Interpreting the Predictions of Prediction Markets

Charles F. Manski
Department of Economics and Institute for Policy Research
Northwestern University

Revised: August 2005

Abstract

Prediction markets are futures markets in which prices are used to predict future events. I derive the equilibrium price when traders are risk-neutral price takers with heterogenous beliefs. I find that price is a particular quantile of the distribution of traders' beliefs. Price does not reveal the mean belief that traders hold but does yield a bound on the mean belief.

Department of Economics, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208, USA,
fax 847-491-8223, email: cfmanski@northwestern.edu

JEL No. D84, G10

This paper is a revision of NBER Working Paper 10359, March 2004. Partial support for this research was provided by National Science Foundation grant SES-0314312. I am grateful for comments from Gadi Barlevy, George Neumann, and Colin Rowat.

1. Introduction

Prediction markets are futures markets in which prices are used to predict future events. Consider an all-or-nothing contract paying a dollar if a specified event occurs and nothing otherwise. Proponents of prediction markets have interpreted the price of such a contract as a "market probability;" that is, a market-generated likelihood that the event will occur. Yet the arguments for this interpretation have been imprecise.

Introducing the Iowa Electronic Markets (IEM) to the research community, Forsythe, Nelson, Neumann, and Wright (1992) sought authority in Hayek (1945), who argued broadly but vaguely that market prices aggregate information. Hayek put it this way (p. 526): "The mere fact that there is one price for a commodity . . . . brings about the solution which . . . . might have been arrived at by one single mind possessing all the information which is in fact dispersed among all the people involved in the process."

In a recent review article, Wolfers and Zitzewitz (2004) wrote (p. 108): "In a truly efficient prediction market, the market price will be the best predictor of the event, and no combination of available polls or other information can be used to improve on the market-generated forecast." The expression "efficient prediction market" refers to the *efficient markets hypothesis* (EMH), which posits that price is a sufficient statistic for all private information held by traders. However, the EMH is only a hypothesis that may hold in some settings; it is not a universal truth.

This report presents the first formal analysis of price determination in prediction markets where traders have heterogeneous beliefs. I consider a simple and illustrative setting, where traders are risk-neutral price takers with finite trading budgets. I report three findings: (a) the equilibrium price of a prediction-market contract is a particular quantile of the budget-weighted distribution of traders' beliefs; (b) price does not reveal the mean belief that traders hold but does yield a bound on the mean belief; (c) the equilibrium price remains the same if traders use price data to revise their beliefs in some ways.

Although these findings are new to the study of prediction markets, finding (a) has previously been obtained in a study of pari-mutuel betting on horse races. Considering races with two horses, Ali (1977) reported the equilibrium condition (1) that I independently derive below. He suggested that this may explain the "favorite-longshot bias," where horses with high equilibrium prices (i.e., favorites) empirically tend to win more often and those with low prices (i.e., longshots) tend to win less often than they should if their prices are interpreted as market probabilities of race outcomes.

2. Price Determination

Consider a prediction market offering all-or-nothing contracts on the occurrence of a binary event; one contract pays a dollar if event m occurs and the other pays a dollar if the contrary event n ≡ (not m) occurs. Let the contract prices be π_m and π_n, where $\pi_m + \pi_n = 1$. It will be shown below that equilibrium prices satisfy this "no-arbitrage" condition.

Suppose that a large population J (formally a continuum) with heterogeneous beliefs participate in this market; without loss of generality, let the unit interval index the members of J. Let each person j have a fixed trading budget of y_j dollars and a subjective probability q_{jm} that event m will occur. Let $P(q_m, y)$ denote the cross-sectional distribution of beliefs and budgets. Let the distribution of beliefs be continuous, with budgets initially being statistically independent of beliefs. Finally, let traders be price takers who maximize the subjective expected value of their contracts.

Under these assumptions, the equilibrium price for contract m in a prediction market such as the IEM solves the equation

$$(1) \quad \pi_m = P(q_m > \pi_m).$$

Thus, most persons have beliefs higher than price when price is above 0.5, and most have beliefs lower than price when price is below 0.5. The equilibrium price is unique and generically equals the $(1-\pi_m)$-quantile of the distribution of beliefs. Equation (1) has a unique solution because $P(q_m > \pi_m)$ is a continuous function of π_m that decreases in value from 1 to 0 as π_m increases from 0 to 1. Price is the $(1-\pi_m)$-quantile of beliefs if the distribution function for beliefs is strictly increasing at $(1-\pi_m)$.

