FROM UNCERTAINTY TO RISK?
PROBLEMS WITH USING PRICES IN PREDICTION MARKETS TO IMPROVE DECISION-MAKING

Ryan Bubb

January 2005

I. INTRODUCTION

With four days to go, there is only a 2% chance that Osama bin Laden will be captured in January 2005. And there is a 13% probability that a Palestinian state will be established by year's end. Whether the Republicans or the Democrats will win the White House in 2008 is basically a toss-up (the Republicans have a slight lead at 50.9% to 48.4%). Or at least so says the prediction market at TradeSports.[1] The price of the last trade on TradeSports for a contract that will pay $10 on Feb. 1 if Osama is captured during the month of January 2005 (and nothing otherwise) was $0.20. Proponents of prediction markets argue that this market price can be interpreted as a meaningful forecast of the likelihood that Osama will indeed be captured in the next several days. A probability of 2% makes this exchange at 20 cents a fair bet. The intuition that the market prices of these contingent claims represent meaningful predictions about future events is based on the theory of efficient capital markets. Markets provide incentives for a large group of individuals to collect information about the event of interest and to reveal their information by placing bets. The price formation process aggregates traders' information, and, if the market is efficient, the resulting price fully reflects all of the information available to traders in the market.

Prediction markets like TradeSports have attracted increasing attention as a potential tool for improving decision-making under uncertainty. Proponents of prediction markets argue that price-based forecasts are more accurate and objective than conventional forecasting approaches and advocate their use by public officials, corporate executives, and the general public to make predictions regarding a broad domain of events. [2]

---

[1] These forecasts reflect market prices on TradeSports at 10:17 AM EST on January 28, 2005.

[2] See, e.g., Abramowicz (2004) ("[T]he objectivity of well-functioning information markets markes them a promising potential administrative decisionmaking tool."); Hahn (2004) ("[I]nformation markets could promote greater transparency in governmental decision making, provide more accurate estimates of the … impacts of different policies, provide a better understanding of uncertainties, … and aid in the design of policies."); Hanson (1999) ("By its nature, a betting-market estimate is decentralized, concise, timely, precise, self-consistent, and responds quickly to new information."); and Berg and Rietz (2002) ("[T]hrough the price formation process, the markets aggreage information across traders, solving what

The purpose of this paper is to point out some basic theoretical problems with using prices in prediction markets to improve decision-making. First, consider events for which an objectivist interpretation of probability is sensible. Prediction markets may be able to solve the problem of aggregating information about such events and to produce a price that reflects the probability estimate that would be made by an agent with access to all of the information available to traders. But whether these markets will produce such a price – that is, whether these markets are efficient – is an empirical question. The potential problem with prediction markets from this objectivist viewpoint is *technical*: given the distribution of information among traders and the cost of obtaining information, will the market price converge to the full information price?

A more fundamental problem with using prediction markets arises in situations when an objectivist interpretation of probability is problematic, namely for unique one-off events for which uncertainty arises not out of a mechanical randomizing device but from more complex and poorly understood systems. For example, it is hard to make sense of an objective probability for any of the events described in the opening paragraph. The uncertainty in these situations is unmeasurable or Knightian uncertainty, and an appeal to efficient markets will not resolve the fundamental *epistemological* problem of resolving Knightian uncertainty into risk.

II. DESCRIPTION OF PREDICTION MARKETS

a. The Mechanics of Prediction Markets

Prediction markets are markets that allow participants to trade contracts that require a payout from the organizer of the market to the holder of the contract as a function of the outcome of some future event. Several different types of contracts are traded on these markets. The simplest are winner-take-all contracts that pay out a set amount if a particular event comes to pass and nothing otherwise. For example, a contract might pay $1 if, say, George W. Bush is reelected, and nothing if not. In contrast, index contracts offer payouts that vary continuously according to the realization of some continuous underlying variable. For example, vote-share contracts pay out according to the fraction of the vote a candidate wins such that if, say, a candidate wins 46% of the vote, that candidate's vote-share contract pays $0.46.[3]

Typically a set of securities are traded that partition the entire state space for some event of interest. Together, the payouts of such a set of partitioning contracts are usually set such that a holder of one of each contract in the set will receive some amount with

---

would otherwise be complex (at best) aggregation problems. … [T]hese forecasts can outperform existing alternatives.").

