NOTE

INFORMATION MARKETS: USING MARKET PREDICTIONS TO MAKE ADMINISTRATIVE DECISIONS

Matthew Einbinder[*]

INTRODUCTION

SOMETHING amazing is happening in Iowa. A few hundred traders have predicted the outcome of the last five presidential elections more accurately than the polls.[1] Trading securities that pay off based on election results, these "well-educated, high income, and young" traders have masterfully predicted elections:[2] in 596 comparisons between market predictions and polls, the market outperformed the polls 451 times.[3] In the week leading up to each of the last four presidential elections, the market's absolute error has been 1.5 percentage points, compared with 2.1 percentage points for the final Gallup poll.[4]

[*] J.D., University of Virginia School of Law, 2005. I am grateful to Molly Bishop, Roger Barkan, Brad Ervin, Professor Albert Choi, my father, Philip McBride Johnson, Raimy Kamons, Professor George Triantis, Professor Bob Sayler, Professor Robert Scott, and fellow classmates in my Business Transactions Seminar for being generous with their time and help. Special thanks to Professor Paul Mahoney whose patience, kindness, and insights always motivated me to work harder.

[1] Joyce Berg et al., Accuracy and Forecast Standard Error of Prediction Markets 13, 33 tbl.3 (July 2003) (unpublished manuscript, on file with the Virginia Law Review Association), available at http://www.biz.uiowa.edu/iem/archive/forecasting.pdf.

[2] Joyce Berg et al., Results from a Dozen Years of Election Futures Markets Research, *in* Handbook of Experimental Economic Results (C.R. Plott & V.L. Smith eds.) (forthcoming) (manuscript at 1, on file with the Virginia Law Review Association) (noting the demographic differences between the average market trader and the average voter).

[3] But see Berg et al., supra note 1, at 9, 33 tbl.3 (noting that a comparison between poll predictions and market predictions is slightly unfair to polls because they do not purport to predict election results, but only to predict results if the election occurred at the time of polling. The market, by contrast, predicts election results on a future election day).

[4] Justin Wolfers & Eric Zitzewitz, Prediction Markets, J. Econ. Persp., Spring 2004, at 107, 112.

The ability to predict elections is more of an art than a science. No formula exists to accurately predict presidential elections. Unlike a simple physics problem, one cannot plug inputs into an equation to compute the percentage of the vote a candidate will garner. Yet a few hundred traders swapping securities in Iowa can predict elections quite accurately. How? In short, each trader individually has some knowledge about how people will vote, and an efficient information market allows the traders to combine their individual knowledge to create a collective body of information sufficient to predict the election.

As of 2000, the Iowa Electronic Markets ("IEM"), conducted by the University of Iowa College of Business, had run forty-nine markets covering forty-one elections in thirteen countries.[5] While the IEM includes a variety of types of markets and securities, vote-share markets are a good illustration of how such markets work. In vote-share markets, a trader purchases a bundle of securities for one dollar. The bundle contains one security for each candidate, and each security pays one penny for each percentage point a candidate garners. Thus, a Kerry security in the 2004 presidential election would pay off forty-seven cents if Kerry won forty-seven percent of the popular vote. Through trading, the market establishes a price for each security that represents the market's prediction of the vote share each candidate will win.

The IEM is an example of an information market, a market type that has received much scholarly attention in recent years.[6] Broadly defined, information markets are those in which participants trade securities that offer returns contingent on the occurrence of an event. In a presidential-election market, the event might be the percentage of the vote each candidate garners, but it could be any definable and observable occurrence—such as whether Charlottesville will experience a hurricane before 2010 or whether the movie *Sahara* will gross $200 million. Because the return on securities in an information market is contingent on an event, the securities re-

[5] Berg et al., supra note 2, at 1.

[6] For a list of papers in the field, see Chris Massie, Papers on prediction markets, http://www.chrisfmasse.com/3/3/papers (last visited February 20, 2006).

semble bets and the market may therefore be described as one in which participants trade bets.[7]

The ability of information markets to accurately predict future events depends upon the existence of two conditions. First, the market must be efficient; at any given time, the market price must fully reflect, or aggregate, all relevant information.[8] Second, market participants must possess a degree of collective knowledge that approaches the level of information needed to predict the event with perfect certainty. In other words, a market price that encapsulates only twenty percent of the information needed to predict an event is less likely to provide an accurate prediction.

This Note will consider information markets of a more limited scope and scale than those that incorporate thousands of traders engaged in large-scale hedging and risk-spreading activities.[9] The markets discussed in this Note only concern events contingent on individual actions performed by large numbers of people, such as voting, and never contain more than a few hundred traders.

Given the successes of information markets, this Note will propose that decisionmakers in an administrative context use these markets to inform their decisions. Empirical research indicates that information markets generally are efficient and therefore provide informed predictions based on aggregated knowledge. Further-

[7] A concerned reader may question the legality of information markets given their similarity to gambling. The formal distinction between gambling and derivatives markets rests on the fact that the latter includes participants wishing to hedge risk while the former does not. Whether the IEM offers suitable hedging opportunities, however, is a moot point legally. The Commodities Futures Trading Commission has provided the IEM immunity from legal action through a "no action letter" based on its role as an educational tool.

[8] Ronald J. Gilson & Reinier H. Kraakman, The Mechanisms of Market Efficiency, 70 Va. L. Rev 549, 554 (1984) (quoting Eugene Fama, Efficient Capital Markets: A Review of Theory and Empirical Work, 25 J. Fin. 383, 383 (1970)); see also F.A. Hayek, The Use of Knowledge in Society, 35 Am. Econ. Rev. 519, 526–27 (1945) (writing about the ability of a market to aggregate information in addition to being able to distribute goods effectively).

[9] For a discussion of such large-scale markets, see Robert J. Shiller, Radical Financial Innovation (Cowles Found. Discussion Paper No. 1461, 2004), at http://cowles.econ.yale.edu/P/cd/d14b/d1461.pdf; see also Robert J. Shiller, The New Financial Order 3 (2003) (discussing these issues); Bertil Lundqvist, Securitization of Risk of Loss from Future Events, *in* New Dev. Securitization 2004, at 1141, 1148 (PLI Com. L. & Practice, Course Handbook Series No. 3023, 2004) (discussing the large-scale market for "catastrophe bonds").

more, theory and empirical research indicate that the market predictions are equivalent to, or better than, predictions by rival mechanisms. This Note will argue that, because of the quality of these market predictions, administrative decisionmakers should utilize information markets as an avenue to more informed decisions.

This Note also will analyze the types of decisions amenable to market predictions, illustrating that markets designed along lines similar to the IEM can provide valuable predictions for policymakers as long as such decisions do not alter the definition of the underlying bet. Decisions that alter the definition of a bet require a more sophisticated market design. Examples include decisions to abandon a policy based on a market prediction that it will lead to poor results, or decisions to select a policy based on a market prediction that it will outperform other proposed policies.

Part I of this Note will discuss the characteristics that define an information market. The first Section will explain how information markets operate, using IEM's vote-share market as an example. The second Section then will compare information markets to stock markets to illustrate how they differ. Both markets allow people to trade securities and rely upon the concept of efficiency, but the concept of accuracy is confined to information markets. Furthermore, price measures something different in each market: in an information market price always concerns an event happening over a fixed period, while in a stock market price reflects the performance of a company over an infinite period.

Part II will review the empirical research into the efficiency and accuracy of information markets, particularly those markets concerning elections, the sale of business products, box-office proceeds, and financial instruments. This analysis will reveal that information markets are fairly efficient and provide accurate predictions when compared with predictions of polls, experts, and other forecasting mechanisms.

Part III then will compare market predictions with those predictions made by group deliberation, voting, and opinion averaging. The comparison will focus on the amount of information underlying a prediction and will reveal that market predictions incorporate the most information. Part III also will analyze the costs and benefits associated with each predictive mechanism. Although information markets provide the most informed predictions, voting and opinion averaging cost less, and group deliberation has the benefit of educating both group members and outsiders. Thus, even though an information market often will lead to the most informed prediction, this benefit might be offset by other considerations that may lead an organization to choose an alternative mechanism.

In light of the discussion in Parts II and III regarding the efficiency and accuracy of information markets, this Note ultimately will suggest that administrative decisionmakers should incorporate such predictive markets in deciding among competing policies. Part IV will propose a market that administrative decisionmakers can use to evaluate the efficacy of different policies prior to implementation. Unlike traditional information markets in which securities pay off based on a single future event such as a presidential election, the securities in a policy market pay off based on different future events, some of which will not occur. For example, in a policy market, the return from Security 1 may be based on the effect of Policy *A* while the return of Security 2 may be based on the effect of Policy *B*. If *A* and *B* are mutually exclusive, then only one future event will occur: the effect of *A* or the effect of *B*. As a result, problems arise in determining how to compensate holders of the security whose payoff is contingent on the event that does not occur.

