

NBER WORKING PAPER SERIES

PREDICTION MARKETS IN THEORY AND PRACTICE

Justin Wolfers
Eric Zitzewitz

Working Paper 12083
<http://www.nber.org/papers/w12083>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2006

This paper was prepared as the entry on “Prediction Markets” for the New Palgrave Dictionary of Economics, 2nd edition, edited by Lawrence E. Blume and Steven N. Durlauf (London: Palgrave MacMillan). The definitive published version of this extract may be found in the complete New Palgrave Dictionary of Economics in print and online, forthcoming. Thanks to Erik Snowberg for very helpful comments. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2006 by Justin Wolfers and Eric Zitzewitz. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Prediction Markets in Theory and Practice
Justin Wolfers and Eric Zitzewitz
NBER Working Paper No. 12083
March 2006
JEL No. C53, D8, G14

ABSTRACT

Prediction Markets, sometimes referred to as "information markets," "idea futures" or "event futures", are markets where participants trade contracts whose payoffs are tied to a future event, thereby yielding prices that can be interpreted as market-aggregated forecasts. This article summarizes the recent literature on prediction markets, highlighting both theoretical contributions that emphasize the possibility that these markets efficiently aggregate dispersed information, and the lessons from empirical applications which show that market-generated forecasts typically outperform most moderately sophisticated benchmarks. Along the way, we highlight areas ripe for future research.

Justin Wolfers
Business and Public Policy Department
Wharton School, University of Pennsylvania
3620 Locust Walk
Room 1456 Steinberg-Deitrich Hall
Philadelphia, PA 19104-6372
and NBER
jwolfers@wharton.upenn.edu

Eric Zitzewitz
Stanford University
Graduate School of Business
518 Memorial Way
Stanford, CA 94305
ericz@stanford.edu

Prediction Markets

Prediction Markets, sometimes referred to as “information markets,” “idea futures” or “event futures”, are markets where participants trade contracts whose payoffs are tied to a future event, thereby yielding prices that can be interpreted as market-aggregated forecasts. For instance, in the Iowa Electronic Market, traders buy and sell contracts that pay \$1 if a given candidate wins the election. If a prediction market is efficient, then the prices of these contracts perfectly aggregate dispersed information about the probability of each candidate being elected. Markets designed specifically around this information aggregation and revelation motive are our focus in this article.

Types of Prediction Markets

The most famous prediction markets are the election forecasting markets run by the University of Iowa (Berg, Forsythe, Nelson and Rietz, 2001). Election forecasting provides a useful way to introduce a variety of different contract types, and Table 1, adapted from Wolfers and Zitzewitz (2004a), shows how different contracts can be designed to reveal various types of forecasts.

<Table 1 here>

The three main types of contracts link payoffs to the occurrence of a specific event (the incumbent wins the election), a continuous variable (the vote share of the incumbent), or to a combination of the two, such as in spread betting. In each case, the relevant contract will reveal the market’s expectation of a specific parameter: a probability, mean, or median, respectively. More complex contract designs can also be used to elicit alternative parameters. For instance, a family of winner-take-all contracts—each linked to different states of nature—can reveal the full probability distribution.

Prediction markets have been used to forecast elections, movie revenues, corporate sales, project completion, economic indicators and Saddam Hussein’s demise. New corporate applications have emerged as firms have looked to markets to predict

research and development outcomes, success of new products, and regulatory outcomes. In the public sector, the Pentagon attempted to use markets designed to predict geopolitical risks, although negative publicity stopped the project (Hanson, 2005). An intriguing attempt to apply prediction markets to forecasting influenza outbreaks is detailed in Nelson, Polgreen and Neumann (2006). Rhode and Strumpf (2004) have detailed the existence of large-scale election betting as far back as the election of Washington.

Prediction market contracts have been traded in a variety of market designs, including continuous double auctions (both with and without market-makers), pari-mutuel pools, bookmaker mediated betting markets, or implemented as market-scoring rules.

