A Proposal for the Use of Prediction Markets to Forecast Influenza Activity

Philip M. Polgreen, M.D.[1]
Daniel J. Diekema, M.D.[1,2,3]
Loreen A. Herwaldt, M.D.[1,3]
Forrest D. Nelson, Ph.D.[4]
George R. Neumann, Ph.D.[4]

[1]Department of Internal Medicine and [2]Department of Pathology
The University of Iowa Carver College of Medicine
[3]Program of Hospital Epidemiology
University of Iowa Hospitals and Clinics
Iowa City, Iowa
[4]Department of Economics
The University of Iowa Henry B. Tippie College of Business
Iowa City, Iowa
Type of Article: Hypothesis

Word Count: 1490

Abstract Word Count: 130

Corresponding Author:

Philip M. Polgreen, M.D.
SW 54 GH, Department of Medicine
University of Iowa Carver College of Medicine
200 Hawkins Drive
Iowa City, IA 52242
Phone: 319-356-67227; Fax 319-356-4600

Abstract

We propose that prediction markets may be useful for predicting influenza activity by aggregating the expert opinion of healthcare professionals quickly, accurately, and inexpensively. Experimental markets have been successfully used by the *Iowa Electronic Markets* (IEM) to predict other outcomes ranging from election results to movie box office receipts. Since 1988, the IEM has developed an impressive prediction record, substantially superior to alternative mechanisms such as polls and surveys. Prediction markets aggregate disparate information from multiple participants who indirectly reveal their personal knowledge when they trade. Available evidence from over 15 years of running experimental markets suggests that prediction markets can outperform both refined statistical forecasting methods and well-designed surveys. In addition, prediction markets could be used for tracking and forecasting other emerging infectious diseases from SARS to avian influenza.

Introduction

Although influenza occurs on an annual basis, unique characteristics particular to each influenza season make forecasting difficult. The geographic locations, rates of increase and decline, duration, and size of each outbreak vary considerably from year to year. Although statistical models using historical data might accurately describe the typical pattern, they do not predict departures from the norm.[1] Yet it is the deviations that are of the most concern and, therefore, the most important to predict.

Many people in a community (e.g. nurses, physicians, school superintendents, pharmacists, and microbiologists) have access to unique information regarding influenza activity that could help predict future influenza activity. However, because this information is disparate, standard research and statistical methods cannot be used to aggregate and analyze it.

One approach to forecasting would be to survey those "experts" about their subjective beliefs regarding future influenza activity. Surveys and polls have been used to predict future events in different contexts, from election outcomes to unemployment rates. To be successful, however, such surveys must be repeated regularly, provide incentives to encourage participants to respond thoughtfully, and have a method for weighting responses. Given the limitations of surveys, we propose to use prediction markets to forecast influenza activity. Prediction markets and surveys have similar advantages. But unlike surveys, prediction markets directly weigh and aggregate information, run continuously (functioning as dynamic, ongoing surveys), provide feedback to respondents (possibly eliciting more informed responses), and use natural incentives to encourage participation.

Futures and Prediction Markets

Futures markets exist for products as different as physical commodities (e.g. grain, lumber) and financial indices (e.g. foreign currencies, interest rates). In recent years, markets in unusual commodities, which we call prediction markets, have been developed to predict future events.[2] Prediction markets are accurate for four reasons: (1) prediction markets aggregate information from all participants, each of whom has different information about the issue in question; (2) the market provides incentives to encourage knowledgeable participants to reveal true information in their trades; (3) the market provides feedback to participants—through market prices, traders learn about the beliefs of others and are motivated to collect more information; and, (4) all trades are anonymous—thus, traders can signal information through the market that they might not want to announce publicly to their peers.