The first Section will show that compensating these security holders with either a fixed sum or the last price at which the security traded normally would alter the definition of the event on which the payoff depends. Instead of the payoff depending on the effect of the policy, it depends on the effect of the policy and the return if the policymakers choose not to implement the policy. The second Section, however, reveals that such a payoff does not necessarily alter the definition of the event. Rather, through clever market design that freezes the market price of those securities whose payoffs are tied to the rejected policy at the time that the agency decides not to implement a policy, it is possible to compensate holders of these securities without altering the definition of the event on which their payoff depends. While the proposed market design adds technical complications, it retains the spirit of IEM's vote-share market and has little to no effect on the thought process of the trader. More importantly, this Note's proposed market allows for accurate predictions within a market of multiple potential events, which is useful in the context of administrative rulemaking—where relatively uninformed decisions among competing policies with unknown effects must often be made. Thus, Part IV will provide a framework for administrative agencies to use market predictions as a decisionmaking tool when faced with a choice between possible future courses of action.

I. WHAT IS AN INFORMATION MARKET?

A. The Operation of an Information Market

While the IEM allows people to bet on various elections and political events, it is best known for its markets concerning presidential elections. In the IEM's vote-share market, participants' purchasing ability is typically limited to five hundred bundles of securities, each worth one dollar.[10] The securities are tradable contracts that return one dollar times the percentage of the vote obtained by a given candidate.[11] Thus, if John Kerry won forty-seven percent of the popular vote in the 2004 presidential election, a Kerry contract would pay off forty-seven cents.

1. Arrow-Debreu Securities

The securities offered in each bundle are called Arrow-Debreu securities.[12] With such securities, the holder of an intact bundle will receive one dollar regardless of the percentage of the vote each candidate garners. To see this in the context of a presidential election market, consider a bundle containing one security for each major candidate and a catch-all security that pays holders based on the percentage of the vote garnered by fringe and write-in candidates. A holder of all four securities will always receive a dollar, but the amount received from each security will vary. Figure 1 below illustrates this point graphically with reference to the 2004 presidential election.

---

[10] Berg et al., supra note 2, at 1 n.1.

[11] Id. at 3; see also Berg et al., supra note 1, at 3.

[12] Charles R. Plott & Kay-Yut Chen, Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem 7 (Cal. Inst. Tech., Social Sciences Working Paper No. 1131, 2002), at http://www.hpl.hp.com/personal/Kay-Yut_Chen/paper/ms020408.pdf (last accessed Nov. 9, 2005).

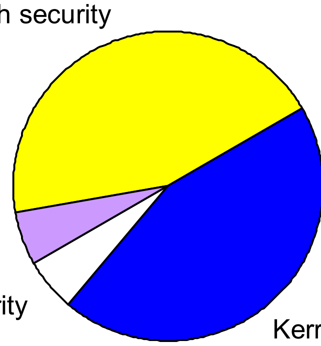

Figure 1: Presidential Election Security Bundle

To take a position on a candidate, a trader must either sell one of the securities in his bundle or purchase a security in the market—otherwise he would merely receive back the original investment. By purchasing a particular security, the trader takes the position that the market price is too low and that the candidate will garner a larger percentage of the vote. By selling a particular security in the bundle (or in his portfolio, assuming he has purchased other securities in the market), the trader takes the position that the market price is too high and that the candidate will garner a smaller percentage of the vote. Through such purchases and sales, the prices for the securities move and the opinions of the traders are aggregated.

Arrow-Debreu securities may be used in contexts other than vote-share markets. For example, imagine that an airline set up a market to predict the number of passengers it will carry in the following month. If the airline knows that its maximum capacity is one hundred passengers, it could offer two Arrow-Debreu securities: one that returned one cent for every passenger carried and another that paid holders one cent for every empty seat. A person holding both securities would always receive one dollar. With Arrow-Debreu securities, at any point in time price movements summed across all the securities in the bundle must equal zero. If one security's price increases there must be a corresponding decrease in the price of other securities, assuming the market is efficient. For example, if one extra passenger decides to fly there will be one fewer empty seat. The security whose payoff depends on the number of passengers will consequently pay off one more cent, while the security depending on the number of empty seats will pay one cent less.

Markets offering Arrow-Debreu securities present several advantages over markets offering only a single security. First, the market sponsor who offers the securities assumes no risk. Though the securities may change hands several times, and the sponsor may have to pay any particular participant more or less that the participant originally paid for the bundle, the sponsor's aggregate liability will always equal the amount that the sponsor originally received—one dollar times the number of bundles sold.[13] The sponsor therefore can benefit from the market's prediction merely for the cost of running the market, provided that traders participate. Compare this setup to one in which the market sponsor offers a single security—a security, for example, that returns one cent for every passenger that flies in a given month. Here, the sponsor will offer the security for some set amount.[14] The market participants will trade the security, establishing a price and thus a prediction. At the end of the month, the sponsor will pay the security holders based on the number of passengers that flew during the month. If more passengers flew than the sponsor expected, the sponsor loses

[13] Technically speaking, the IEM makes money running the market at the expense of market participants because it receives the interest on the money from the bundle sold before it has to pay out.

[14] The following discussion assumes that the security represents a contract between the market sponsor and the holder, entitling the holder to a payment from the market sponsor based on the outcome of an event. The contract, however, could be between the holder and the person who sold the security to the holder, entitling the holder to payment from the seller. In this case, the contract must include provisions to ensure that the seller has proper funds to satisfy his obligation to the buyer. The market sponsor could operate such a market in two ways. First, the market sponsor could be a party to the contract as the seller, until the initial buyer sells the contract and assumes the market sponsor's obligation. Second, the market sponsor could forego being an initial party to the contract and instead only act as an intermediary between the initial buyer and seller. The latter is superior to the former in that the market sponsor assumes no risk because it is never a party to the contract, but is inferior in that the market may not be liquid if an initial buyer must wait for another party to act as the seller.

money because it sold the security for a set amount that proved less than required to pay out to security holders. Alternatively, the sponsor makes money if fewer passengers flew than the sponsor's offering price predicted. Thus, the single security setup requires the sponsor to bear some risk, as it can either make or lose money based on the relationship between the number of passengers and its offering price. With Arrow-Debreu securities, in contrast, the market sponsor always pays out the same amount. Moreover, the sponsor need not worry about making an initial prediction—it merely offers the bundle for one dollar and the market makes the initial prediction.

Arrow-Debreu securities present the additional advantage of making it is easy to identify arbitrage opportunities. If at any point the prices of the contracts do not add to one dollar, an arbitrage opportunity exists because one security's price movement has not caused a corresponding price counter-movement in the price of the other securities. To identify an arbitrage opportunity in a vote-share market, one would look for the last price movement in one candidate's security that did not result in an opposite price movement in another candidate's security.

2. Establishing a Market Price and a Market Prediction

Once participants in the IEM receive security bundles, they may begin to trade. Information aggregation begins once a person buys or sells a contract. Suppose that a security bundle for the 2004 presidential election contains a George W. Bush contract and that traders are willing to buy or sell at forty-seven cents. This amount represents an equilibrium price based on all available information regarding the election. If a trader discovers positive news about Bush, he would buy Bush contracts at forty-seven cents, driving up the price of a Bush contract. He would continue to buy as long as he believed the price was too low. Eventually, the price would incorporate the value of this new information about Bush and the trader would stop buying. Each successive trade would disclose more of the trader's private information until it was completely aggregated into the market. The prices paid along the way would represent the market's original information plus the incremental portion of the new information disclosed by each of the previous purchases.[15]

B. Information Markets Compared with Stock Markets

Stock markets are a familiar reference point for discussing the operation of markets. Accordingly, this Section compares information markets to stock markets to explore the extent of similarity between the markets. In addition, this Section provides a technical overview of efficiency and accuracy, concepts that describe the amount of information encapsulated by price in a given market, a central concern in devising a viable information market for administrative decisionmaking.

At first glance, information markets and stock markets appear similar because participants trade securities in each. Further, the concept of efficiency extends to both types of markets, with each producing a price-time series for securities. Nevertheless, information and stock markets are different in two important respects. First, price measures something different in each market. In an information market, price measures the expected outcome of an event. In a stock market, by contrast, price measures the net present value of a company from the present into infinity. To put it another way, price in an information market is a measurement of a specific, tangible outcome, whereas price in a stock market is a measurement of the future going forward.[16]

The second salient difference between information and stock markets is that, although the concept of efficiency extends to both, the concept of accuracy extends only to information markets. Accuracy refers to the ability of an information market to predict the outcome of an event correctly. An accurate information market must be efficient, meaning it is able to aggregate all known information and the difference between the known information, and the information theoretically necessary to predict the event with certainty must be negligible—that is, the degree of knowledge about

[15] Although this description is illustrative of the way in which the information aggregation process works, "the exact method by which information gets into the market is unknown." Charles R. Plott, Markets as Information Gathering Tools, 67 S. Econ. J. 1, 8 (2000).