Prediction Markets in Theory: Information Aggregation

The claim that prediction markets can efficiently aggregate information is based on the Efficient Market Hypothesis. In certain cases, existing theoretical results regarding efficient capital markets can be applied directly. Grossman (1976) documents a set of sufficient conditions for the equilibrium price of index futures to summarize private information perfectly: In a market where traders with CARA utility functions each receive independent draws from a normal distribution about the true value of the asset, the market price fully summarizes their information.

Manski (2004) noted that much of the analysis of the price of binary options had simply assumed that these revealed a market-based probability estimate, but that appropriate theoretical results were lacking. He illustrates the importance of this issue by way of an example where prediction market prices fail to aggregate information appropriately. In his model all traders are willing to risk exactly \$100. Thus if a contract paying \$1 if an event occurs, is selling for \$0.667, then buyers each purchase 150 contracts, while sellers can afford to sell 300 contracts (at a price of \$0.333). This can only be an equilibrium if there are twice as many buyers as sellers, implying that the market price must fall at the 33rd percentile of the belief distribution, rather than the

mean. The same logic suggests that a prediction market price of π implies that $1-\pi\%$ of the population believes that the event has less than a $\pi\%$ chance of occurring. Clearly the driving force in this example is the assumption that all traders are willing to risk a fixed amount.

Wolfers and Zitzewitz (2005a) provide sufficient conditions under which prediction market prices coincide with average beliefs among traders (and hence aggregate all information in the Grossman setup). They consider individuals with log utility and initial wealth, y , who must choose how many prediction market securities, x , to purchase at a price, π , given that they believe that the probability of winning their bet is q :

$$\text{Max}_{\{x\}} EU_j = q_j \text{Log}[y + x_j(1 - \pi)] + (1 - q_j) \text{Log}[y - x_j\pi]$$

$$\text{yielding : } x_j^* = y \frac{q_j - \pi}{\pi(1 - \pi)}$$

The prediction market is in equilibrium when supply equal demand:

$$\int_{-\infty}^{\pi} y \frac{q - \pi}{\pi(1 - \pi)} f(q) dq = \int_{\pi}^{\infty} y \frac{\pi - q}{\pi(1 - \pi)} f(q) dq$$

If beliefs (q) and wealth (y) are independent, then this implies:

$$\pi = \int_{-\infty}^{\infty} q f(q) dq = \bar{q} .$$

Thus under log utility, the prediction market price equals the mean belief among traders. If wealth is correlated with beliefs, then the prediction market price is equal to a wealth-weighted average belief. This finding is general in the sense that no assumptions are required about the distribution of beliefs, but it is also quite specific, in that it holds only under log utility. Experimenting with a range of alternative utility functions and distributions of beliefs typically yields prediction market prices that diverge from the mean of beliefs by only a small amount.

Both the Manski and Wolfers-Zitzewitz models are silent as to the sources of the different beliefs across traders, which allows them to sidestep the theoretical difficulty posed by Milgrom and Stokey (1982): that under common beliefs, no trade will occur. The logic of the “no trade theorem” is simply that each trader should always be wary that

anyone seeking to trade with them possesses an information advantage, and hence should moderate their beliefs accordingly. Explaining why there is any trade in prediction markets remains an important open theoretical question. Wolfers and Zitzewitz (2006) provide a simple adaptation of the Kyle model in which trade is driven by uninformed outsiders with either hedging- or entertainment-driven demand for the prediction security, or by manipulators attempting to influence market prices.

Another important role of prediction markets is that potential trading profits provide an incentive for *information discovery*. Grossman and Stiglitz (1976) consider the case where information is expensive to garner. They point to the impossibility of prices being fully efficient: if prices fully reflect information, then there is no incentive for any trader to gather that information. Instead, they construct a model in which prices never fully reflect all of the information possessed by informed traders; in equilibrium the inefficiency in pricing is just sufficient to induce a proportion of traders to become informed.