The *Iowa Electronic Market* (IEM) is a well-known prediction market, and the only one in the U.S. that is allowed to trade for cash by the Commodity Futures Trading Commission (CFTC) for educational and research purposes**.** Since 1988, the IEM has run prediction markets for elections in the U.S. and other countries (Figure 1), and for currency prices, stock options, and movie box office receipts. The IEM has developed a prediction record superior to alternative mechanisms such as opinion polls.[3, 4, 5] In his acceptance speech after winning the Nobel Prize in Economics in 2002, Vernon Smith cited the IEM as one of the best demonstrations of how efficiently markets aggregate information about uncertain future events.[6]

How Markets Work

Whenever uncertainty regarding future events exists, markets can be used to aggregate information for forecasting purposes. This process is predicated on both the uncertainty surrounding an event *and* on the existence of disparate information about the event. A simple example, patterned after one provided by Eisenberg and Gale[7], illustrates how markets can aggregate information. In this case, there are two investors and three competing technologies that could produce a new product. Only one technology will ultimately succeed. With no other information, each investor might regard each technology's likelihood of success as 1/3 and thus make offers to buy the rights to all three. Suppose instead that different information is given to each investor. Investor 1 is told that Technology A has an inherent flaw, and Investor 2 is told that Technology B cannot succeed. Each investor still is uncertain about the outcome. However, Investor 1 believes that Technologies B and C each have a 50% chance of success and Investor 2 believes that Technologies A and C each have a 50% chance of success. On the basis of their privileged knowledge, Investor 1 pursues Technologies B and C, and Investor 2 pursues Technologies A and C. Competition will drive the price of Technology C higher, while the lack of competition will drive the price of Technologies A and B to zero. Thus, the prices investors are willing to pay for the rights to the technologies will reveal the successful outcome.

In a less abstract example, Roll[8] demonstrated that contract prices in the orange juice futures market tend to move in expected directions even before the US National Weather Service releases their predictions of freezing weather. Apparently, the money they have at stake motivates traders to gather and use information from diverse sources effectively and efficiently.

Pilot Market for Influenza Activity in the State of Iowa

The usefulness of prediction markets for forecasting influenza activity has yet to be tested. However, we recently conducted a pilot study to determine if healthcare workers with no experience with futures markets could trade in an influenza prediction market. We opened our demonstration market in mid-January of 2004 and invited a limited number of healthcare workers (N = 52) from a variety of backgrounds to participate. Each trader was granted 100 units of a valueless currency ("Flu Dollars") with which to trade. Our software provided a simple web-based interface for "24 and 7" access.

Contracts in this market were based on the Centers for Disease Control's (CDC) color-coded system: RED = Widespread, BLUE = Regional, PURPLE = Local, GREEN = Sporadic, YELLOW = No Activity. A new set of contracts was introduced for each odd-numbered week, beginning with the third week of 2004. The contracts were liquidated when the CDC released its activity report the following week. Figure 2 is an image of the trading screen, showing the set of five contracts for week 3. The contract (color) that denoted the actual outcome for a particular week had a liquidation value of FLU$1.00; all other contracts for that week had a value of FLU$0.00.

Traders bought and sold contracts, based on their beliefs about future influenza activity, at prices in accordance with their expectations. The interaction of all these traders created a set of market prices. The prices for the different contracts reflected the traders' consensus of the probability that future influenza activity will be at a particular level. Because the market was open 24 hours each day, prices were updated with each new trade.

Week 3 contracts traded for one week, January 17 to January 25. Other contract sets were traded over longer periods of time – three weeks for week 5 contracts, five weeks for week 7 contracts, and six weeks for weeks 9 and 11. Prices of trades in those contracts are summarized in Figure 3, and Table 1 illustrates the prediction of influenza activity implied by market prices during each week of trading. As seen in the table, the market correctly anticipated the future level of influenza activity in 19 of 21 cases.

Despite the low level of trading and lack of actual financial incentives, predictions from the market were quite good, with the weekly average prices predicting the outcome over 90% of the time. Furthermore, our market heavily and correctly discounted the possibility that influenza B activity would increase, as is often observed, several weeks after the initial peak of influenza A cases occurred.

Future Directions

We propose a new method for predicting influenza activity using prediction markets. Our pilot project suggests that such markets are feasible. The predictions in the pilot market were accurate and markets conducted by the IEM in other contexts have been very successful. Thus, we believe that similar markets could predict influenza activity accurately several weeks in advance.