[16] See Berg et al., supra note 1, at 6–7.

the event must be high. Suppose the information aggregated in a given market represents only twenty percent of the information needed to predict an event with certainty. While the market may be efficient in that it successfully aggregates all available information into the prediction, it might not be accurate. Without the remaining eighty percent, any correct predictions made by the market would result mostly from chance.

With an information market, it is possible to quantify the knowledge that the market lacks by measuring the difference between an event's predicted outcome and its actual outcome. With a stock market, by contrast, it is impossible to quantify the knowledge absent from the market because there is no way to measure the difference between an asset's trading price and its "true" value, which is the price at which market participants would buy or sell if they knew everything about the asset.

With assets traded in an information market, it is always eventually possible to determine the assets' "true" value because at some point the event of interest occurs. For example, the vote-share market of 2004 ended when officials released the final vote count. The final vote count is an observable event from which one can determine the "true" value of each candidate's securities. Because Kerry garnered forty-eight percent of the vote, the true value of a Kerry security was forty-eight cents. The difference between the "true" value and the security's trading price at any point in time represents what traders did not know at the time. The key to measuring our degree of knowledge is the existence of an observable event against which we can later evaluate earlier market predictions.[17]

Stock markets, however, do not have observable events, aside from bankruptcy, that reveal the "true" stock price of a company. There is no discrete moment at which one can observe the "true" price of a stock because the price always depends on the unobservable future. Consequently, it is impossible to measure whether the

[17] This Note does not address the potential use of information markets as tools for understanding the extent of our knowledge. For example, it may be feasible to measure the extent to which we understand global warming by measuring the difference between the trading price and the true price of a security that is contingent on average global temperature.

trading price of a stock differs from its "true price," and hence, the accuracy of the price.

As a result, one must use less sophisticated tools to assess the significance of a stock's trading price. Instead of measuring accuracy, analysts measure market efficiency, which indicates the extent to which price incorporates known information. If a market is efficient then confidence in market prices is justifiable because market prices are based on all known relevant information, even if there is no way ex post to examine whether those prices accurately reflected the stocks' "true" prices.

To measure market efficiency, analysts generally test whether the time series of a stock's price follows a "random walk,"[18] which means that for any time *t* the price at time *t*+1 depends only on the price at time *t* plus a random amount. If a stock price's time series exhibits random-walk behavior, it suggests the market is efficient. In an efficient market, future price movements depend only upon the random discovery and direction of new information because by definition the current price incorporates all known information. When the time series of a stock's price exhibits random-walk behavior, it indicates that price movement depended on the random discovery of information rather than market inefficiency, which one must assume would produce non-random price movements. Thus, the random-walk test suggests efficiency indirectly by showing that price movements at *t*+1 were random, which supports the conclusion that the price at time *t* incorporated all known information.

One can apply the random-walk test to any market in which trading produces a price-time series for securities. Because both information and stock markets meet this requirement, it is possible to measure the efficiency of each type of market using the random-walk technique. Evidence that price moves randomly creates confidence that no profitable trading strategies exist because those strategies require traders to take set actions based on patterns in the time series of a stock's price. If the time series is a random walk then patterns do not exist and trading strategies based on price patterns will not be profitable.

[18] Ronald J. Gilson, The Law and Finance of Corporate Acquisitions 158–59 (1986).

Two important observations about the accuracy and efficiency of information and stock markets emerge. First, trading prices in information and stock markets measure different phenomena. Prices in information markets measure what will happen during a fixed time period. Prices in a stock market measure what will happen over an infinite and changing period. Therefore, one can assess the accuracy of information markets because, unlike "true" stock prices, the events predicted by information markets are observable. Second, the concept of efficiency extends to both information and stock markets, and it is possible to measure the efficiency of each market by studying whether the time series of prices follows a random walk.

II. EMPIRICAL STUDIES OF THE EFFICIENCY AND ACCURACY OF INFORMATION MARKETS

A. Efficiency

A number of recent studies address the efficiency of information markets. Several studies of markets predicting presidential elections have shown that they are efficient. For instance, Professors Berg, Nelson, and Rietz concluded that the time series of prices in the IEM are "consistent with efficient random walks."[19] Professors Rhode and Strumpf found similar results in their analysis of presidential betting markets between 1868 and 1940.[20] A time-series analysis of prices indicates that presidential betting markets in this period exhibited weak-form efficiency, fell slightly short of semi-strong efficiency, and did not demonstrate strong efficiency.[21]

[19] Berg et al., supra note 1, at 25.

[20] Paul W. Rhode & Koleman S. Strumpf, Historical Presidential Betting Markets, J. Econ. Persp., Spring 2004, at 127, 127, 136–38.

[21] Id. at 136–38. The terms weak, semi-strong, and strong efficiency delineate three forms of market efficiency. Weak efficiency represents the situation in which price reflects all information indicated by past price movement. Semi-strong efficiency exists where the price reflects all public information, such as the information found in annual reports. Strong efficiency describes the more extreme circumstance in which the market price incorporates all private information, including information known by a single non-trading individual. James H. Lorie & Mary T. Hamilton, The Stock Market: Theories & Evidence 70–71 (1973), *reprinted in* Gilson, supra note 18, at 158, 158–59; see also Gilson & Kraakman, supra note 8, at 558, *reprinted in* Gilson, supra note 18, at 172, 174. Typical capital markets do not exhibit strong-form efficiency, but they come close. Rhode & Strumpf, supra note 20, at 138.

Professors Rhode and Strumpf also investigated the existence of arbitrage opportunities in the early presidential betting markets. They found virtually identical prices within various New York markets, with only small price variations across cities.[22] As the authors point out, the results are remarkable given the limitations on long-distance communication during the era.[23] The investigations of Professors Wolfers and Zitzewitz buttress these arbitrage findings. Studying the securities offered by Tradesports.com and World Sports Exchange that paid off if Arnold Schwarzenegger won the 2003 California gubernatorial race, Wolfers and Zitzewitz found price movements nearly in sync across the exchanges and virtually no arbitrage opportunities.[24] The same authors also investigated the possibility of manipulation of election markets. Ultimately, they concluded, the several known manipulation attempts failed to succeed for more than a brief period, as illustrated by the effects of random bets placed on the IEM and failed attempts to manipulate the early presidential betting markets studied by Rhode and Strumpf.[25]

Results are mixed outside the electoral arena. On one hand, results are encouraging for information markets in which traders bet on financial instruments. For example, in one analysis by Professors Wolfers and Zitzewitz comparing Tradesports.com bets on movements of the Standard and Poor's ("S&P") index with actual S&P options offered on the Chicago Mercantile Exchange, the researchers found that extremely unlikely outcomes were overpriced on Tradesports.com.[26] Research by Professor Tetlock, however, challenges this as an exceptional result. Professor Tetlock discovered no long-shot bias on Tradesports.com contracts concerning financial events, while finding such bias for contracts concerning sporting events[27]—a phenomenon observed by other scholars at

[22] Rhode & Strumpf, supra note 20, at 136–37, 137 n.3.

[23] See id. at 135 (describing the lack of national information sources as a barrier to accurate forecasts).

[24] Wolfers & Zitzewitz, supra note 4, at 116 & fig.4.

[25] Id. at 119.

[26] Id. at 117–18. Long-shot biases are common, however. Such biases have been observed in horse-track betting and option markets for options that are strongly "in" or "out" of the money. Id. at 117.

[27] Paul C. Tetlock, How Efficient are Information Markets?: Evidence from an Online Exchange 34 (Jan. 2004) (unpublished manuscript, on file with the Virginia

race tracks.[28] Moreover, Tetlock concludes that betting markets on financial instruments are efficient despite very low liquidity and volume compared to traditional financial markets.[29]

On the other hand, where the incentives offered by a market are low or the number of traders few, the results have not been as positive for information markets. In an analysis of the Hollywood Stock Exchange ("HSX"), where participants bet on box office sales using play money, researchers found the market less efficient than the IEM.[30] In another study by Kay-Yut Chen and Professor Plott conducted at Hewlett-Packard ("HP"), the researchers compared small-scale information markets to HP forecasting procedures.[31] In their experiments the number of market participants varied between seven and twenty-six people,[32] and in each experiment participants were given Arrow-Debreu securities[33] with payoffs contingent on upcoming sales numbers.[34] As noted earlier, the market price of Arrow-Debreu securities should always sum to the price of the bundle. In all twelve experiments, the summed prices of the securities exceeded the cost of the bundle.[35] Thus, in all the experiments theoretical arbitrage opportunities existed, indicating market inefficiency. Yet the question arises whether the inefficiencies of the HSX and HP experiments stem from inadequate incentives. In the HSX, participants used play money and in the HP experiment participants were given the securities for free. Perhaps a larger stake in the outcome would eliminate the observed inefficiency. In addition, it is possible that the inefficiencies would disappear if the software trading platform notified participants of ar-

Law Review Association), at http://www.mccombs.utexas.edu/faculty/Paul.Tetlock/papers/Tetlock-Efficient_Info_Markets-01_02.pdf.