Another key advantage of prediction markets over alternative approaches to information aggregation is that they provide incentives for *truthful revelation* of beliefs. If prediction markets are to be used as inputs into future decisions, this may provide a countervailing incentive to trade dishonestly to manipulate prices. While such manipulation would typically lead the manipulator to lose money, Hanson and Oprea (2005) have shown that these losses increase the rewards for informed trading, which may ultimately increase the accuracy of prediction market prices.

Prediction Markets in Practice

While we are still accumulating evidence on the behavior of prediction markets in different contexts, there are already a few generalizations that can be drawn from existing, albeit piecemeal evidence.

First, market prices tend to respond rapidly to new information. Figure 1 draws an interesting example from Snowberg, Wolfers and Zitzewitz (2006): movements in the price of the Tradesports contract on the re-election of President Bush, around election

day, 2004. Early exit polls suggesting a Kerry victory were leaked at around 3pm, and prices started to move immediately. Indeed, the figure shows that they moved in lockstep with prices on the much larger equity markets. As the count proceeded, it became clear that these early polling numbers were wrong, and the market reversed course sharply. This is only a single anecdote but is representative of the rapid incorporation of new information by prediction markets observed in many domains.

<Figure 1 here>

Second, in most cases, the time series of prices in these markets appears to follow a random walk, and simple betting strategies based on publicly available information appear to yield no profit opportunities. That is, these markets appear to meet the standard definition of weak-form efficiency.

Third, the law of one price appears to (roughly) hold, and the few arbitrage opportunities that arise in these markets are fleeting, and involve only small potential profits.

Fourth, attempts at manipulating these markets typically fail. Camerer (1998) attempted to manipulate pari-mutuel betting on horse races by canceling \$500 bets at the last moment. Rhode and Strumpf (2005) report attempts by specific political campaigns to manipulate the election betting odds on their candidates in the large-scale betting markets operating in the early 20th century. They also analyze an attempt to manipulate the price of a Kerry victory on Tradesports in 2004, as well as their own attempts to manipulate prices on the Iowa Electronic Markets in 2000. None of these attempts at manipulation had a discernible effect on prices, except during a short transition phase.

Finally, prediction markets typically provide quite accurate forecasts and have typically outperformed alternative prediction tools.

Figure 2 shows evidence collected by Gürkaynak and Wolfers (2005) on the relative performance of a prediction market (the “Economic Derivatives” market established by Goldman Sachs and Deutsche Bank), and a survey of economists, in predicting economic outcomes. They show that the market-based forecast encompasses

the information in the survey-based forecasts. Moreover the behavioral anomalies that have been noted in survey-based forecasts are not evident in the market-based forecasts.

<Figure 2 here>

Figure 3 compares the forecasting performance of the Iowa Electronic Markets and the Gallup Poll in predicting the outcomes of Presidential elections in the United States. Over the 13 candidacies from 1988-2004, the average absolute error of the market-based forecasts was 1.6 percentage points, while the corresponding number for the Gallup Poll was 1.9 percentage points. As Berg, Nelson, and Rietz (2003) discuss, the forecasting advantage of markets over the polls is probably even larger over long horizons, as polling numbers tend to be excessively volatile through the electoral cycle. The initial success of these forecasting methods in the United States has led to similar analysis of the forecasting power of prediction markets in Austria, Australia, Canada, Germany, the Netherlands and Taiwan.

<Figure 3 here>

Tests of prediction markets and expert opinions have also been conducted in a range of other domains. The Hollywood Stock Exchange has generated forecasts of box office success and of Oscar winners, that have been more accurate than expert opinions (Pennock, Lawrence, Nielsen and Giles, 2001). Both real and play-money markets have generated more accurate forecasts of the likely winners of NFL football games than all but a handful among 2000 self-professed experts (Servan-Schreiber, Wolfers, Pennock and Galebach, 2004). In the corporate context, the market established by Chen and Plott (2002) within Hewlett-Packard yielded more accurate sales forecasts than the firm's internal experts. Similarly, Ortner (1998) reports that an internal market correctly predicted that the firm would definitely fail to deliver on a software project on time, even when traditional planning tools suggested that the deadline could be met.