To see why equation (1) holds, consider a person with trading budget y who places subjective probability q_m on occurrence of event m. Price taking and maximization of subjective expected value imply that this person invests his entire budget in contract m if $q_m > \pi_m$ and in contract n if $q_m < \pi_m$. Thus, the person purchases y/π_m units of contract m if $q_m > \pi_m$ and y/π_n units of contract n if $q_m < \pi_m$. All portfolios are optimal if $q_m = \pi_m$, but the assumption that $P(q_m)$ is continuous implies that this equality almost never occurs among the members of J. Hence, aggregating individual behavior across the population, the market demands for contracts m and n are $(1/\pi_m)\cdot E\{y\cdot 1[q_m > \pi_m]\}$ and $(1/\pi_n)\cdot E\{y\cdot 1[q_m < \pi_m]\}$ respectively.

In a market such as the IEM, a person with budget y is given y units of contract m and y units of contract n. Hence, the market supply of each contract is E(y). Equilibrium requires that the demand for each contract equal this common supply. Thus, (π_m, π_n) is an equilibrium price vector if and only if

$$(2) \quad E(y) = (1/\pi_m)\cdot E\{y\cdot 1[q_m > \pi_m]\} = (1/\pi_n)\cdot E\{y\cdot 1[q_m < \pi_m]\}.$$

Statistical independence of y and q_m implies that equation (2) reduces to

$$(3) \quad 1 = (1/\pi_m)\cdot P(q_m > \pi_m) = (1/\pi_n)\cdot P(q_m < \pi_m).$$

The first equality is equation (1). Continuity of $P(q_m)$ implies that $P(q_m < \pi_m) + P(q_m > \pi_m) = 1$, so the two equalities together yield the no-arbitrage condition.

Observe that equation (1) corresponds to equation (4) of Ali (1977), writing on pari-mutuel betting markets. Pari-mutuel markets differ institutionally from prediction markets such as the IEM, but the same equilibrium condition holds if the track returns all money to the bettors. Consider a race with two horses, m and n. The price of a "win" bet on horse m is the fraction of the betting pool that is wagered on m. In the notation of this paper, the aggregate amount bet on m is $E\{y\cdot 1[q_m > \pi_m]\}$ and the betting pool is $E(y)$. The price of a bet on m solves the equation $\pi_m = E\{y\cdot 1[q_m > \pi_m]\}/E(y)$. If y is statistically independent of q_m, this reduces to (1).

Finally, suppose that trading budgets and beliefs are statistically dependent rather than independent. Then equation (2) shows that the equilibrium price depends on the conjunction of traders' budgets and beliefs rather than on their beliefs alone. Although equation (1) does not hold as stated, this equation does hold if one redefines J to be the population of dollars traded rather than the population of traders.

3. Bound on Mean Beliefs

It follows from equation (1) that the mean belief $E(q_m)$ held by traders must lie in the open interval $(\pi_m^2, 2\pi_m - \pi_m^2)$. This interval has midpoint π_m and width $2(\pi_m - \pi_m^2)$. Thus, prices near zero or one are very informative about the mean beliefs of traders, but prices near 0.5 are much less informative.

To obtain the upper bound on the mean belief, observe that the distribution that satisfies equation (1) and has the largest possible mean is the two-point distribution with $P(q_m = \pi_m) = 1 - \pi_m$ and $P(q_m = 1) = \pi_m$. This distribution has mean $2\pi_m - \pi_m^2$. The assumption that $P(q_m)$ is continuous implies that this two-point distribution is not feasible. However, for $0 < \delta < \min(\pi_m, 1 - \pi_m)$, all continuous distributions placing probabilities $1 - \pi_m$ and π_m on the intervals $(\pi_m - \delta, \pi_m)$ and $(1 - \delta, 1)$ satisfy (1). Such distributions have means that are smaller than $2\pi_m - \pi_m^2$ but that approach this value as $\delta \to 0$.

To obtain the lower bound, observe that the distribution that satisfies equation (1) and has the smallest possible mean is the two-point distribution with $P(q_m = 0) = 1 - \pi_m$ and $P(q_m = \pi_m) = \pi_m$. This distribution has mean π_m^2. For $0 < \delta < \min(\pi_m, 1 - \pi_m)$, all continuous distributions placing probabilities $1 - \pi_m$ and π_m on the intervals $(0, \delta)$ and $(\pi_m, \pi_m + \delta)$ satisfy (1). Such distributions have means that are larger than π_m^2 but that approach this value as $\delta \to 0$.