[28] See, e.g., Richard H. Thaler & William T. Ziemba, Anomalies: Parimutuel Betting Markets: Racetracks and Lotteries, J. Econ. Persp., Spring 1988, at 161, 162–67, 170.

[29] Tetlock, supra note 27, at 34.

[30] David M. Pennock et al., The Power of Play: Efficiency and Forecast Accuracy in Web Market Games 10–11 (NEC Research Inst. Tech. Report 2000-168, 2001), at http://artificialmarkets.com/am/pennock-neci-tr-2000-168.pdf (last accessed Nov. 9, 2005).

[31] Plott & Chen, supra note 12, at 5.

[32] Id. at 19.

[33] Id. at 7.

[34] Id. at 19.

[35] Id. at 16.

bitrage opportunities. But as the number of traders decreases there are fewer people to catch and correct the mistakes of others, making longer stretches of inefficiency likely.[36]

B. Accuracy

There are two ways to evaluate the accuracy of an information market. One is to determine whether a particular market predicts an event correctly. The other is to assess the general ability of markets to predict a certain type of event correctly. In the context of market predictions of presidential election results, the first method would look at whether the market accurately predicted the presidential election of 2004; the second, whether markets for presidential elections generally provide accurate predictions. This Section focuses mostly on the latter.

Observers have long recognized the accuracy of election markets. In 1904, Andrew Carnegie commented on the election, "[f]rom what I see of the betting . . . I do not think that Mr. Roosevelt will need my vote. I am sure of his election."[37] In fifteen elections between 1884 and 1940, the voting odds in mid-October have failed to predict the November winner only once.[38] Given that polls did not emerge until the 1930s, these results indicate that the market can predict elections accurately without piggybacking on polling data.[39]

The results are also impressive in more recent presidential election markets. The average error of the last four IEM presidential election markets was 1.37% on the election's eve[40] and 1.5% in the week before the election.[41] In 596 comparisons of the accuracy of the markets and polls, the markets proved more accurate than the polls 451 times.[42] The superior predictive power of the market holds regardless of the timing of the comparison—whether made

[36] Evidence on the efficiency of thin markets is scant. See infra note 59 and accompanying text.
[37] Rhode & Strumpf, supra note 20, at 132 (citing Carnegie, Returning, Talks on Roosevelt, N.Y. Times, Oct. 24, 1904, at 1).
[38] Id. at 129.
[39] Id.
[40] Berg et al., supra note 1, at 5.
[41] Wolfers & Zitzewitz, supra note 4, at 112.
[42] Berg et al., supra note 1, at 33.

months or days before the election.[43] This comparison, however, is slightly unfair to polls, which do not purport to predict election results on election day—they are only intended to predict election results if the election were to occur at the time of polling. Market predictions, by contrast, predict election results on a future election day.[44]

Information markets in other contexts also appear to be mostly accurate, though less so than presidential election markets—and overall, information markets normally outperform rival predictive mechanisms. In the summer of 2004, the IEM ran two information markets to predict Google's stock price at the close of the first day of trading.[45] In order to avoid the problem of underpricing shares, Google used an auction to set its initial public offering ("IPO") price.[46] In the auction, potential purchasers bid on shares, which were allocated to the highest bidders, and the sale price was set at the amount paid by the lowest bidder who received a share.[47] Google ultimately underpriced its offering by just over fifteen percent, costing it $300 million.[48] In contrast, the information market's price proved quite accurate. The night before Google filed its final prospectus, the information market predicted a price that was within 3.8% of Google's share price after the first day of trading.[49] Of course, this result provides only a single data point and does not suggest that information markets will generally predict appropriate IPO prices. But it is a salient data point that fits nicely with Professor Tetlock's observation that information markets concerning financial events are remarkably efficient given their low liquidity.[50]

Similar research has examined the accuracy of the HSX, which uses play money in trading, and the HP experiment, which involved a limited number of traders. In both cases, the markets per-

[43] Id. at 1.

[44] Id. at 6.

[45] Joyce Berg et al., Searching for Google's Value: Using Prediction Markets to Forecast Market Capitalization Prior to an Initial Public Offering 9–12 (2005) (unpublished manuscript, on file with the Virginia Law Review Association), available at http://www.biz.uiowa.edu/faculty/trietz/papers/google.pdf.

[46] The underpricing of initial public offerings is a well-known phenomenon that costs companies about 15% of their potential IPO proceeds. Id. at 2.

[47] Id. at 3.

[48] Id. at 5, 22.

[49] Id. at 17.

[50] Tetlock, supra note 27, at 34–35.

formed acceptably, making predictions that nearly equaled or surpassed the accuracy of rival mechanisms. The HSX's predictions of box office receipts for 109 movies from March 3, 2000, to September 1, 2000, had an average error rate of 31.5%,[51] as compared to the rival expert's rate of 27.5%.[52] Although the market's 31.5% average error rate appears high, the correlation between the market's predictions and the box office receipts was 0.94, meaning the market's predictions explained 94% of the variation in box office results.[53] In the HP experiment, eight of the twelve studied markets included an HP sales forecast against which to compare the markets' predictions. In four of the eight experiments, the price of the last trade outperformed the official HP forecasts.[54] Across the eight experiments, the market predictions had an average error rate of 20.8%; in contrast, the HP forecasts had an average error rate of 21.96%.[55]

The HP and HSX results may be misleading, however, given the possibility that one mechanism incorporated or "piggybacked" on the prediction of another mechanism. For example, in the HP experiment, the company's official forecasts were made after the market closed, and evidence indicates that these forecasts took the price produced by the information market into account.[56] Without the information provided by the market, it is likely that the official forecasts would have been less accurate. In the case of HSX, the market may have relied on expert predictions, the experts on the market, or both upon each other.[57] Given the presence of mutual

[51] David M. Pennock et al., Extracting Collective Probabilistic Forecasts from Web Games 174, 177 (Proc. 7th ACM SIGKDD Int'l Conference on Knowledge Discovery & Data Mining, 2001), at http://portal.acm.org/citation.cfm?id=502512 (last accessed Nov. 9, 2005); see also Martin Spann & Bernd Skiera, Internet-Based Virtual Stock Markets for Business Forecasting, 49 Mgmt. Sci. 1310, 1321–23 tbls.3–5 (2003) (looking at a later and larger data set of HSX's prices and reporting similar results).

[52] Pennock et al., supra note 51, at 177; Spann & Skiera, supra note 51, at 1320–23.

[53] Spann & Skiera, supra note 51, at 1320–23; see also David S. Moore & George P. McCabe, Introduction to the Practice of Statistics 144, 168–70 (2d ed. 1993) (explaining correlation in regression analysis).

[54] Plott & Chen, supra note 12, at 20 tbl.2.

[55] Id. The average rates are derived by averaging the percentages of error for events two through nine.

[56] Id. at 5.

[57] Pennock et al., supra note 51, at 176–78.

reliance on the other's outcome, it is likely that the accuracy of each mechanism would be quite similar.

Ultimately, however, the empirical results suggest that information markets are generally efficient and somewhat accurate, particularly when compared to rival mechanisms. Markets concerning presidential elections and financial events in particular appear to be efficient and provide accurate predictions. Markets such as HSX, in which participants trade in play money, are less efficient and provide less accurate predictions. The same holds for the HP experiments in which there are only a few traders who do not purchase, but rather are given, securities. Nevertheless, all the markets studied provide predictions that nearly equal or exceed the predictions of rival mechanisms.

In theory, it makes sense that an efficient information market should be able to outperform experts. The information-aggregating property of such markets suggests that they should include all known relevant information. A rival mechanism that lacks such a capability is unlikely to equal an efficient market's prediction. Particularly in the context of events contingent on large groups where information about the behavior of many individuals is likely to be distributed broadly across society, both experts and small groups will have trouble collecting the information and will thus be handicapped. In these contexts, however, information markets are likely to be efficient because information about other people is easy to interpret, so traders can incorporate it into the market without extensive analysis or cost. Additionally, information about others is easy to acquire because for every person, it is likely that another exists with information about the future actions of the first.