Despite this impressive evidence, there still remain a number of documented pathologies in prediction markets. Figure 4 shows evidence from Snowberg and Wolfers

(2005) of the “favorite-longshot bias,” which describes a tendency to over-price low probability events. A similar tendency has been documented in a range of other market contexts, suggesting that some caution is in order in interpreting the prices of low probability events.

<Figure 4 here>

Laboratory experiments also point to the possibility that in some contexts prediction markets will fail to aggregate information as efficiently as alternative procedures. Sunder (1995) provides an excellent review of experimental prediction markets, including experiments showing market designs that fail to aggregate information, and that lead to the appearance of bubbles, false equilibria, or excess volatility.

Economic Analysis of Prediction Market Prices

Prediction markets are a useful way to elicit predictions, but how might they be used? The most direct form of inference involves simply using these predictions directly. For instance, forecasts of election outcomes may be of intrinsic interest.

Some analyses have tried to link the time series of expectations elicited in prediction markets with time series of other variables, so as to isolate a causal influence. For instance, Roberts (1990) analyzes changes in the betting odds posted by Ladbrokes on Ronald Reagan’s re-election and the returns to holding stocks in defense firms, inferring that Reagan led to more robust defense spending. Likewise, Herron, et. al. (1999) and Knight (2005) analyze the correlation of individual stocks and industry indices with movements in the 1992 and 2000 Iowa Electronic Markets U.S. Presidential election markets. Snowberg, Wolfers and Zitzewitz (2006) conduct a similar analysis for the aggregate equity and bond markets at an intraday frequency, using the data shown in Figure 1, to infer partisan impacts of the 2004 election. Slemrod and Greimel (1999) examine the effect on municipal bond prices of changes in the probability of the

Republicans nominating Steve Forbes, whose “flat tax” would have eliminated the tax exemption for municipal bond interest.

Moving beyond ex-post studies of elections, Wolfers and Zitzewitz (2005b) report on an ex-ante analysis of the co-movement of oil and equity prices with a contract tracking the probability of a U.S. attack on Iraq in 2002-3 (Figure 5). The results suggest that a substantial war premium was built into oil prices (and a discount built into equities).

<Figure 5 here>

The contracts we have described thus far have depended on only one outcome. The same principles can be applied to contracts tied to the outcomes of more than one event. These contingent contracts potentially provide insight into the correlation between events. For instance, Wolfers and Zitzewitz (2004b) ran experimental markets on the online betting exchange Tradesports.com in the run-up to the 2004 presidential election. In one example, they ran markets linked to whether George W. Bush would be re-elected, whether Osama bin Laden would be captured prior to the election, and whether *both* events would occur. These markets suggested a 91 percent chance of Bush being re-elected *if* Osama had been found, but a 67 percent unconditional probability. Berg and Reitz (2003) report on contracts whose payoff was linked to 1996 Democratic vote shares conditional on different potential Republican nominees; on the basis of these prices they argue that alternative nominees, such as Colin Powell, would have outperformed Bob Dole.

The potential to apply these markets to determine the consequences of a range of contingencies has led Hanson (1999) to term these “Decision Markets”. Indeed, Hanson (2000) has suggested that such markets could be used to remove technocratic policy implementation issues from the bureaucracy, a suggestion endorsed in Hahn and Tetlock (2006). Moreover, while the previous example involves only one contingency, Hanson (2003) suggests that market scoring rules can allow traders to simultaneously predict many combinations of outcomes. The basic intuition of his proposal is that rather than

betting on each contingency, traders bet that the sum of their errors over all predictions will be lower.

However while contingent markets can be used to estimate the joint probability of choice A and outcome B, care must be taken before inferring that choice A should be made because it will maximize the probability of outcome B. That is, while these markets can highlight the correlation between events, the difficulty of inferring causation remains.