In most markets, transactions made for the purpose of maximizing profits are the means by which traders reveal their private information to the marketplace. Consequently, to test our hypothesis, we must address the issue of incentives for trading in the influenza market. One option is for participants to trade in an imaginary currency without investing their own money. When the markets end, we could translate the account balances into small educational grants that traders could use to buy reference books, pay for journal subscriptions, or to defray conference fees.

Prediction markets for infectious diseases have applications beyond influenza. For example, we believe the influenza market could be expanded to help predict which influenza strains should be included in the subsequent year's vaccine. This can be accomplished by a market that aggregates information from influenza experts in the fields of virology, epidemiology, and clinical microbiology. With a trading system and an appropriate panel of traders in place, new prediction markets could be created almost overnight to address emerging infectious diseases even before a specific etiologic agent is identified. For example, information about SARS existed in the spring of 2003, and a prediction market for SARS could have quickly, accurately, and inexpensively aggregated expert opinion about this new infection. Such a prediction market would have been especially useful, given the absence of preexisting SARS surveillance systems.

References

1. Stroup D, Thacker S, Herndon J. Application of Multiple Time Series Analysis to the Estimation of Pneumonia and Influenza Mortality by Age 1962-1983. *Statistics in Medicine* 1988; 7: 1045-1059.

2. Berg JE, Forsythe R, Rietz TA. "The Iowa Electronic Market." In: Paxson D, Wood D, editors. *Blackwell Encyclopedic Dictionary of Finance*. Oxford: Blackwell Press; 1997.

3. Forsythe R, Nelson FD, Neumann GR, Wright J. "Anatomy of an Experimental Political Stock Market." *American Economic Review* 1992; 82:1142-1161.

4. Shaw D, Roberts B. "Campaign Events, the Media and the Prospects of Victory The 1992 and 1996 US Presidential Elections," *British Journal of Political Science* 2000; 259-293.

5. Berg, JE, Forsythe R, Nelson FD, Rietz T. Results from a Dozen Years of Election Futures Markets Research. In: Plott C, Smith V, editors. *The Handbook of Experimental Economic Results*. Forthcoming (2004)

6. Smith V. Constructivist and Ecological Rationality in Economics. *American Economic Review* 2000; 30: 465-508.

7. Eisenberg E, and Gale D. Consensus of Subjective Probabilities: The Pari-Mutuel Method. *The Annals of Mathematical Statistics* 1958; 30: 165-168.

8. Roll R. Orange Juice and Weather. *The American Economic Review* 1984; 74: 861-880.

Table 1. Predictions from the market. Entries in the body of the table indicate the color of the contract and the highest average price for contract during a specific week, which can be interpreted as the predicted probability that the level of influenza activity will be that color. If there was no trading in a contract (color) in a given week, the price of the most recent transaction in that contract is used as its standing prediction rather than the average price. Markets for target weeks 3, 5 and 7 ran for fewer than six weeks.

Weeks in Advance	Target Week 3	5	7	9	11	Fraction of correct predictions
5	n.a.	n.a.	n.a.	green 0.533	yellow 0.797	1 / 2
4	n.a.	n.a.	green 0.525	yellow 0.882	yellow 0.900	3 / 3
3	n.a.	n.a.	green 0.710	yellow 0.844	yellow 0.946	3 / 3
2	n.a.	purple 0.540	green 0.723	yellow 0.720	yellow 0.975	4 / 4
1	n.a.	purple 0.690	green 0.671	yellow 0.807	yellow 0.975	4 / 4
0	blue 0.943	green 0.624	green 0.715	yellow 0.810	yellow 0.975	4 / 5
Actual Level	blue	purple	green	yellow	yellow	total 19 / 21