III. COMPARING INFORMATION MARKETS TO OTHER PREDICTIVE MECHANISMS

While information markets provide better predictions, other predictive mechanisms—namely, group deliberation, voting, and opinion aggregation—offer other benefits that cannot be achieved through markets. For example, the implementation of a market requires investments in technology and trader time, while voting requires less technology and time. Further, neither voting nor information markets educate people to the same extent as group deliberation. In practice, a decision to implement one mechanism over another will depend on the cost of operation weighed against the value of participant education and predictive accuracy. Ultimately, when accuracy is the paramount concern, a market is the best choice.

To compare the quality of predictions made by different mechanisms, the first step is to determine the measure of quality. Two possible approaches stand out: measuring predictive accuracy, or measuring the amount of information on which the prediction is based. Although accuracy is the most important characteristic of a prediction, it is not the best metric to assess quality for two reasons. First, one mechanism might be more accurate than another merely because it relied on data provided by the other mechanism. For example, the expert in the HSX discussion might look at market performance prior to making his prediction, thereafter making a small refinement. In such a case, the market is the better predictor because it carries most of the predictive burden. Second, it is possible that one mechanism enjoys greater access to information than another. Say we compared a market prediction regarding whether the United States would invade Iraq in 2003 to a group prediction made by President Bush's senior advisors with access to classified information. If the senior advisors outperform the market, it does not indicate that group predictions are better than market predictions, but only that the senior advisors possessed superior information. Given these problems, it is more useful to assess the quality of prediction based on the mechanisms' relative ability to aggregate information. The mechanism that generates predictions based on the best information will be superior.

By definition, an efficient market aggregates all information, rendering it impossible for other mechanisms to outperform the market prediction. In situations involving a small number of traders, however, such as the HP experiment, the market may not be efficient.[58] Although evidence about the efficiency of markets with

[58] Market efficiency is not a purely academic issue; companies such as Microsoft, Eli Lilly, Hewlett-Packard, and Intel all have experimented with small-scale information markets in place of traditional forecasting mechanisms. See Barbara Kiviat, The End of Management? With Experimental Markets, Workers Are Betting on Their Company's Future—and Moving in on the Boss's Domain, Time, July 12, 2004, at A4, 1, 2; see also Thomas W. Malone, Bringing the Market Inside, Harv. Bus. Rev., Apr. 2004, at 106, 107–11 (providing an excellent description of some uses of information markets in the business context). For a business considering whether to use consultants or

few traders is scant, scholars have proposed market-like mechanisms that mitigate the thin-market problem, in which a buyer or seller is not always readily available to take the opposite side in a transaction.[59] Thus, even in situations where there are too few traders to implement a successful market, information aggregating mechanisms similar to information markets may allow for the efficient aggregation of information.

Two factors distinguish information markets and related information aggregating mechanisms from other predictive mechanisms: participants in a market can both reveal private information and stress its importance. The absence of these two factors in the other predictive mechanisms accounts for their less efficient information aggregation.

Consider the following hypothetical as a framework for comparing the predictive mechanisms. An airline CEO would like to know how many passengers the airline will carry over Christmas week and must select the mechanism that will provide the best predic-

a market to predict sales in the following year, market efficiency is critical, and the value of the market's prediction will hinge on its efficiency.

[59] One scholar has argued on mathematical grounds that a market with as few as ten participants should provide good results. Robin Hanson, Combinatorial Information Market Design, 5 Info. Sys. Frontiers 107, 110, 110 fig.1 (2003) (showing in a chart that when an individual trader provides one-tenth of the total number of estimates, market accuracy is good, meaning that even if there is only one estimate, a market containing ten traders should provide sufficiently accurate data). While this evidence indicates that information markets can make good predictions even in thin markets, the performance of such markets is still an open empirical question. It is important to note that information markets are not the only market-based mechanisms that can aggregate information and perform well in thin markets. For example, market-scoring rules expand on conventional scoring rules, in which experts are paid increasing amounts as their predictions about an event approach the actual outcome. Thus, when the temperature is 72 degrees, a guess of 70 degrees will result in higher pay than a guess of 64 degrees. Under market-scoring rules, subsequent experts can make additional predictions to further approximate the actual answer. For the right to make a new prediction, subsequent experts must pay the house an amount equal to the reward the first expert would have received. Michael Abramowicz, Information Markets, Administrative Decisionmaking, and Predictive Cost-Benefit Analysis, 71 U. Chi. L. Rev. 933, 959 (2004); Hanson, supra, at 110-11. If the new prediction is worse than a previous prediction, the expert must pay the house the difference between his lower reward and the higher reward associated with the previous answer. In essence, the mechanism is a clever way to guarantee market activity even when the market is thin.

tion. Which of the four following mechanisms would the CEO choose?[60]

1. A market like the IEM vote-sharing market in which the public trades Arrow-Debreu securities offered at one dollar per bundle; the first security in the bundle pays one cent for each passenger carried and the other pays one cent for every empty seat;
2. A vote in which each pilot is asked to identify the range of passengers that the airline will carry, from among zero to twenty, twenty to forty, forty to sixty, sixty to eighty, and eighty to one hundred, and where the final prediction is the range selected by the most pilots;[61]
3. A prediction obtained by asking each pilot to predict a specific number of passengers and averaging the results; or
4. The consensus of the airline's ten smartest employees following group deliberation.

Assume that the maximum capacity of the airline is one hundred passengers and that the market, pilots, and employees all have the same information in aggregate.

In deciding between the mechanisms, the CEO must engage in a two-step analysis. First, the CEO must determine which mechanism will provide the most informed answer. Second, the CEO must analyze the costs and benefits associated with each mechanism.

A. The Most Informed Prediction

Of the four choices, the second, pilots voting for a range, provides the least informed prediction. This becomes evident with two

[60] This inquiry is inspired by Cass Sunstein. See Cass R. Sunstein, Group Judgments: Statistical Means, Deliberation, and Information Markets, 80 N.Y.U. L. Rev. 962, 963 (2005).

[61] Pilots vote on ranges instead of whole numbers to avoid the odd results that occur when many pilots have similar estimates but select different numbers. For example, say the first pilot selected one, the second pilot two, and so on, until the eighth pilot selected eight. Then, imagine that the ninth and tenth pilots each selected forty. The prediction of forty would win the vote, but the outcome of the vote would actually indicate that the majority of the pilots believed the answer was between one and eight.

assumptions. First, assume that there are three relevant pieces of information needed to determine how many passengers the airline will carry. Next, assume that the information is distributed among a sample of ten pilots as follows: the first eight pilots know the first two pieces of information, each of which suggests that the correct range is zero to twenty passengers, while the last two pilots know the final piece of information, which is critical and suggests a range of sixty to eighty passengers. In the event of a vote, the first eight pilots would lead the CEO astray, while the last two will be unable to stress that the final piece of information is more important than the first two pieces. Thus, a straightforward vote is problematic because it does not allow a voter to reveal his information or its quality.

What if the company polled the pilots first, showed them the results, and then had the pilots vote? The pilots would know that eighty percent of their colleagues believed that the lowest range was correct, while twenty percent thought the correct range was sixty to eighty. After the poll, each of the pilots could deduce that he did not have all the information, but the polling mechanism would not allow them to communicate this private information to each other. Pilots 9 and 10 would wonder why pilots 1 through 8 voted for the lowest range, while pilots 1 through 8 would wonder why pilots 9 and 10 voted for the sixty to eighty range. Moreover, there is no way for the pilots to deduce the importance of the missing information. After seeing the results of the poll, pilots 9 and 10 would suspect that the number of passengers would be fewer than sixty, but they would be unable to quantify the difference. Similarly, the first eight pilots would suspect that the number of passengers was greater than twenty, but not by how much. In the end, voting coupled with polling does a poor job of aggregating information because it reveals to voters only that they lack information, not the importance of the missing information.

In contrast, the third method, asking each pilot for an opinion and then averaging those guesses, is likely to produce a good prediction. Numerous social-science studies demonstrate that the average of group guesses will be close to the correct answer and will exceed the accuracy of the vast majority of individual guesses.[62]

[62] Sunstein, supra note 60, at 971–72.

This result is not surprising. If information indicating that the answer is higher or lower than the true answer is distributed somewhat evenly among the pilots, the high guesses should counteract the low guesses. If critical information rests in the hands of a few pilots, however, problems emerge because those pilots have no ability to stress the importance of their information; their vote carries the same weight as the votes of the other pilots. For instance, assume that there are three relevant pieces of information; that the first two pieces of information indicate that the airline will carry ten passengers and the third indicates seventy passengers; that the third piece of information is five times more important in predicting the actual number of passengers than are the first two; and that pilots one through eight know about the first two pieces of information, while pilots nine and ten know only of the third piece. Under these assumptions, the actual number of passengers the airline will carry would be fifty-three,[63] but the pilots' prediction would be twenty-two.[64] Thus, a prediction reached by averaging the pilots' opinions may be superior to a prediction based only on individual guesses; this technique becomes less accurate, however, when valuable information rests with only a few individuals who cannot communicate the importance of that information to others.

