

Methods to Elicit Forecasts from Groups: Delphi and Prediction Markets Compared

Kesten C. Green
Business and Economic Forecasting Unit, Monash University,
kestenc@kestencgreen.com

J. Scott Armstrong,
The Wharton School, University of Pennsylvania
armstrong@wharton.upenn.edu

Andreas Graefe
Institute for Technology Assessment and Systems Analysis,
Research Center (Forschungszentrum) Karlsruhe
graefe@itas.fzk.de

Forthcoming in *Foresight: The International Journal of Applied Forecasting* (Fall 2007)

This working paper offers a more detailed treatment of the issue than was provided in the published version.

September 22, 2007

Abstract

Traditional groups meetings are an inefficient and ineffective method for making forecasts and decisions. We compare two structured alternatives to traditional meetings: the Delphi technique and prediction markets. Delphi is relatively simple and cheap to implement and has been adopted for diverse applications in business and government since its origins in the 1950s. It can be used for nearly any forecasting, estimation, or decision making problem not barred by complexity or ignorance. While prediction markets were used more than a century ago, their popularity waned until more recent times. Prediction markets can be run continuously, and they motivate participation and participants to reveal their true beliefs. On the other hand, they need many participants and clear outcomes in order to determine pay-offs. Moreover, translating knowledge into a price is not intuitive to everyone and constructing contracts that will provide a useful forecast may not be possible for some problems. It is difficult to maintain confidentiality with markets and they are vulnerable to manipulation. Delphi is designed to reveal panelists' knowledge and opinions via their forecasts and the reasoning they provide. This format allows testing of knowledge and learning by panelists as they refine their forecasts but may also lead to conformity due to group pressure. The reasoning provided as an output of the Delphi process is likely to be reassuring to forecast users who are uncomfortable with the "black box" nature of prediction markets. We consider that, half a century after its original development, Delphi is under-utilized.

Much can be done to improve upon traditional group meetings. As Armstrong (2006) showed, it is difficult to think of a structured approach that would not improve on the predictions and decisions made in traditional meetings.

Rowe (2007) presents evidence that, in comparison with traditional meetings, the Delphi technique can improve forecasting and decision making. How does it do that? If conducted properly, Delphi greatly improves the chances of obtaining unbiased estimates and forecasts that take full account of the knowledge and judgment of experts. A third method for aggregating individual judgments, prediction markets, is based on the idea that markets provide forecasts by giving people an opportunity to benefit financially or otherwise by buying, or selling, when the current price seems too low, or high, given the knowledge they possess. The market price thus reveals the aggregate knowledge of market participants. We discuss the relative merits of Delphi and prediction markets.

How Delphi Has Been Used

The Delphi procedure has been around since the 1950s. To assess its use, we conducted a Google search for “Delphi AND (predict OR forecast)”. This yielded 805 unique sites out of a total of 1.4 million, showing that some people have paid attention.

Using the same keywords, we conducted searches in the Social Sciences Citation Index and the Science Citation Index Expanded to assess what has been happening to researcher interest in Delphi over the years. We identified altogether 65 relevant items; 1 from the 1960s, 8 from the 1970s, 3 from the 1980s, 21 from the 1990s, and 32 so far this decade.

When we searched for “Delphi forecast of” and “Delphi forecasts of”, we found 42 unique applications of the Delphi technique. Most of them (43%) were business applications. These included forecasts for:

- the Argentine power sector
- broadband connections
- dry bulk shipping
- leisure pursuits in Singapore
- rubber processing
- Irish specialty foods, and
- oil prices.

Forecasts of technology were also popular (36%), these included forecasts about intelligent vehicle-highway systems, industrial robots, intelligent internet, and technology in education. Finally, 21% of applications were concerned with broader social issues such as the “urban future” of Nanaimo in British Columbia, and the future of law enforcement.

We also found about 1,000 unique items using a Google Scholar search for the “Delphi method” and either “forecast OR predict OR estimate.” This also suggests that Delphi is widely used.

We have ourselves employed Delphi for problems ranging from forecasting prisoner number, to choosing between regional development options, to predicting outcomes of political elections, to deciding which applicants should be hired for academic positions, to predicting how many meals to order at conference luncheons.

How Delphi Might Be Used

Delphi can be used for nearly any problem involving forecasting, estimation, or decision making—as long as complexity and ignorance do not preclude the use of expert judgment. In short, it could be used to replace most face-to-face meetings other than those involving negotiations or selling.

The issue of ignorance is important. If the individuals in a group are misinformed about a topic, the use of Delphi will, like a traditional group meeting, add only confidence to their ignorance. However, uncovering disparity among the experts might help to alert decision makers to this problem. For example, in a Delphi study of economic growth, support for higher education was believed to be a positive factor by most participants, while a small minority claimed it was negative. This issue should have been decided by reference to the research literature rather than by asking experts.

People are not good at thinking through complex situations, such as those that involve several rounds of interactions with others. Green and Armstrong (2007) showed that unaided experts are unable to provide valid forecasts about the outcomes of negotiations and other conflict situations. The Delphi process cannot improve forecasts when the individual panelists are incapable of providing valid forecasts.