Conclusion

The healthy bibliography below attests to the fact that interest in prediction markets has boomed in recent years. Many questions remain. Theoretical research holds the promise of better understanding the institutional design features that yield optimal information aggregation and efficient pricing. The practical agenda includes developing new ideas about how and when prediction markets can aid decision-making by business and government.

Bibliography

- Berg, Joyce, Robert Forsythe, Forrest Nelson and Thomas Rietz, 2001. "Results from a Dozen years of election Futures Markets Research," in *Handbook of Experimental Economic Results*. Charles Plott and Vernon Smith, eds. Amsterdam: Elsevier, forthcoming.
- Berg, Joyce, Forrest Nelson, and Thomas Rietz, 2006. "Accuracy and Forecast Standard Error in Prediction Markets," *mimeo*, University of Iowa.
- Berg, Joyce and Thomas Rietz, 2003. "Prediction Markets as Decision Support Systems," in *Information System Frontiers*, 5(1), 79-93.
- Berg, Joyce and Thomas Rietz, 2006. "The Iowa Electronic Market: Lessons Learned and Answers Yearned," in *Information Markets: A New Way of Making Decisions in the Public and Private Sectors*, ed. Robert Hahn and Paul Tetlock, AEI-Brookings Joint Center, Washington D.C.
- Camerer, Colin, 1998. "Can Asset Markets be Manipulated? A Field Experiment with Racetrack Betting," *Journal of Political Economy*, 106(3), 457-482.
- Chen, Kay-Yut and Charles Plott, 2002. "Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem," CalTech Social Science Working Paper No. 1131.
- Grossman, Sanford J., 1976. "On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information", *The Journal of Finance*, 31(2), 573-585.
- Grossman, Sanford J. and Joseph E. Stiglitz, 1976. "Information and Competitive Price Systems," *The American Economic Review*, 66(2), 246-253.
- Gürkaynak, Refet and Justin Wolfers, 2005. "Macroeconomic Derivatives: An Initial Analysis of Market-Based Macro Forecasts, Uncertainty, and Risk", in *NBER International Seminar on Macroeconomics*, ed. Christopher Pissarides and Jeffrey Frankel, NBER.
- Hahn, Robert and Paul Tetlock, 2006. "Using Information Markets to Improve Decision Making," *Harvard Journal of Law and Public Policy*, forthcoming.
- Hanson, Robin, 1999. "Decision Markets," *IEEE Intelligent Systems*, 14(3), 16-19.
- Hanson, Robin, 2000. "Shall We Vote on Values, But Bet on Beliefs?" *mimeo*, George Mason University.
- Hanson, Robin, 2003. "Combinatorial Information Market Design," *Information Systems Frontiers*, 5(1) 105-119.
- Hanson, Robin, 2006. "Designing Real Terrorism Futures," *Public Choice*, forthcoming.
- Hanson, Robin and Ryan Oprea, 2005. "Manipulators Increase Information Market Accuracy", *mimeo*, George Mason University.
- Herron, Michael, Donald Cram, James Lavin and Jay Silver, 1999. "Measurement of Political Effects in the United States Economy: A Study of the 1992 Presidential Election," *Economics and Politics*, 11(1) 51-81.
- Knight, Brian, 2005. "Are Policy Platforms Capitalized into Equity Prices? Evidence from the Bush/Gore 2000 Presidential Election," *Journal of Public Economics*, forthcoming.
- Manski, Charles, 2004. "Interpreting the Predictions of Prediction Markets," *NBER Working Paper #10359*.