In contrast to voting and averaging mechanisms, group deliberation has the ability to match the predictive value of an efficient market. The communication required to generate a group prediction facilitates information exchange and the creation of new ideas. Thus, in theory, a group prediction has the ability to exceed a market prediction, at least in the short term, until traders are able to discover the newly generated information and restore the market's efficiency. Unfortunately, studies indicate that this optimism regarding group deliberation is unwarranted. In practice, "groups tend to do about as well as or slightly better than their average member, but not as well as their best members," suggesting that groups aggregate information poorly. Groups typically fail to aggregate information for two broad reasons: members do not make private information known, and they may misinterpret the significance of certain information. Individual group members may re-

[63] That is, $(10 + 10 + (5 \times 70)) \div 7 = 53$.

[64] $((8 \times 10) + (2 \times 70)) \div 10 = 22$. This represents the average of the pilots' votes.

frain from sharing private information due to informational influences that convince them that their information is incorrect, social pressures resulting from fear of being disliked, or aversion to being a sole dissenter.[65] Uneven distribution of information among the group may also contribute to the lack of information aggregation as groups may give more weight to information held by more people.[66] In addition, the order in which the members contribute may result in failure to emphasize information appropriately, as the collective force of previously stated opinions may induce subsequent speakers to change their minds.[67] Thus, while groups have the potential to provide a better answer than a market, they are unlikely to do so.[68]

In the final comparison, the most informed answer will come from an efficient information market. That said, as the number of participants decreases, the market may become less efficient, and group deliberation or consultation with an expert might produce superior predictions.[69]

B. The Costs and Benefits Associated with Each Predictive Mechanism

The cost-benefit analyses for both the voting and averaging mechanisms are similar. Both share two sets of costs: the time it takes the pilots to participate and for the results to be counted, and the incentives needed to encourage the pilots to disclose informa-

[65] Sunstein, supra note 60, at 984–86.

[66] Id. at 994–97.

[67] Id. at 999–1000.

[68] Id. at 982. The most optimistic line of studies indicates that only highly competent groups answering factual questions with demonstrably correct answers perform better than individual members. Id. at 1007–09. These studies do not contend that groups aggregate all or nearly all information; rather, they argue that under limited conditions a group decision may outperform the best individual decision. Another line of research indicates that groups perform slightly better if members are told that other members have particular skills and expertise. Id. at 1019–20. Again, however, there is no claim that the group aggregates all information. Rather, results improve because when an individual offers information in his area of expertise other group members may give that information additional weight. Id.

[69] The CEO also might use market-scoring rules as an alternative to the information market; however, it is unclear whether such an approach would produce a prediction superior to one produced through group deliberation. See Abramowicz, supra note 59, at 959; Hanson, supra note 59, at 110–11.

tion. Both also offer the benefit of allowing employees to participate in the decisionmaking process. The averaging technique provides predictive value superior to the voting approach, however, thus making it the preferable mechanism.

The costs associated with using information markets include building an infrastructure to support the market, providing incentives to encourage employees to trade, and the costs of either employees' time (if the employees are trading directly) or lawyers' fees (if trading is open to the public).[70] In addition to the time consumed when employees act as traders, time is also consumed when traders contact employees to learn more about an issue.[71] The chief benefit of an information market is the quality of the predictions it produces. Beyond the prediction, an additional benefit is that the success of employees who trade in the market provides an objective measure of those employees' abilities to make business predictions. Companies could use employees' trading histories to identify those who make good predictions. Like the other internal mechanisms, group deliberation consumes employee time and thus generates costs. Unlike the other mechanisms, however, group deliberation educates the participating employees, thereby benefiting the organization generally. Quantitative predictions, such as those provided by the market, voting, and averaging mechanisms, are atomic—that is, they are pieces of information not easily broken into constituent parts. A person who knows a market price does not know the pieces of information that caused individual traders to buy and sell. Thus, price is like a black box, and considerable effort is required to understand the way in which a particular piece of

[70] Companies that have experimented with information markets, such as Microsoft, Eli Lilly, HP, and Intel, have all used employees as traders. See Kiviat, supra note 58, at 1–3. One possible way for companies to reduce costs is to use open-source information market software or build proprietary software, as Microsoft has done. See also Masse, supra note 6 (describing various ongoing open-source projects).

[71] Another plausible way to reduce costs is to use market-scoring rules. One might argue that because market-scoring rules work in thin markets, a company could save employee time by adopting market scoring rules instead of an information market. Professor Hanson seems to suggest, however, that by the time the information market contains ten or more participants it will perform slightly better than an approach based on market-scoring rules. Hanson, supra note 59, at 110. Thus, the argument in favor of using market-scoring rules over an information market may have merit, but its scope is limited and it is applicable only under conditions when a group might be a better alternative.

information may have affected price.[72] In contrast, following group deliberation, the participants may write up a report explaining the reasons behind their decisions, or they may convey their ideas verbally to other employees. Quantifying the benefits of such information exchange and internal education is difficult, but they are clearly significant and valuable.

In sum, the CEO must weigh three factors—cost, quality of the answer provided, and the value of internal education—in selecting an appropriate mechanism. Averaging employee opinions is an inexpensive way to produce a decent prediction. Both information markets and group deliberation cost more than averaging mechanisms, but these mechanisms offer countervailing benefits. Efficient information markets provide high-quality answers, while group deliberation educates the organization as a whole. Ultimately, the CEO's decision will depend upon the relative significance of these various costs and benefits in the context of the firm's overall business objectives.

IV. USING MARKET PREDICTIONS TO MAKE DECISIONS TODAY

A. A Problem: Decisions That Alter a Bet's Definition

The two previous Parts addressed the justifications for using information markets to make predictions. Part II reviewed empirical evidence showing that information markets are efficient and provide accurate predictions in certain contexts. Part III illustrated that information markets are superior to other predictive mechanisms, such as group deliberation, voting, and averaging guesses. Given the predictive value of information markets, a practical

[72] Consider securities litigation in which a false merger denial by a target artificially depressed its stock price. See Levinson v. Basic Inc., 786 F.2d 741, 745 (6th Cir. 1986) (illustrating endorsement of the fraud on the market theory), *vacated*, 485 U.S. 224 (1988). To argue that the denial depressed the stock price, plaintiffs often rely on event studies, statistical analyses that examine the effect of an event, such as a corporate statement, on a dependent variable, such as a corporation's stock price. Beyond event studies, few other tools reveal precise information about the events that affected a stock's price from the price signal of the stock. Thus, it suffices to say that price is almost atomic. See Gilson, supra note 18, at 214 (discussing the use of event studies and cumulative abnormal return analysis in judging the impact of information disclosures on security prices).

question arises: How can these predictions about the future be used to make today's decisions?[73]

Three types of decisions might benefit from the use of market predictions. The first two types illustrate cases in which the market's prediction directly enables a decision. The third type of decision presents difficulties, however, because once the decision occurs, the definition of the bet on which the prediction was made changes.

The first type of decision amenable to market prediction is a contingency decision. Such decisions are incidental to the prediction, and thus do not affect the prediction itself. For example, a movie theater might buy more concessions for the next month because the HSX predicted three upcoming blockbuster movies. In this scenario, the decision to buy the concessions does not affect the underlying bet regarding the success of the movies. Another example might be Google using a prediction of its stock price to set a better offering price for its IPO. The price at which the stock trades at day's end should be independent of the price at which Google offers the stock on the morning of its IPO. Google's decision, though it depends on the prediction, does not affect the underlying bet. Individuals who bet on Google's share price do not care whether Google offers its stock at forty dollars a share or sixty dollars a share; they care about the stock price in the secondary market at the end of the first day, the event that determines the success of their bets. Because contingency decisions do not affect the market's prediction, decisionmakers can act based on such pre-

[73] There are a number of ways to use information markets that this Note does not discuss. One significant potential use would provide hedging opportunities to traders. For example, imagine a market in which securities paid holders contingent on a weather event, such as rain. A business that received most of its revenue on sunny days could then hedge against revenue loss by buying securities that paid upon rain. See Melanie Cao et al., Weather Derivatives: A New Class of Financial Instruments 1–4 (Apr. 2003) (unpublished manuscript, on file with the Virginia Law Review Association), available at http://www.mgmt.utoronto.ca/%7Ewei/research/JAI.pdf (introducing the concept of weather derivatives markets). Information markets could also enable traders to determine the effects of a particular event by comparing prices across a given span of time. For example, if a presidential candidate wanted to know the impact of his debate performance on his chances of winning the election, the candidate could simply compare the price of a security in the vote-share market prior to the debate to the price of the same security following the debate.

dictions without fear that the act of making the decision will alter the market's predictive capabilities.