With Delphi, experts are asked to provide reasons for their forecasts, and to respond to the forecasts and justifications given by the other experts. This information is then used by the experts to revise their forecasts or estimates. By this line of reasoning, the feedback from additional rounds would only be useful when the participants can offer new information to one another. If they are all drawing upon a common pool of knowledge, a one-round expert survey would be sufficient. The latter is often referred to as the Nominal Group Technique or NGT.

In our experience, feedback provided by experts is attractive to those clients who are skeptical of forecasts from a statistical model. Hoffmann et al. (2007) observed that the findings of their survey of expert opinions on the distribution of food-borne illnesses in the US were met with skepticism until their audiences saw the list of expert participants.

Rowe (2007) indicates that Delphi can be expensive; but is it expensive in comparison with traditional group meetings? We like the taxi-meter solution to meetings: Each person attending a meeting enters a billing rate into a computer and the computer shows the mounting cost as the meeting grinds on.

When high expert status is not needed to help sell the forecasts, only modest expertise is required. This means that expenses can be kept low and that forecasts can be made rapidly.

Freeware for conducting Delphi sessions is available at forecastingprinciples.com (under Software). When the forecast question is clear and panelists are cooperative, the software helps the administrator to complete a session in quick time. The software is used to compile questions, store a list of potential panelists and their email addresses, send appeals to panelists, and compile responses. The software also provides guidance on how to use Delphi. The forecastingprinciples.com directors continue to increase the flexibility of the Delphi software to allow greater customization.

One way to reduce the cost of traditional group meetings is to use Delphi procedures within the meeting, a process known as *Mini-Delphi* or *estimate-talk-estimate*. This also helps to ensure that people provides their estimates during the “meeting.” To further speed up traditional meetings, Gordon and Pease (2006) developed *Real-Time Delphi*, a web-based approach that automatically aggregates participants’ judgments and allows them to reassess their positions. Real-Time Delphi appears promising, but it has not yet been evaluated.

Resurgence of Prediction Markets

In recent years, there has been a resurgence of interest in prediction markets, which were quite popular in the late-1800s and early-1900s (Rhode and Strumpf 2004). In her *Businessweek* article, King (2006) claimed that at least 25 companies had started to experiment with prediction markets. Prediction market forecasts have proved to be accurate in limited tests to date. An internal market at Hewlett-Packard on future product sales, for example, beat the official forecasts of the company in 6 out of 8 events (Chen & Plott 2002). Researchers are also doing more in this area and in response to this interest the *Journal of Prediction Markets* was launched in 2007.

Participants in prediction markets buy and sell contracts. These contracts promise a payoff if an event occurs. In their entry on prediction markets for the *New Palgrave Dictionary of Economics* (2nd ed.), Wolfers and Zitzewitz (2006) provided a useful summary of the method. They tabulated three different types of contract—binary option, index futures, and spread betting—each designed to provide a different kind of forecast. In the case of a binary option market, the price at which a contract most recently traded (or an average of the most recent prices) is interpreted as the market's assessment of the probability that the event will occur. For example, if a contract will pay \$1 in the event of Britain withdrawing more than 50% of her troops from Iraq before the end of 2007 and nothing otherwise and the contract last traded at 22 cents, then the market's assessment is that the likelihood of that event is 0.22.

Delphi vs. Prediction Markets

Prediction markets are similar to Delphi in that they are both methods for aggregating diverse opinions. The relationship is illustrated by the Methodology Tree at forecastingprinciples.com. Although both provide forecasts that are more accurate than those from unstructured group meetings, there is no empirical evidence on the relative accuracy of forecasts from the two methods and the relative effects of different conditions.

Prediction markets may avoid some potential drawbacks of Delphi:

1. While Delphi can help panelists reach an understanding, such an understanding can arise as a result of group pressure to conform and may conceal disagreement (Woudenberg 1991). By contrast, participants in markets can only benefit if they challenge the group opinion and so they are motivated to reconsider their judgments and to reveal their true beliefs.
2. Delphi is typically conducted as a one-off activity providing results at a certain point in time. Although new forms like Real-time Delphi enable live feedback, the question remains how panelists can be motivated to participate over longer periods. Participants are motivated to keep participating in markets and markets can be run continuously so that new information is instantly and automatically incorporated into the forecast.
3. It can sometimes be difficult to recruit suitable experts to participate on a Delphi panel. In the case of markets, participants themselves choose to take part if they think that their private information is not yet incorporated in the forecast.

Yet, the applicability of prediction markets is limited insofar as they require a situation whose outcome will eventually be known. Without a clear outcome, such as the percentage of votes gained by a candidate, the sales figures for a given time period, or the annual growth in GDP, participants could not be appropriately rewarded or punished.

Delphi seems to have these advantages over prediction markets:

1. It can be used for a much broader range of problems, since there is no need to judge the outcome of a situation in order to determine pay-offs for participants.