[3] Interestingly, winner-take-all contracts are typically more popular with traders than index contracts. When both winner-take-all contracts and vote-share contracts are offered for an election, the trading volume of winner-take-all contracts exceeds that of vote-share contracts. See Berg, Nelson and Rietz (2003, p. 11). This suggests that participants in these markets are rather risk-loving since winner-take-all contracts are generally more risky than index contracts. It also highlights the entertainment value of these markets; presumably it is more fun to bet on your candidate winning rather than to bet on the more arcane vote shares of candidates.

certainty (typically $1, $10 or $100). For example, winner-take-all contracts are often traded for each major candidate running in an election, as well as an additional catch-all contract for all other candidates. Each winner-take-all contract pays out the same amount if the contract's candidate wins, thus ensuring that a holder of one of each contract in the set will receive that amount with certainty.[4] This feature provides a simple way to seed the market; the market-maker sell bundles of securities that pay some amount with certainty.

Participants trade contracts in prediction markets just as they do securities in conventional securities markets. Holders of contracts announce ask prices, prospective buyers announce bid prices, and when a bid price meets or exceeds an ask price, a trade is executed.[5]

The Iowa Electronic Market (IEM) is the only legal prediction market in the U.S.[6] It was established in 1988 by the College of Business at the University of Iowa and focuses primarily on U.S. elections and economic forecasts. Dublin-based TradeSports is a popular offshore site offering contracts on diverse topics including sports, politics, and current events.[7]

b. Proposed Applications of Prediction Markets

The primary reason people believe that these markets are socially useful, and the reason they are called "prediction" markets, is that many believe that the market prices at which contracts are traded can be used to infer the probability of future events. Proposed applications of the supposed predictive power of these markets are numerous.

i. Politics

One potential use of prediction markets is to supplement or even supplant polling as a source of information on political issues and candidates. The price of a winner-take-all contract for a political candidate is believed to reflect the probability that the candidate will be elected. The prices of vote-share markets are interpreted as predictions of the ultimate vote-share a candidate will receive (thus, a Kerry vote-share contract trading at $0.49 is interpreted as predicting that Kerry will receive 49% of the vote). During the 2004 election cycle, price-watching almost eclipsed poll-watching as many regularly monitored the prices of Bush and Kerry contracts on IEM and TradeSports in order to

---

[4] Vote-share markets are set up similarly such that each candidate's vote-share contract (and a catch-all "other" contract) pays out the product of the fraction of votes the candidate wins and some sum that is uniform across all candidates' contracts. A holder of one of each candidate's vote-share contracts then receives that sum with certainty.

[5] The precise details of the trading mechanism vary from market to market. For a detailed description of how trading works on IEM, see the IEM Trader's Manual at http://www.biz.uiowa.edu/iem/trmanual/.

[6] IEM is online at http://www.biz.uiowa.edu/iem/.

[7] TradeSports is online at http://www.tradesports.com.

glean insight into the fortunes of the candidates.[8] Berg, Nelson and Rietz (2003) compare the prices of vote-share contracts on IEM to polls for U.S. presidential elections in 1988, 1992, 1996 and 2000 and conclude that IEM prices significantly outperformed polls in predicting final vote-shares.