The second type of decision amenable to market prediction includes those that affect the prediction, but not the definition of, the underlying bet. For example, a presidential candidate might elect to change campaign tactics based on a vote-share market's prediction. A subsequent decision to increase political advertising or appeal to a special interest group will affect the market's prediction about the candidate's success in the upcoming election. The decision will not affect the underlying bet, however, which is still the percentage of the vote the candidate will garner. As with the first type of decision, an individual may use the prediction directly to make a decision about future events.

The third type of decision affects both the prediction and the definition of the underlying bet because the decision prevents occurrence of the event on which the bet depends. Consider, for example, a government agency that establishes a market in May to predict the number of passengers airlines would carry in June if the agency were to implement a particular policy. The agency's goal is to use the market's predictions about the policy's success to decide whether to implement it before the country experiences any of its effects. Suppose that after a week of trading, the agency finds the market prediction about the policy unfavorable and consequently decides not to implement it. If the agency were to reject the policy in favor of a different policy and let trading continue, it would alter the definition of the underlying bet, because traders anticipating the agency's change in policy ex ante would not bet based on the predicted effects of the original policy. Instead, these traders' bets would reflect the belief that airline passenger volume would decline under any policy the agency adopted. In such an instance, the event anticipated by traders was the success of any policy, not the success of this particular policy. Thus, market price throughout the trading period would predict that the agency would fail to attract passengers under any policy, and not that the original policy, in particular, would fail to entice passengers. The agency therefore would have decided to reject the policy based on the wrong prediction.

To avoid redefining the bet, the agency could stop trading after it changes its policy. If traders expect that the agency might stop trading in the event of unfavorable predictions, there would be no effect on the bet's definition; at all times, the bet would reflect the predicted effects of the original policy. Moreover, the possibility that the agency might stop trading prior to the event is unlikely to substantially reduce the quality of the market prediction. Consider other betting environments where there is a possibility that a bet may not be paid at all, such as a baseball game that is likely to be rained out. In these situations, predictions about the event are still good because bettors face the risk of losing money if the event occurs. A bettor who knows that a game only has a fifty-percent chance of occurring still must make a sincere bet based on all available information or else suffer increased risk of financial loss in the event that the game occurs. Of course, as the likelihood of the event approaches zero, so does the threat of losing, and hence at some point the market's prediction cannot be trusted. Here though, assuming the likelihood that the agency will implement the policy is non-negligible, it is reasonable that the market's predictions are trustworthy.

While halting trading solves the problem of redefining the underlying bet, such a move creates another problem.[74] How much money should security holders receive if there is no observable event on which to base a payout? There is no clear answer.[75] The two obvious solutions, paying the security holders a fixed amount or paying them the market price at the time trading is stopped, are not satisfactory. In either case, the amount returned would distort the price, which would no longer measure the effect of the original policy, but rather would represent an expected value, equal to the number of passengers under the original policy multiplied by the probability of its acceptance, plus the amount returned on the security multiplied by the probability that the agency will change its policy.

[74] By stopping the market and returning traders' money, the market sponsor wastes traders' time and effort. Such a market may have trouble attracting traders.

[75] One solution not explored in this Note is the possibility of undoing all of the transactions undertaken with regard to the securities in question once the agency decides not to implement the policy. The main drawback to this solution arises from the amount of record keeping it entails. Every single transaction between two security holders from the date of issuance until the moment trading stops must be recorded. Then, after trading stops, each transaction must be undone, clearly a non-trivial task if parties to a transaction have already exchanged money.

If the agency set the amount returned at a fixed number, it would act like a magnet pulling the security's price toward it and away from what it should be measuring—namely, the number of passengers under the original policy. For example, suppose that the market believes the probability that the agency will adopt the policy is fifty percent and that the number of passengers under the policy would be one hundred, and the security pays one dollar for every passenger. Also, suppose that if the agency does not implement the policy it will return ten dollars for each security. Ideally, the security should trade at one hundred dollars, reflecting the number of passengers the market predicts will fly under the policy. But because of the fixed return in the event that the policy is not implemented, the security will trade at fifty-five dollars.[76] The fixed payout distorts the prediction by pulling price toward the fixed amount and away from a price representing the number of passengers.

Problems would also ensue if the market price at the end of trading were equal to the amount returned. As noted above, as the probability of an event occurring diminishes, so does the threat of losing, and thus the trustworthiness of the market's predictions declines. If two traders knew the agency would pay an amount equal to the market price, and knew that the event had little chance of occurring, they could cooperate to profit at the agency's expense. A trader who owned a security could sell it to a second trader for an exorbitant amount. As a result of this trade, the market price would rise, leading the agency to pay the second trader the new market price and allowing that trader to recoup his costs. The first trader could then split the profit made by selling the security to the second trader. In the extreme case, the market price would rise to the highest amount that the agency would be willing to pay. As a result, the security would trade for an amount exceeding the price that corresponded to the actual value of the proposed policy. Thus, as with the fixed payoff, the meaning of the price would be changed if the agency committed to paying an amount equal to the

[76] This result is reached by multiplying the prediction of 100 passengers by the 50% likelihood that the policy will be adopted, plus the $10 payout if the policy is not adopted multiplied by the 50% likelihood of that occurring. That is, $(100 \times 50\%) + (10 \times 50\%) = 55$.

market price when it stopped trading before the realization of the policy's effects.

In the above example, the agency's objective was to use the market's prediction about a policy's effect on passenger volume to decide whether to implement the policy. Unfortunately, if the agency uses the market's prediction as grounds for rejecting or changing the policy, it alters the underlying bet. If traders come to expect such a policy change, then the prediction becomes unreliable. Alternatively, if the agency stops trading upon its decision not to implement the policy that underlies the market bet, the market's prediction becomes reliable, but a new problem emerges: how to pay security holders. Two potential solutions, awarding security holders either a fixed sum or the current market price, are not viable because they alter the definition of the bet.

To ensure that the definition of the underlying bet is not altered and to compensate holders of securities whose payoffs depend on the effect of a non-implemented policy, this Note offers two solutions. First, to ensure that the bet's definition is not altered, this Note proposes to pair the original bundle, whose securities provide a return based on a policy which the agency may or may not implement, with a complementary bundle containing securities that pay holders based on the effects of any policy other than the one contemplated by the original bundle. The complementary bundle acts as a catch-all. Once the agency takes an action that alters the definition of the bet underlying the securities in the original bundle, the securities in the complementary bundle will provide payoffs based on the effects of the policy actually implemented. Second, to compensate the holders of securities in the original bundle after the agency decides not to implement the original policy, this Note proposes rules to prevent traders from manipulating the market prices of these securities, thereby allowing the market sponsor to return an amount equal to the securities' market price.

B. A Solution to Make a Policy Market Viable

Businessmen and policymakers frequently request input about proposed policies and then must decide whether to undertake those policies. Markets can be designed to help these individuals make their decisions. The following market model resolves the problems that arose in the above example when the agency sought to use the market's prediction of the policy's effect in its decision whether to adopt a policy. In particular, the following model ensures that an agency's decision to reject a policy does not alter the definition of the underlying bet and that the holders of these securities receive appropriate compensation once an agency decides not to implement the policy.

Assume that maximum airline capacity for the month of June is one hundred. Instead of offering one security or one bundle, the agency offers two bundles of Arrow-Debreu securities with two securities in each bundle. Each bundle costs one dollar, and the agency sells the bundles together for two dollars.[77] The first bundle of Arrow-Debreu securities pays its owners if the agency adheres to the policy. Security 1 pays one cent for each passenger carried, while Security 2 pays one cent for each empty seat. If the agency abandons the policy, holders of Security 1 and Security 2 each receive a sum, described in detail below, from the agency, assuming that the agency is the market sponsor. The second bundle of Arrow-Debreu securities pays its holders if the agency changes its policy. Security 3 pays one cent for each passenger carried under a changed policy, while Security 4 pays one cent for each empty seat under a changed policy. If the agency adopts the policy, it will pay holders of Security 3 and Security 4 a sum described in detail below.

The second set of Arrow-Debreu securities has a catch-all quality. Unlike the first set of securities, which pays off only if the agency retains the original policy, the second set does not require the agency to implement a particular policy. Once the agency changes its policy, the second set of securities pays its holders, regardless of any other policy change made by the agency. Thus, the bet encompasses any actions taken by the agency, just as a bet on an election encompasses any action taken by a candidate. The event on which the bet hinges does not require the agency to retain a particular policy, nor is it contingent on any other condition.