2. Many people lack the understanding of how markets work or how to translate their expectations into market prices. It is easier for people to reveal their opinions in Delphi.
3. It can be challenging, if not impossible, to formulate some problems as contracts in prediction markets. It is easier to address complex issues and to obtain predictions by asking direct questions of a Delphi panel.
4. It is easier to maintain confidentiality with Delphi. For markets, it may be morally objectionable to benefit from trading on the outcome of critical issues. For example, the policy analysis market set up by the DARPA to predict events like regime changes in the Middle East or the likelihood of terrorist attacks was cancelled one day after it was announced (Looney 2004). Concerns may also arise over the use of markets within businesses, for example to decide who to hire or fire, or where the forecast may de-motivate participants who are also employees of the business.
5. Prediction markets are vulnerable to speculative attacks mounted in order to manipulate the results. A Delphi administrator, on the other hand, can choose panelists who are likely to reveal their true beliefs and exclude extreme values from the calculation of the Delphi forecast either directly or by calculating a median rather than a mean.
6. The opportunity to provide comments or reasons for judgments allows Delphi participants to introduce new ideas into the discussion. And the transparent exchange of knowledge allows experts to learn while participating in the Delphi process.
7. The transparent exchange of knowledge reveals information that has already been taken into account. This helps Delphi panels avoid two undesirable features of predictions markets: the inefficiency of each participant independently searching for information, and *cascades*. A cascade is a cumulative and excessive price movement that occurs when some participants, assuming that a shift in price is due to new information, react; and then others react to the reactions.
8. Delphi requires only 5 to 20 experts who have agreed to participate and should therefore be superior to thin markets (those with few participants) where the incentive to trade, and thereby reveal information, is weak (Abramowicz 2004).

In sum, by providing incentives for participation and automatically aggregating information, we believe that in situations where markets are feasible they have some advantages over Delphi. Markets should be superior for short-term problems that are easy to address, when it is possible and desirable to involve a large number of participants, and when the primary goal is obtaining a group response. In contrast, the uses of Delphi are not limited to providing reliable and accurate forecasts. Delphi is also valuable because it can be used in situations of high uncertainty by enabling communication, encouraging discussion, and generating new ideas.

We believe that Delphi (or, on occasion, the simpler NGT) should be much more widely used than it is today. It should replace many traditional meetings. Provided that it does not drive out other valid structured methods, it is unlikely to cause harm and will likely improve forecasting and decision making—and thus increase the efficiency and effectiveness of organizations.

References

- Abramowicz, M. B. (2004). Information Markets, Administrative Decisionmaking, and Predictive Cost-Benefit Analysis. *University of Chicago Law Review*, 71, 933-1020.
- Armstrong, J. S. (2006). How to Make Better Forecasts and Decisions: Avoid Face-to-Face Meetings. *Foresight*, 5, 3-8.

- Chen, K.-Y. & Plott, C. R. (2002). Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem. *Social Science Working Paper No.1131*, California Institute of Technology, Pasadena.
- Gordon, T. & Pease, A. (2006). RT Delphi: An Efficient, “Round-Less” Almost Real Time Delphi Method. *Technological Forecasting and Social Change*, 73, 321-333.
- Green, K. C. & Armstrong, J. S. (2007). The Value of Expertise for Forecasting Decisions in Conflicts. *Interfaces*, 37, 287-299.
- Hoffmann, S., Fischbeck, P., Krupnick, A. & McWilliams, M. (2007). Elicitation from Large, Heterogeneous Expert Panels: Using Multiple Uncertainty Measures to Characterize Information Quality for Decision Analysis. *Decision Analysis*, 4 (2), 91-109.
- King, R. (2006). Workers, Place Your Bets, *BusinessWeek*, August 3, http://www.businessweek.com/technology/content/aug2006/tc20060803_012437.htm.
- Looney, R. E. (2004). DARPA's Policy Analysis Market for Intelligence: Outside the Box or Off the Wall? *International Journal of Intelligence and Counterintelligence*, 17, 405-419.
- Rhode, P. W. & Strumpf, K. S. (2004). Historical Presidential Betting Markets. *The Journal of Economic Perspectives*, 18, 127-141.
- Rowe, G. (2007). A Guide to Delphi. *Foresight*, 8, forthcoming.
- Rowe, G. & Wright, G. (2001). Expert opinions in Forecasting: The Role of the Delphi Technique. In: J. S. Armstrong (Ed.), *Principles of Forecasting - A Handbook for Researchers and Practitioners*. Boston, MA; Kluwer Academic Publishers, pp. 125-144.
- Wolfers, J. & Zitzewitz, E. (2006). Prediction Markets in Theory and Practice. NBER Working Paper 12083, [http://bpp.wharton.upenn.edu/jwolfers/Papers/PredictionMarkets\(Palgrave\).pdf](http://bpp.wharton.upenn.edu/jwolfers/Papers/PredictionMarkets(Palgrave).pdf).
- Woudenberg, F. (1991). An Evaluation of Delphi, *Technological Forecasting and Social Change*, 40, 131-150.