Predictions regarding political events may be useful to both candidates and voters. Berg and Rietz (2003) suggest that campaign managers gauge the effectiveness of particular campaign strategies by measuring the change in prices in response to new campaign approaches, such as the introduction of a new ad or stump speech. They also argue that voters should assess the electability of primary candidates by noting the prices of contracts that pay out according to the success of a primary candidate in the general election. They point to the relatively low values of Bob Dole general election winner-take-all contracts during the Republican primary race in 1996 as evidence that Dole was a weak candidate against Bill Clinton.

ii. Policy

An application with potentially significant social welfare ramifications is in government decision-making. Government decision-makers must make a vast number of decisions under uncertainty, and they typically utilize a number of techniques to reduce uncertainty and improve decision-making under uncertainty, such as employing intelligence services to gather and analyze information, commissioning studies by experts, and requiring administrative agencies to follow rules in their decision-making process that allow interested parties an opportunity to furnish relevant information. Hanson (2004) argues that decision markets can be used to improve estimates of the consequences of adopting proposed policies. Consider a policy P that a public official is considering adopting. Suppose a critical criterion for the decision whether to adopt P is the policy's expected effect on some outcome variable that can be measured by a statistic O (normalized to lie in [0, 1]). Thus, the public official would like to estimate the "treatment effect": E[O | P] – E[O | not P]. Hanson proposes that this effect be measured using the prices of contingent claims traded in a prediction market. In particular, consider the following sets of securities, each of which pays $1 with certainty:

1. [Pays $1 if P] and [Pays $1 if not P]
2. [Pays $O if P], [Pays $(1-O) if P], [Pays $O if not P], and [Pays $(1-O) if not P]

Each set of securities could be sold as a bundle by the market-maker for $1 to seed the market. Then, a market could be created that would allows traders to exchange the security [Pays $O if P] for a fraction of the security [Pays $1 if P]. The fraction is the market price and represents the market estimate for E[O | P]. Similarly, another market could be established that would allow traders to exchange [Pays $O if not P] for a fraction of [Pays $1 if not P], and the fraction at which contracts trade is the market

[8] See, e.g., Pamela Gaynor, "Kerry Leads Bush in Online Trading," Pittsburgh Post-Gazette, Nov. 3, 2004, p. B13; and Daniel Gross, "Economic View: Polls Say Kerry. Futures Say Bush." New York Times, Aug. 8, 2004.

estimate of E[O | not P]. Such an approach, Hanson argues, would effectively aggregate information relevant to the estimate and provide an unbiased estimate that is likely to be more accurate than conventional approaches.

A recent attempt to employ prediction markets by the U.S. government ended in a public relations debacle. The Defense Advanced Research Projects Agency (DARPA) funded a program to develop prediction markets that would provide forecasts related to various geopolitical issues, such as the likelihood of a bioweapons attack against Israel, and the stability of Saudi Arabia conditional on a pull-out of U.S. troops from the country. National security-related decision-making has been under particular scrutiny in the U.S. following the failure of the U.S. intelligence community to anticipate and prevent the attacks of September 11, 2001, and the failure of U.S. intelligence related to WMD in Iraq.[9] Despite this perceived need for intelligence reform, prediction markets were judged beyond the pale, and the program was shutdown in July 2003 amid outcries from Congress that the program allowed people to bet on tragedy.[10]

iii. Business

Prediction markets could also potentially be used by managers to improve business decision-making under uncertainty. Chen and Plott (2002) report on the use of a prediction market by Hewlett-Packard Corporation to make sales projections. Another potential business use is by securities analysts to make predictions about events relevant to the valuation of securities, such as whether a particular drug will be approved by the FDA.[11]

III. THEORETICAL PROBLEMS USING PREDICTION MARKETS TO PREDICT

Wolfers and Zitzewitz (2004) summarize the rationale for prediction markets as follows: "[T]he power of prediction markets derives from the fact that they provide incentives for *truthful revelation*, they provide incentives for research and *information discovery*, and the market provides an algorithm for *aggregating opinions*." However, the usefulness of prediction markets as a tool to improve forecasts (and thus decision-making) in each of these domains depends critically on whether the prices of contingent claims traded in prediction markets can indeed be interpreted as meaningful predictions. The intuition that the prices on prediction markets can be interpreted as predictions is rooted in the theory of efficient capital markets. Generally, a market is deemed

---

[9] Interestingly, contracts that paid out if WMD were found in Iraq by various dates were traded on TradeSports in 2003 and the probabilities implied by the market prices of the contracts reached as high as 80% before plummeting. See Wolfers and Zitzewitz (2004), who conclude that "the public information on the probability of weapons of mass destruction in Iraq appears to have been of dubious quality…"

[10] Senator Hillary Rodham Clinton described the program as "a futures market in death." See Robert Schlesinger, "Plan Halted for a Futures Market on Terror," The Boston Globe, July 30, 2003, p. A1.