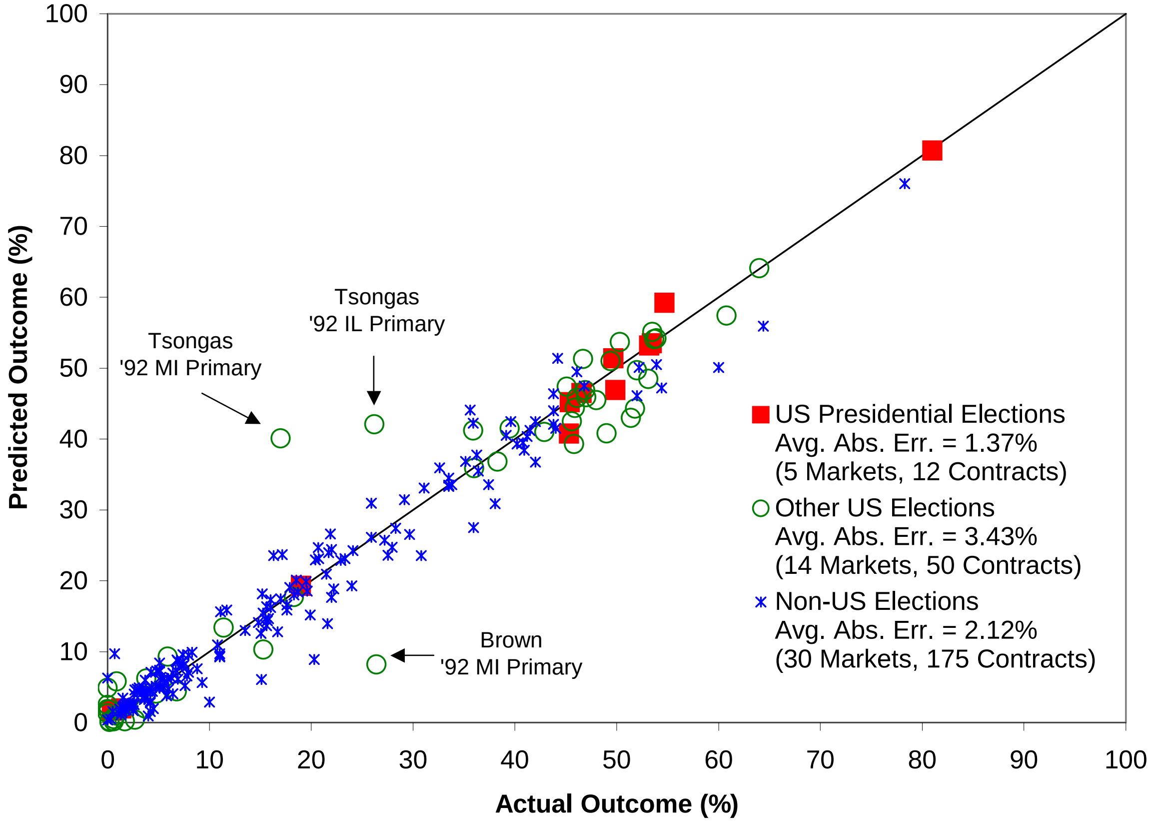

Figure 1: Political futures market predicted versus actual outcomes for vote-share and seat-share markets. Vote-share market predictions are for percentages of votes received by parties or candidates. Seat-share market predictions are for percentages of seats in congress or parliament held by parties. Predictions are based on normalized (to sum to 100%), last-trade prices as of midnight the night before each election. The 45-degree line represents perfect accuracy [5].