[77] Each bundle of Arrow-Debreu securities accounts for all of the possible outcomes associated with a particular policy. The first bundle corresponds to outcomes under the original policy, while the second bundle corresponds to outcomes under a changed policy. By packaging the bundles together, the agency offers a second-order bundle of Arrow-Debreu securities that accounts for all possible outcomes under any policy.

To understand the operation of these securities in agency decisionmaking, imagine two hypothetical scenarios: one in which the agency adopts the policy and another in which the agency rejects the initial policy and adopts a changed version. In the first scenario, suppose that midway through May, the market predicts that the policy will bring in eighty passengers and that a changed policy would bring in sixty. Based on this prediction, the agency will adopt the proposed policy. At the end of June, Securities 1 and 2 would pay out based on the number of passengers carried and the number of empty seats, respectively. Securities 3 and 4 would continue to be traded through the end of June because it is possible that the agency could amend its policy at any time. At the end of the month, the agency would return some amount of money to the holders of Securities 3 and 4. In the second scenario, suppose that midway through May, the market predicts that the proposed policy will bring in sixty passengers and that a changed policy would bring in eighty. As a result, the agency decides to change its policy. The agency immediately stops trading Securities 1 and 2 and returns some amount of money to those traders. Trading in the second set of securities continues. At the end of June, the agency pays out the appropriate amount to holders of Securities 3 and 4 based on the numbers of passengers and empty seats.

A potential objection to this market is that it allows traders to increase the price of Security 1 or Security 3 in order to encourage the agency to adopt a particular policy. For example, an airline that preferred the original policy might continuously buy Security 1 to keep the price high. Although theoretically possible, such a situation remains unlikely in practice. First, the offending airline would have to acquire the resources to necessary to keep the market price artificially high. Second, the airline would suffer a large financial loss because the amount it paid for Security 1 would not reflect the number of passengers airlines would carry under the policy, and thus the airline could not actually capture the benefits in the form of passenger volume.

There still remains, however, the difficulty of how the market sponsor will return money owed to traders who hold securities contingent on an event that does not occur. As discussed above, returning either the market price or a fixed amount can distort the price itself. The solution offered by this Note is that the security holders must receive the original price of the bundle minus the price of the other security within that bundle. For example, suppose that the agency adopts the policy, thereby eliminating any event for which Securities 3 and 4 might pay holders. In response, the holders of Security 3 should receive one dollar, corresponding to the original price of the bundle, minus the market price of Security 4. Likewise, the holders of Security 4 should receive one dollar minus the market price of Security 3. This solution is based upon the intuition that traders should receive the market price for their security, but they should have no ability to raise the price artificially and thereby manipulate the market.

To implement this solution, first it is necessary to ensure that the proposal does not change the price itself. Notice that, similar to the proposals to return a fixed price or the market price, the definition of price changes to an expected value. The price of Security 3, for example, adjusts to

$$(p \times (1 - S4)) + ((1 - p) \times (X \div 100)),$$

where p is the probability that the agency adopts the proposed policy, $S4$ is the price of Security 4, X is equal to the number of passengers under the changed policy, and division by one hundred ensures that the security pays one cent rather than one dollar per passenger. The first term of this formula represents the payoff for Security 3 if the agency selects the proposed policy, while the second term represents the pay-off under a changed policy.

Added together, the probability that the agency adheres to the original policy and the probability that the agency changes the policy must equal one because there are no other possibilities. Also, notice that one dollar minus the price of Security 4 equals the price of Security 3, which measures the number of passengers carried under the changed policy. The formula for the price of Security 3 thus can be rewritten as

$$(p \times X) + ((1 - p) \times (X \div 100)).$$

Further simplification of this formula demonstrates that the price of Security 3 depends on the number of passengers carried under the changed policy. In other words, Security 3 measures exactly what it is designed to measure.

Second, it is critical to ensure that this proposal does not encourage strategic behavior. Suppose that a trader holds Security 3, it is late in June, and there is almost no chance that the agency will reject the policy in favor of another policy. A trader therefore would be quite certain that the agency would return money. The trader knows that Security 3 will pay one dollar minus the price of Security 4, and because the agency will not change fares, the threat of losing money from a bad bet is low. Thus, if the price of Security 3 no longer predicts reality, the trader would not mind. The question is whether such a trader can take strategic action to increase the money received when the agency ultimately adopts the policy.

Fortunately, this trader cannot take strategic advantage of the market. For the trader to receive a greater amount of money for Security 3, the price of Security 4 must decrease because Security 3's payoff is one dollar minus the price of Security 4 when the probability that the agency adopts the policy is 100%. But there is no clear way for a holder of Security 3 to lower the price of Security 4. Holders of Security 4 would like to see the price rise and they would have no incentive to sell their securities for less than they were worth when the threat of a change in agency policy was higher. The only strategic behavior available to the trader would be to artificially bid up the price of Security 3, resulting in a corresponding decrease in Security 4's price. It is not clear, however, why a Security 3 holder would bid up its price when such action provides no benefit and only harms a Security 4 holder.

Instead, it seems more reasonable that, as the possibility of the agency changing its policy diminishes, Security 3 holders will lose interest in the market and trading will slow, perhaps even to a stop. In the event that Security 3 holders did bid up the price, the market sponsor easily could recognize and stop such behavior by monitoring the prices of both securities. If Security 3 holders bid up the price of Security 3, Security 4 holders would likely retaliate and bid up Security 4. As a result of this bidding warfare, the prices of both securities would rise, no longer adding to $1. Consequently, a market sponsor could easily determine that traders were acting strategically, and the sponsor then could undo the strategic trades that caused the combined price of Securities 3 and 4 to rise above a threshold, say $1.05.

The main drawback to this Note's proposed market lies in its complexity. Offering four securities instead of one is not a particularly elegant solution and the proposed scheme for returning money to shareholders is difficult to explain. The technical details of the scheme, however, are not important to a trader, who only must predict correctly a policy's effect and then have faith that the scheme prevents opportunistic behavior.[78] The proposed market is also beneficial in that it extends easily. For example, imagine that an agency sought to construct a market to predict future states of the world under different policies in order to select the optimal policy. If the agency were concerned primarily with the effects of two different policies, it might offer three bundles of Arrow-Debreu securities: Bundle 1 would pay based on the effects of Policy *A*; Bundle 2 would pay based on the effects on Policy *B*; and Bundle 3 would pay based on the effects of a policy that differed from both Policies *A* and *B*. Of course, the quality of the market's predictions still would depend on market efficiency and traders' knowledge about the subject. Yet, when the success of the policies depends on the individual actions of many people, the predictions are likely to be accurate. As with markets generally, however, in specialized areas where the knowledge needed to make a prediction rests in the hands of one or two people, it might be more cost-effective to ask those people directly.

CONCLUSION

The basic aim of this Note has been to introduce the concept of information markets to discussions concerning administrative decisionmaking, particularly rulemaking. Information markets constitute an intriguing tool for administrative agencies to evaluate the uncertain effects of various potential policies. This Note has argued for the use of information markets in the administrative context

[78] See Robert W. Hahn & Paul C. Tetlock, Using Information Markets to Improve Public Decision Making, 29 Harv. J.L. & Pub. Pol'y 214, 218, 225–36 (2005) (discussing a market with similar characteristics designed to improve the way the government awards contracts); see also Robin Hanson, Decision Markets, 14 IEEE Intelligent Sys. 16, 16–18 (1999), at http://hanson.gmu.edu/decisionmarkets.pdf (proposing a market to help policymakers select among competing policies).

based on both empirical and theoretical grounds, and most importantly, has explained the primary features of a market suitable for evaluating policies prior to their implementation. Questions concerning the incorporation of information markets into the existing administrative law regime, founded in the Administrative Procedure Act, are left for later discussion and articulation.

Decisionmakers can and should use the predictions of information markets to make administrative decisions. Both empirical research and economic theory suggest that information markets provide predictions that equal or outperform predictions made by groups, voting, experts, and opinion averaging. In cases in which the decision does not affect the definition of the underlying bet, market predictions may be used directly to make decisions. Complications emerge when decisions are made that alter the underlying bet, such as an agency decision to undertake a policy that prevents the occurrence of the event on which the bet depends. When the agency changes its policy based on the market's prediction, trading in the security contingent on the effects of the policy may either continue or stop. Should trading continue, it remains unclear whether traders are betting on the effects of the policy in question or the effects of any policy the agency may implement. If trading stops, it is equally unclear how the bets should pay security holders in the absence of an observable event. Ultimately, this Note has presented a market that allows an agency to change a policy without altering the underlying bet, while simultaneously providing a way to return money to security holders even if the event on which their bet is based does not occur. In this way, decisionmakers gain a powerful tool to evaluate competing policies of uncertain effect prior to implementation.