[11] Interestingly, such an application would involve using a market to make a prediction that can be used in making trading decisions in yet another market. The prices in the stock market would presumably already reflect the market prediction concerning the underlying event of interest, but a prediction market would potentially enable a securities analyst to disaggregate a particular event from other events that affect the value of a security.

"efficient" if the market prices at which securities are traded "fully reflect" all information available to traders in the market. (Fama (1970)). Thus, the thinking goes, if a prediction market is efficient, the prices in the market reflect probabilities calculated using all information available to traders.

The advocates of prediction markets implicitly adopt an objectivist view of probability. The market estimate is deemed to be objective in that it reflects the probability that any agent with access to all of the information available to traders in the market would make. Two potential problems arise with this perspective. First, even assuming that an objectivist framework is appropriate for the event underlying a contingent claim, the market may fail to converge to the objective probability-determined fair bet price. This is essentially a technical problem concerning the price formation process. The second problem is more fundamental: if the events one is interested in forecasting are not susceptible to an objectivist interpretation of probability, the best one can hope for from a prediction market is to provide information concerning the distribution of subjective probabilities among traders. The market cannot resolve this fundamental epistemological problem by transforming subjective beliefs into a price.

a. Prediction Markets Under Objective Risk – The Efficient Prediction Market Hypothesis

First, let us adopt an objectivist interpretation of probability and consider potential problems using prices in prediction markets as objective probability estimates.[12] To correctly compute the objective probability distribution over a state space, one generally needs information about the relevant events. The role of prediction markets in such situations is to provide incentives to collect that information, and to provide an algorithm for combining the beliefs that traders form based on their individual information. Most advocates of prediction markets assume that prediction markets will produce prices that correspond to the probability assessment of an agent with access to all of the information available to traders in the market. Fama (1970) provides three sufficient conditions for capital market efficiency that can be applied to prediction markets: (1) there are no transaction costs for trading; (2) all available relevant information is costlessly available to all traders; and (3) all traders agree on the implications of current information for the distributions of future prices of each security.

Consider a prediction market where these conditions are (to a first approximation) met. As a simple concrete example, consider whether the Philadelphia Eagles or the New England Patriots will win the opening coin toss in the Super Bowl on February 6, 2005. There is virtually unanimous agreement that this will be determined by the toss of a fair coin.[13] Suppose a prediction market is established to predict which team will win the

---

[12] An objectivist interpretation is appropriate for events which can easily be repeated under similar conditions such that a relative frequency interpretation of probability can unproblematically be applied. This is certainly not an accurate characterization of most events for which prediction markets have been proposed, and subjectivist objections are considered in Section III.b. below.

[13] Conspiracy theories are imaginable, but let's exclude these potential sources of uncertainty and consider only the randomness in a fair coin toss.

toss. Fama's sufficient conditions (2) and (3) would likely hold reasonably well in such a market, and with Internet-based trading technology, trading would be cheap. It seems likely that any trades in such a market would be at a price that implies a 50/50 probability distribution. Note, however, that a prediction market would be of little value in such a situation as all would agree on the probability estimate. Furthermore, there would likely not be much trading in such a market given the lack of disagreement among traders regarding the probability distribution.

More interesting and potentially useful applications of prediction markets are in cases where Fama's conditions (2) and (3) do not hold – where information is costly to obtain, asymmetrically distributed, and has ambiguous implications. As Fama (1970) emphasizes, these conditions are only sufficient, not necessary, for efficient pricing. When conditions (2) and (3) do not hold, prediction markets may be able to incentivize information collection and to aggregate information about such events to produce an efficient price that implies the probability estimate that would be made by an agent with access to all of the information available to traders. Whether prediction markets are able to achieve this is ultimately an empirical question.