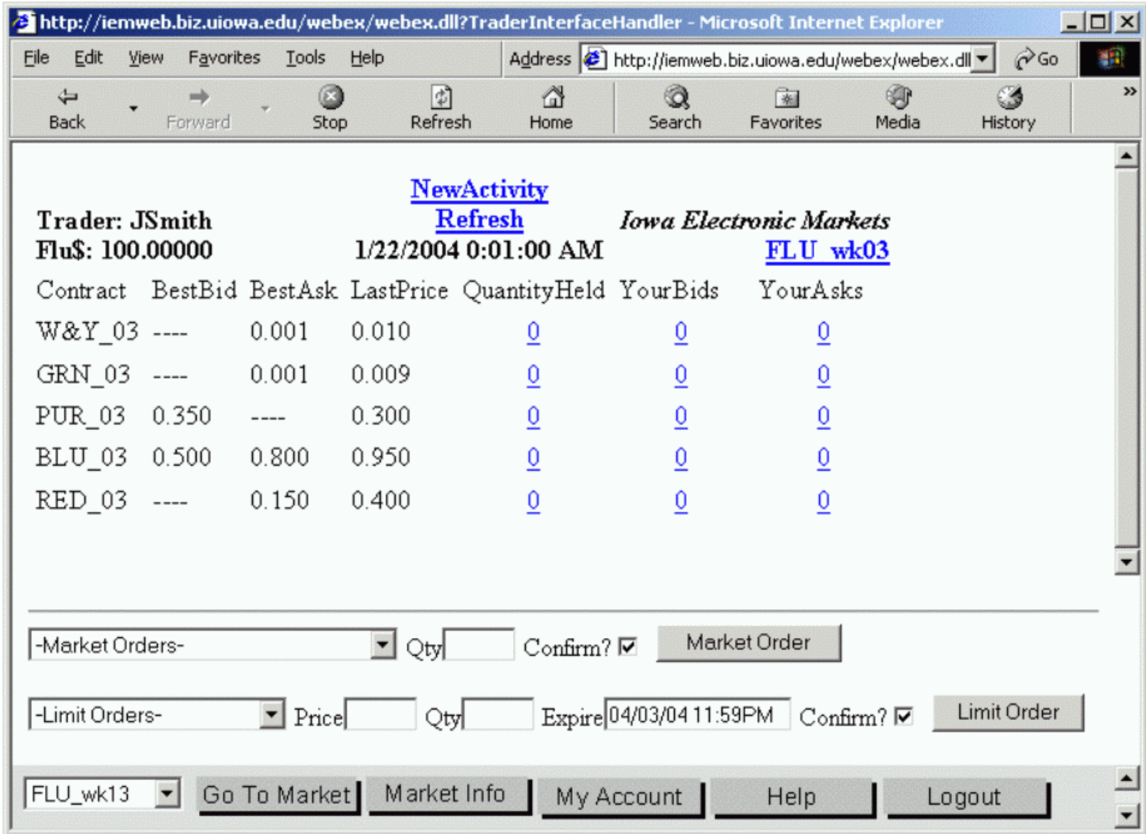

Figure 2: Image of Trading Screen. Participants in the flu market accessed the market and executed transactions from this screen. The image is from early on January 22, 2004, during the first week of trading in the Week 3 contracts.

