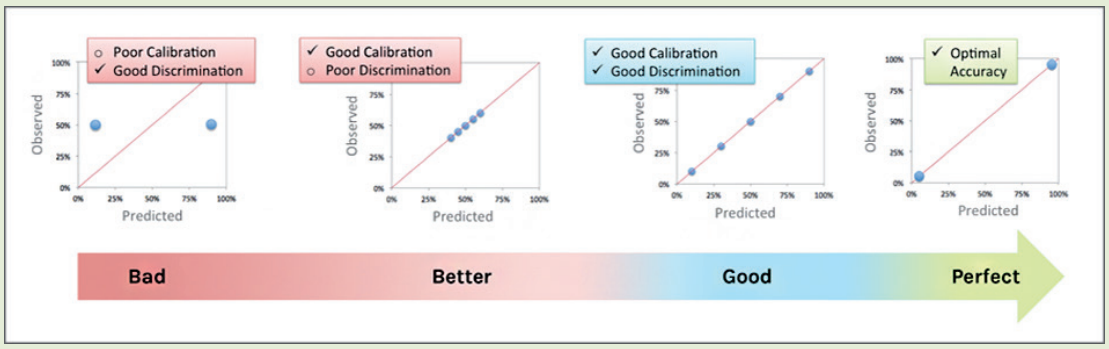

Traditionally prediction markets have demonstrated remarkable calibration; whereas, the level of discrimination attained has depended on the difficulty of the questions. ■

Source: https://hypermind.com/hypermind/app.html?fwd=#research. Accessed January 18, 2017.

Further Reading

Related Techniques

Despite various attractive features, all prediction markets share one significant drawback: the intrinsic complexity of the trading paradigm. This often proves a hurdle to participation, especially in a corporate setting. Therefore, several practitioners have developed simpler types of betting contests that still reward participants for their level of accuracy, precision and timeliness yet set aside the market features.

This competitive-forecasting approach includes the Good Judgment Project's "prediction polls", which collect forecasts through a classic survey and assess them with Brier scores,

or Lumenogic's CHIPS software platform, which allows participants to distribute a limited number of tokens among various outcome bins. One advantage of this approach over prediction markets is that the consensus forecast can be computed in the background and treated statistically without ever being published or revealed to the participants. Decision-makers in corporate settings often find this feature attractive. A drawback for participants, however, is that scoring and performance assessment are less transparent and intuitive than in a prediction market. (For more information, see the 2016 article by Atanasov *et al.* in the bibliography.)

Current Research

The first 25 years of research in prediction markets focused on documenting their prediction accuracy and applicability in different contexts ranging from academic laboratories to corporate settings but also including general public applications in various areas. Currently, the most interesting research highlights the following areas:

- **Combinatorial markets:** The classical prediction markets discussed above become cumbersome quickly when instead of predicting independent events, the goal becomes computing a probability for an outcome that is conditional on other outcomes also occurring. Combinatorial markets seek to overcome this practical limit by using algorithmic market makers that link bets in various related questions through Bayesian inference networks. (For more information, see work by Hanson, and Lahaie *et al.*, and the SciCast website listed in the bibliography.)

- **Individual differences:** The proliferation of public prediction markets that forecast the same events, *e.g.*, the US elections and Brexit, has shown that their forecasts may differ significantly, just as the polling number put forward by various institutes may differ for the same election. Research is needed to understand what causes these differences. Key questions remain, such as *are these differences caused by a market mechanism, trader profiles, the reward system? How predictable might they be? How best to aggregate the forecasts of different types of prediction markets?* (For more information, see the PredictWise website listed in the bibliography.)

- **Hybrid markets:** In many sectors, *e.g.*, chess, financial trading or weather forecasting, the smartest agents are neither human brains nor artificial intelligence systems but rather a combination of both. Prediction markets have relied thus far almost exclusively on human collective intelligence. Hence the question how might artificial intelligence and computer models best be combined with human judgment in hybrid markets? (For more information, see the IARPA/HFC website in the bibliography.)

- **Argumentative markets:** To the frustration of decision-makers, prediction markets generate well calibrated probabilities but remain unable to explain their reasoning. A new line of applied research seeks to augment prediction markets with formal argumentative capabilities. (For more information, see the IARPA/CREATE website in the bibliography.)

Selective Bibliography

Principles and Mechanisms

Arrow Kenneth, Forsythe Robert, Gorham Michael, Hahn Robert, Hanson Robin, Ledyard John, Levmore Saul, Litan Robert, Milgrom Paul, Nelson Forrest, Neumann George, Ottaviani Marco, Schelling Thomas, Shiller Robert, Smith Vernon, Snowberg Erik, Sunstein Cass, Tetlock Paul, Tetlock Phillip, Varian Hal, Wolfers Justin, and Zitzewitz Eric, "The Promise of Prediction Market", *Science*, 320(5878), May 2008, pp. 877-878, DOI: 10.1126/science.1157679.

Atanasov Pavel, Rescober Phillip, Stone Eric, Swift Samuel, Servan-Schreiber Émile, Tetlock Phillip, Ungar Lyle, and Mellers Barbara, "Distilling the Wisdom of Crowds: Prediction Markets vs. Prediction Polls", *Management Science* (Articles in Advance), 2016.