The voluminous literature on the efficient capital market hypothesis is directly relevant to this "efficient prediction market hypothesis," and analogous forms of the strong, semi-strong, and weak forms of the efficient capital market hypothesis can be formulated and tested for prediction markets. The strong form of the efficient market hypothesis states that no trader possesses information that allows him to receive higher expected returns than other traders. The semi-strong form maintains that there is no trading strategy based on public information that produces higher than market expected returns. The weak form holds that prior price histories of securities do not contain information about future returns. A comprehensive review of the literature on the efficient market hypothesis is beyond the scope of this paper[14] – my aim is simply to point out that many advocates of prediction markets simply assume (sometimes only implicitly) that this hypothesis is true, in the strong form, with respect to prediction markets without offering sufficient justification or evidence. A few further observations are also possible.

First, Wolfers and Zitzewitz (2004) point out a particular information structure in which prediction markets are unlikely to be successful – events for which a small group of traders is highly informed but public information is poor and the general public is aware of the existence of the informed group. In such a situation, traders with access only to public information would be unwilling to trade, knowing that insiders would likely win the implicit bets, and the market would be unlikely to form. In support of this argument, they note that TradeSports contracts on the next Supreme Court retirement and the next Pope have generated little volume. They argue that prediction markets will work better in cases where there is ambiguous public information and not accurate private information since there will be more information on whose interpretation traders can disagree.

[14] For a recent, pro-efficient market hypothesis review of the empirical literature, see Malkiel (2003).

Second, there is tension between the efficiency of a prediction market and the incentive to acquire costly information. Grossman and Stiglitz (1980) make this observation with respect to capital markets, arguing that an efficient market (in the strong form sense) when information is costly to obtain results in a paradox: efficient pricing implies that there is no return to obtaining costly information, yet efficient pricing requires (at least some) traders to acquire that information. The same argument applies to prediction markets. The implication is that, at a minimum, there must be some adjustment period before the prices in a prediction market reflect new information obtained by traders.

b. Prediction Markets Under Knightian Uncertainty

A more fundamental problem with using prediction markets arises when the objectivist interpretation of probability is challenged. Consider the identity of the Democratic nominee for President of the U.S. in 2008. Current prices on TradeSports put Hillary Clinton way out in front at 35% with Indiana Senator Evan Bayh a distant second at 10% and seventeen other potential candidates in the running at below 10%. Can the prices for these contingent claims really be interpreted as objective probabilities? I submit that the identity of the eventual Democratic nominee in 2008 is uncertain in a Knightian sense – it is unknown and no objective probability distribution over possible outcomes can be determined. The oft-quoted formulation of Keynes (1937) is useful:

> By "uncertain" knowledge, let me explain, I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty... Or, again, the expectation of life is only slightly uncertain. Even the weather is only moderately uncertain. The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence...About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know.

Similarly, we simply do not know who will be nominated by the Democratic party in 2008. Prediction markets provide no resolution to this fundamental epistemological problem.

That many, if not most, questions to which prediction markets are applied are subject to this more fundamental uncertainty is not addressed in the literature on prediction markets. For example, Abramowicz (2004) asserts that, "[w]hile information markets cannot allow us to select a single future, [assuming that information markets aggregate information as well as alternative approaches] they can permit us to select from probability distributions of possible futures." Most proponents of prediction markets similarly seem to believe that by filtering subjective beliefs through a market, the individually subjective becomes objective, regardless of the nature of the underlying event in question.

For decision-making under Knightian uncertainty, one must acknowledge the subjective nature of probability assessments. Prediction markets can still potentially perform a useful function – they provide a mapping from distributions of subjective

probability beliefs to prices and thus may provide information about underlying subjective beliefs. The extent to which prediction markets reveal information about underlying beliefs is similar to the extent to which prediction markets are efficient – it depends on the details of the price formation process. Manski (2004) uses a model of risk neutral price takers with heterogeneous beliefs and fixed budgets to investigate the mapping of prediction markets from subjective belief distributions to prices. In his model, the equilibrium price only pins down the mean belief of the population of traders to within some (potentially large) interval, and reveals nothing about the higher moments of the distribution. His results remain unchanged even when he incorporates a simple belief revision process by which traders update beliefs upon learning the market price.[15] Further modeling and empirical investigation is needed to better understand the price formation process of prediction markets.