It remains to show that $E(q_m)$ can take any value in the interval $(\pi_m^2, 2\pi_m - \pi_m^2)$. The argument above shows that, for small positive values of λ, there exist continuous distributions that satisfy (1) and that have means $\pi_m^2 + \lambda$ and $2\pi_m - \pi_m^2 - \lambda$. Let $0 < \alpha < 1$ and consider a $(1 - \alpha, \alpha)$ mixture of such distributions. The mixture is continuous, satisfies (1), and has mean $(1 - \alpha)(\pi_m^2 + \lambda) + \alpha(2\pi_m - \pi_m^2 - \lambda)$. Considering all $\alpha \in (0, 1)$ and letting $\lambda \to 0$ proves the result.

4. Using Price Data to Revise Beliefs

I have thus far assumed that persons hold fixed beliefs about the likelihood of event m. A prominent theme of information economics, taken to the limit in the efficient markets hypothesis, has been that prices may reveal private information held by market participants. Hence, much research in information economics conjectures that traders use price data to revise their beliefs and studies the implications for market outcomes.

It is not known whether and how traders in prediction markets use price data to revise their beliefs. However, I can show that the equilibrium price derived above is unchanged if traders do not revise their beliefs too much, in a sense made specific below.

Let each trader j hold prior subjective probability q_{jm} when the market opens and revise this belief to $q_{jm}(\pi_m)$ after observing trades take place at price π_m. Assume that prior and posterior beliefs have the same ordinal relationship to price; thus, $\text{sgn}[q_{jm}(\pi_m) - \pi_m] = \text{sgn}[q_{jm} - \pi_m]$. This assumption does not mandate a particular rule for revision of beliefs, but it is consistent with many possible rules. For example, the assumption holds if $q_{jm}(\pi_m)$ is a weighted average of π_m and q_{jm} of the form $q_{jm}(\pi_m) = \theta_j q_{jm} + (1 - \theta_j)\pi_m$, where $\theta_j \in (0, 1]$.

To see that the equilibrium price is unchanged, replace the equilibrium condition (2) that held using prior beliefs with the analogous condition using posterior beliefs; that is,

$$(4) \quad E(y) = (1/\pi_m)\cdot E\{y\cdot 1[q_m(\pi_m) > \pi_m]\} = (1/\pi_n)\cdot E\{y\cdot 1[q_m(\pi_m) < \pi_m]\}.$$

If prior and posterior beliefs have the same ordinal relation to price, then $1[q_m(\pi_m) > \pi_m] = 1[q_m > \pi_m]$ and $1[q_m(\pi_m) < \pi_m] = 1[q_m < \pi_m]$. Hence, the equilibrium price using posterior beliefs is the same as the one using prior beliefs.

5. Postscript

The analysis of this paper demonstrates the danger of loosely interpreting prices in prediction markets as "market probabilities" that aggregate the information held by traders. I suggest that my specific findings about the relationship between price and the distribution of traders' beliefs should be applied cautiously to actual prediction markets, where traders may or may not behave in the manner assumed here.

Circulation of the first draft of this paper in early 2004 stimulated other researchers to study price determination when the assumption of risk neutrality in my analysis is replaced by various assumptions of risk aversion. Working independently of one another, Gjerstad (2005) and Wolfers and Zitzewitz (2005) show that equilibrium price is the mean belief $E(q_m)$ if all traders maximize the subjective expected value of the logarithm of their contracts. Other forms of risk aversion imply other relationships between price and the distribution of trader beliefs. These findings, in conjunction with my own, make plain that interpretation of prices in actual prediction markets requires knowledge of traders' risk preferences.

References

Ali, M. (1977), "Probability and Utility Estimates for Racetrack Bettors," *Journal of Political Economy*, 85, 803-815.

Forsythe, R., F. Nelson, G. Neumann, and J. Wright (1992), "Anatomy of an Experimental Political Stock Market," *American Economic Review*, 82, 1142-1161.

Gjerstad, S. (2005), "Risk Aversion, Beliefs, and Prediction Market Equilibrium, " Economic Science Laboratory, University of Arizona.

Hayek, F. (1945), "The Use of Knowledge in Society," *American Economic Review*, 35, 519-530.

Wolfers, J. and E. Zitzewitz (2004), "Prediction Markets," *Journal of Economic Perspectives*, 18, 107-126.

Wolfers, J. and E. Zitzewitz (2005), "Interpreting Prediction Market Prices as Probabilities," Wharton School, University of Pennsylvania.