- Milgrom, Paul and Nancy Stokey, 1982. "Information, Trade and Common Knowledge," *Journal of Economic Theory*, 26(1), 17-27.
- Nelson, Forrest, George Neumann and Philip Polgreen, 2006. "Operating With Doctors: Results from the 2004 and 2005 Influenza Markets", *mimeo*, University of Iowa.
- Ortner, Gerhard, 1998. "Forecasting Markets—An Industrial Application," *mimeo*, Technical University of Vienna.
- Pennock, David, Steve Lawrence, Finn Årup Nielsen and C. Lee Giles, 2001. "Extracting Collective Probabilistic Forecasts from Web Games," in *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 174-183.
- Rhode, Paul W. and Koleman S. Strumpf, 2004. "Historical Presidential Betting Markets," *Journal of Economic Perspectives*, 18(2), 127-142.
- Rhode, Paul W. and Koleman S. Strumpf, 2005. "Manipulation in Political Stock Markets", *mimeo*, University of North Carolina.
- Roberts, Brian, 1990. "Political Institutions, Policy Expectations, and the 1980 Election: A Financial Market Perspective," *American Journal of Political Science*, 34(2), 289-310.
- Servan-Schreiber, Emile, Justin Wolfers, David Pennock and Brian Galebach, 2004. "Prediction Markets: Does Money Matter?" *Electronic Markets*, 14(3), 243-251.
- Slemrod, Joel, and Timothy Greimel (1999) "Did Steve Forbes scare the municipal bond market?" *Journal of Public Economics*, 74, 81-96.
- Snowberg, Erik and Justin Wolfers, 2005. "Explaining the Favorite-Longshot Bias: Is it Risk-Love, or Misperceptions?" *mimeo*, University of Pennsylvania.
- Snowberg, Erik, Justin Wolfers and Eric Zitzewitz, 2006. "Partisan Impacts on the Stockmarket: Evidence from Prediction Markets and Close Elections," *mimeo*, University of Pennsylvania.
- Sunder, Shyam, 1995. "Experimental Asset markets: A Survey," in *Handbook of Experimental Economics*, ed. Al Roth and John Kagel, New Jersey: Princeton University Press.
- Wolfers, Justin and Eric Zitzewitz, 2004a. "Prediction Markets," *Journal of Economic Perspectives*, 18(2), 107-126.
- Wolfers, Justin and Eric Zitzewitz, 2004b. "Experimental Political Betting Markets and the 2004 Election," *The Economists' Voice*, 1(2), article #1.
- Wolfers, Justin and Eric Zitzewitz, 2005a. "Interpreting Prediction Market Prices as Probabilities," *NBER Working Paper No. 10359*.
- Wolfers, Justin and Eric Zitzewitz, 2005b. "Using Markets to Inform Policy: The Case of the Iraq War," *mimeo*, University of Pennsylvania.
- Wolfers, Justin and Eric Zitzewitz, 2006. "Five Open Questions About Prediction Markets," in *Information Markets: A New Way of Making Decisions in the Public and Private Sectors*, ed. Robert Hahn and Paul Tetlock, AEI-Brookings Joint Center, Washington D.C.

Table 1: Contract Types: Estimating Uncertain Quantities or Probabilities

Contract	Details	Example	Reveals market expectation of...	More general application
Binary option	Contract costs \$ p Pays \$1 if and only if event x occurs.	Event x : George Bush wins the popular vote	Probability that event y occurs, $p(x)$.	Defining many events, x_1, x_2, \dots, x_n reveals probability distribution $F(x)$.
Index futures	Contract pays \$ x .	Contract pays \$1 for every percentage point of the popular vote won by George Bush	Mean value of outcome x : $E[x]$	Contract pays some function of x : $\$g(x)$. Reveals specific moments, $E[g(x)]$
Spread betting	Contract costs \$1 Pays \$2 if $x > x^*$ Pays \$0 otherwise. Bid according to the value of x^* .	Contract pays even money if Bush wins more than x^* % of the popular vote.	Median value of outcome, x .	\$1 contract pays $\$(1/q)$ if $x > x^*$. Reveals specific quantile, $F_{1-q}(x)$.

Figure 1: 2004 Election




Figure 2




Figure 3




Figure 4




Figure 5