But even if the price formation process worked such that you could, perhaps through a series of different contracts, pin down the entire distribution of subjective probabilities among traders, this does not then give you any objective or "accurate" probability for the event in question. Either the event will come to pass, or it will not, and with one-off, unique events subject to Knightian uncertainty, there is no meaningful way to speak of an objective probability for the event, much less an aggregation rule that translates a distribution of subjective probability assessments into an objective probability. In short, prediction markets offer no resolution to the basic epistemological problem of transforming uncertainty into objective risk.

IV. CONCLUSION

Prediction markets fall into a general family of social institutions that incentivize the discovery and revelation of information. Some of its cousins avoid the technical and epistemological problems of prediction markets. For example, “information prizes” are a simple way to incentivize the collection and revelation of information; people who provide information that leads to an arrest and conviction for a crime commonly receive monetary rewards and patents provide strong (if crude) incentives to develop new knowledge.

The relationship between market prices of contingent claims traded in prediction markets and the likelihood that the underlying event will occur has generally been undertheorized. That the existing literature on prediction markets does not thoroughly consider the problems raised in this paper is perhaps a tribute to the lasting hegemony of the efficient market hypothesis, despite recent assaults by behavioral economists.[16] Theoretical and empirical investigation of the “efficient prediction market hypothesis” is

[15] Manki assumes that prior and posterior beliefs bear the same ordinal relation to price, and argues that this is a reasonable way to model rational expectations.

[16] Malkiel (2004) writes, “Many of us economists who believe in efficiency do so because we view markets as amazingly successful devices for reflecting new information rapidly and, for the most part, accurately.” This belief in amazingly efficient markets has clearly become part of the deposit of faith among economists. McCloskey’s (1983) observation that much argument in economics is based not on its formal modernist methodological underpinnings but rather on a looser informal methodology is reflected in the faith in efficient markets evinced by the literature on prediction markets.

in order, as well as consideration of the epistemological problems posed by many of the events for which prediction markets are hoped to reduce uncertainty.

REFERENCES

Abramowicz, Michael, "Information Markets, Administrative Decisionmaking and Predictive Cost-Benefit Analysis," *University of Chicago Law Review,* 71 (2004) pp. 933-1020.

Berg, Joyce A. and Thomas A. Rietz, "Prediction Markets as Decision Support Systems," *Information System Frontiers*, 5 (2003), pp. 79-93.

Berg, Joyce A., Forrest Nelson, and Thomas Rietz, "Accuracy and Forecast Standard Error of Prediction Markets," Working Paper (July 2003).

Fama, Eugene, "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance*, 25 (1970), 383-423.

Grossman, Sanford and Joseph Stiglitz, "On the Impossibility of Informationally Efficient Markets," *American Economic Review* 70 (1980), 393-408.

Hahn, Robert W., "Using Information Markets to Improve Policy," AEI-Brookings Jont Center Working Paper 04-18 (November 2004).

Hanson, Robin D., "Decision Markets," *IEEE Intelligent Systems Vol.* 14 (May/June 1999), pp. 16-17.

Hanson, Robin D. "Impolite Innovation: The Technology and Politics of 'Terrorism Futures' and Other Decision Markets," unpublished working paper (Nov. 16 2004).

Keynes, J.M., "The General Theory of Employment," *The Quarterly Journal of Economics* 51 (Feb. 1937), pp. 209-223.

Malkiel, Burton G., "The Efficient Market Hypothesis and Its Critics," *Journal of Economic Perspectives* 17 (2003), pp. 59-82.

Manski, Charles F., "Interpreting the Predictions of Prediction Markets," Working Paper (Feb. 2004).

McCloskey, Donald N., "The Rhetoric of Economics," *Journal of Economic Literature* 21 (1983), pp. 483-517.

Plott, Charles R. and Kay-Yut Chen, "Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem," Working Paper (March 2002).

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