Hahn Robert, and Tetlock Paul (Eds), *Information Markets: A New Way of Making Decisions,* Washington, D.C.: AEI (American Enterprise Institute)-Brookings Joint Center for Regulatory Studies, 2006.

Hanson Robin, "Combinatorial Information Market Design", *Information Systems Frontiers,* 5(1), 2003.

Lahaie Sébastien, Dudík Miro, Rothschild David, and Pennock David, "A Combinatorial Prediction Market for the U.S. Elections", *in Proceedings of the Fourteenth ACM Conference on Electronic Commerce, Philadelphia, June 16-20, 2013,* New York: ACM (Association for Computing Machinery), 2013, pp. 341-358.

Servan-Schreiber Émile, Wolfers Justin, Pennock David, and Galebach Brian, "Prediction Markets: Does Money Matter?", *Electronic Markets,* 14(3), 2004.

Servan-Schreiber Émile, "Prediction Markets: Trading Uncertainty for Collective Wisdom", *in* Hélène Landemore, and Jon Elster (Eds), *Collective Wisdom: Principles and Mechanisms,* New York: Cambridge University Press, 2012.

Wolfers Justin, and Zitzewitz Eric, "Prediction Markets", *Journal of Economic Perspectives,* 18(2), Spring 2004, American Economic Association.

Applications

Cowgill Bo, and Zitzewitz Eric, "Corporate Prediction Markets: Evidence from Google, Ford, and Firm X", *The Review of Economic Studies,* 82(4), 2015.

Elberse Anita, and Anand Bharat, "The Effectiveness of Pre-release Advertising for Motion Pictures: An Empirical Investigation Using a Simulated Market", *Information Economics and Policy,* 19(3-4), 2007.

Polgreen Phillip, Nelson Forrest, and Neumann George, "Use of Prediction Markets to Forecast Infectious Disease Activity", *Clinical Infectious Diseases,* 44(2), 2007.

Rothschild David, "Forecasting Elections: Comparing Prediction Markets, Polls, and their Biases", *Public Opinion Quarterly,* 73(5), 2009.

Hypermind (Case Study)

Atanasov Pavel, and Joseph Regina, "Which Election Forecast Was the Most Accurate? Or Rather: The Least Wrong?", *The Washington Post,* November 30, 2016.

Servan-Schreiber Émile, "Hypermind Wins the 2016 Republican Nomination Race", *Hypermind Gazette,* June 2016, URL: https://blog.hypermind.com/2016/06/01/hypermind-wins-the-2016-republican-nomination-race/. Accessed January 18, 2017.

Servan-Schreiber Émile, "Brexit Pundits Embarrass Themselves with Misguided Attacks on Prediction Markets", *Hypermind Gazette,* July 2016, URL: https://blog.hypermind.com/2016/07/04/brexit-pundits-embarrass-themselves-with-misguided-attacks-on-prediction-markets/. Accessed January 18, 2017.

Servan-Schreiber Émile, and Atanasov Pavel, "Hypermind vs. Big Data: Collective Intelligence Still Dominates Electoral Forecasting", *in Proceedings of the 2015 Collective Intelligence Conference, Santa Clara, May 31-June 2, 2015*. URL: http://sites.lsa.umich.edu/collectiveintelligence/wp-content/uploads/sites/176/2015/02/Servan-Schreiber-and-Atanasov-CI-2015-Abstract.pdf. Accessed January 19, 2017.

Notable Websites (Alphabetical Order)

- Betfair: http://www.betfair.com

The world's largest betting exchange. Covers mostly sports, but also some US, UK, and European elections.

- Hollywood Stock Exchange: http://www.hsx.com

Play-money market specialized in the US box office and other entertainment forecasts, since 1996.

- Hypermind: https://hypermind.com

High-performance play-money prediction market specialized in business-relevant geopolitics and economics.

- IARPA/CREATE: https://www.iarpa.gov/index.php/research-programs/create

Presentation of the CREATE research program (CRowdsourcing Evidence, Argumentation, Thinking and Evaluation) on the website of the Intelligence Advanced Research Projects Activity (IARPA).

- IARPA/HFC: https://www.iarpa.gov/index.php/research-programs/hfc

Presentation of the HFC research program (Hybrid Forecasting Competition) on the website of the Intelligence Advanced Research Projects Activity (IARPA).

- Iowa Electronic Markets: https://tippie.biz.uiowa.edu/iem/

The original real-money political market run by the University of Iowa's business school. Covers mostly US national politics.

- PredictIt: https://www.predictit.org

The largest US-based real-money political market. Specializes in American politics.

- PredictWise: http://www.predictwise.com

A prediction-market data aggregator (mostly Betfair, Hypermind and PredictIt). Focuses on U.S. politics, sports and entertainment.

- SciCast: https://scicast.org

A rare combinatorial prediction market (play money) focusing on science and technology. Run by George Mason University and initially sponsored by IARPA.

Interview

Servan-Schreiber Émile, and Marks Aubrey, "Live Webinar Featuring Dr. Emile Servan-Schreiber", *PredictIt Blog*, April 2016, URL: http://blog.predictit.org/post/143214662052/webinar-with-prediction-market-expert-dr-emile. Accessed January 24, 2017.