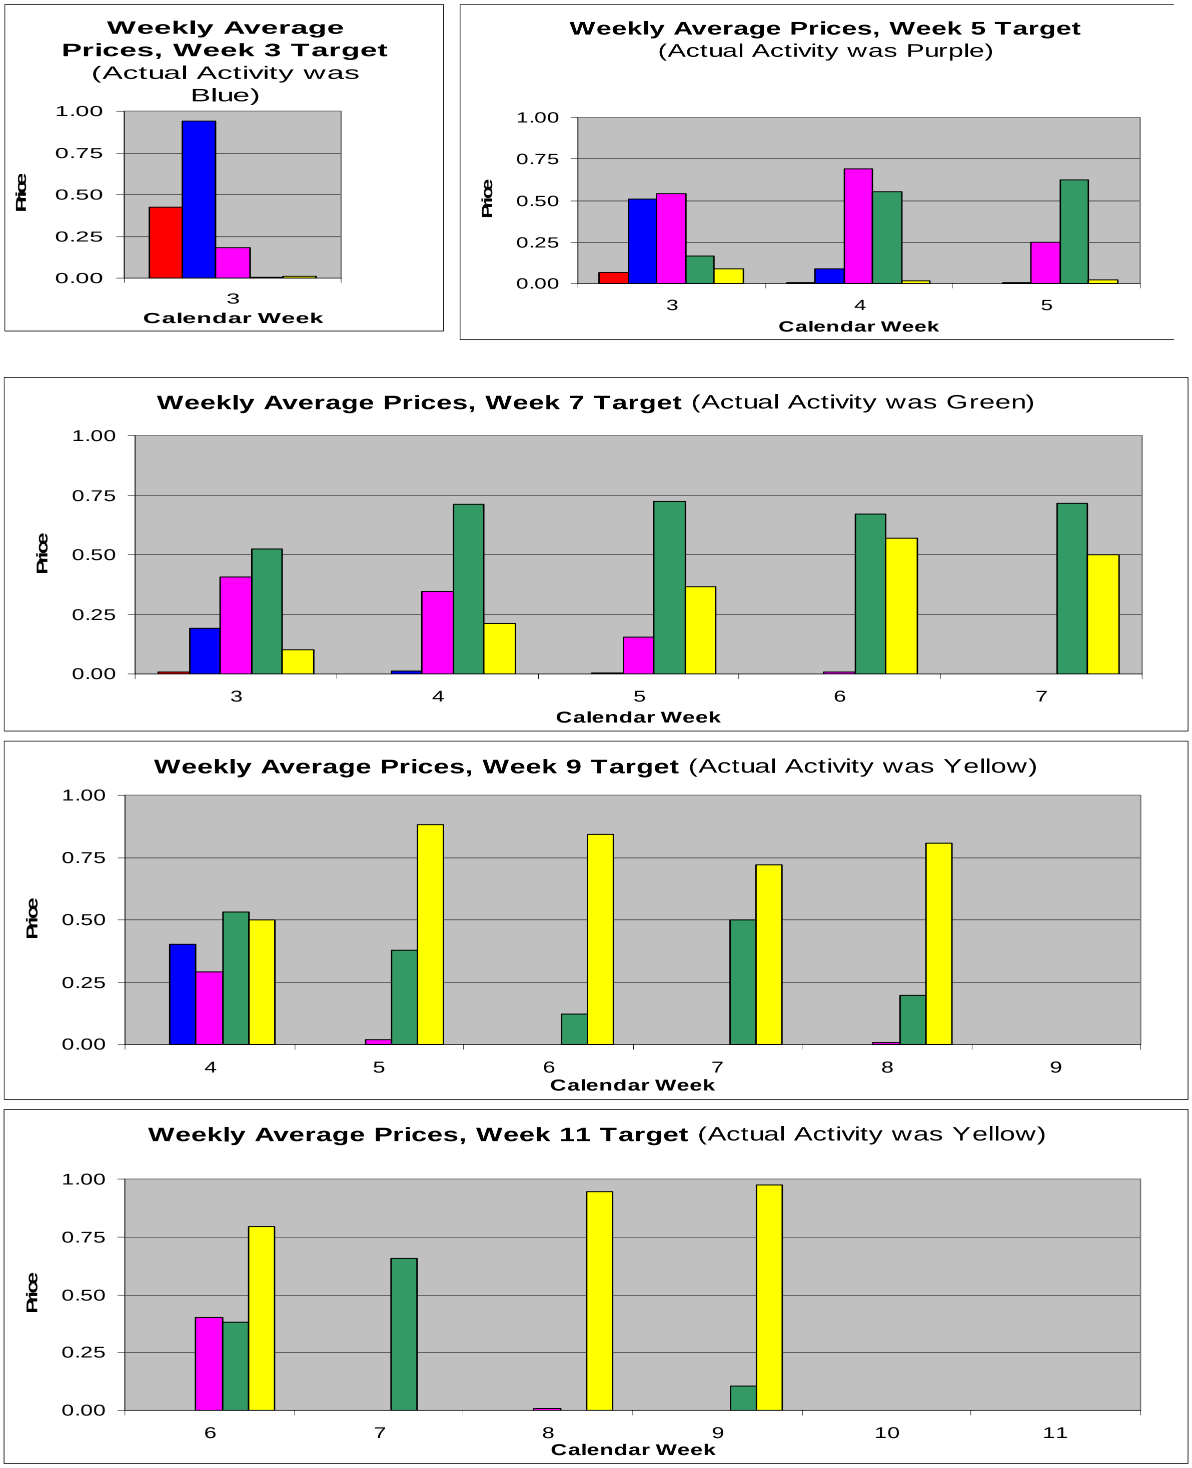

Figure 3: Weekly Average Prices in Demonstration Flu Prediction Market. The height of a bar indicates the average transaction price across one week for one contract. Each panel refers to a different set of contracts; all weeks of trading in one set of contracts are represented in each panel; and the color of a bar indicates the color code of the associated contract. The absence of a bar for a particular contract in any week indicates that no transactions were executed that week in that contract.