



GW Law Faculty Publications & Other Works

Faculty Scholarship

2004

Information Markets, Administrative Decisionmaking, and Predictive Cost-Benefit Analysis

Michael B. Abramowicz

George Washington University Law School, abramowicz@law.gwu.edu

Follow this and additional works at: http://scholarship.law.gwu.edu/faculty_publications

 Part of the [Law Commons](#)

Recommended Citation

Michael Abramowicz, Information Markets, Administrative Decisionmaking, and Predictive Cost-Benefit Analysis, 71 U. Chi. L. Rev. 933 (2004).

This Article is brought to you for free and open access by the Faculty Scholarship at Scholarly Commons. It has been accepted for inclusion in GW Law Faculty Publications & Other Works by an authorized administrator of Scholarly Commons. For more information, please contact spagel@law.gwu.edu.

Information Markets, Administrative Decisionmaking, and Predictive Cost-Benefit Analysis

by Michael Abramowicz*

FutureMAP, a project of the Defense Advanced Research Projects Agency, was to involve experiments to determine whether information markets could improve Defense Department decisionmaking. Information markets are securities markets used to derive information from the prices of securities whose liquidation values are contingent on future events. The government intended to use such a market to assess the probabilities of potential political assassinations, and the indelicacy of this potential application contributed to a controversy leading to the cancellation of the program. In this Article, Professor Abramowicz assesses whether information markets in theory could be useful to administrative agencies, and it concludes that information markets could help discipline administrative agency predictions, but only if a number of technical hurdles such as the danger of manipulation can be overcome. Because the predictions of well-functioning information markets are objective, they function as a tool that exhibits many of the same virtues in predictive tasks that cost-benefit analysis offers for normative policy evaluation. Both approaches can help to overcome cognitive errors, thwart interest group manipulation, and discipline administrative agency decisionmaking. The Article suggests that the two forms of analysis might be combined to produce a “predictive cost-benefit analysis.” In such an analysis, an information market would predict the outcome of a retrospective cost-benefit analysis, to be conducted some years after the decision whether to enact a particular policy. As long as the identity of the eventual decisionmaker cannot be anticipated, predictive cost-benefit analysis estimates how an average decisionmaker would be expected to evaluate the policy. Because the predictive cost-benefit analysis assessment is not dependent on the identity of current agency officials, they cannot shade the numbers to justify policies that the officials prefer for idiosyncratic or ideological reasons.

I. THE ECONOMICS OF INFORMATION MARKETS	11
A. The Mechanics	12
B. Accuracy	18
C. Market Design Challenges	21
1. Conditional Markets.....	21
2. Thin Markets.....	26
3. Subsidized Markets.....	28
II. INFORMATION MARKETS AND ADMINISTRATIVE DECISIONMAKING	30

* Associate Professor of Law, George Mason University. J.D., Yale Law School; B.A., Amherst College. For comments and suggestions, I thank Terrence Chorvat, Emily Frye, Robin Hanson, Shi-Ling Hsu, Stephen Parr, Dan Polsby, and Steven Schooner. Anthony Ellis provided outstanding research assistance.

A. The Objectivity of Information Markets.....	31
1. Why Objectivity Matters.....	31
2. Are Information Markets Objective?	39
B. Some Potential Governmental Uses for Information Markets.....	51
1. Homeland Security Vulnerabilities.....	52
2. Solvency Regulation	56
3. Budget Forecasts.....	58
4. Interim Regulation	60
C. Institutionalizing Information Markets	62
III. PREDICTIVE COST-BENEFIT ANALYSIS AND BEYOND	66
A. The Mechanics	68
B. Comparisons	72
1. Normative vs. Positive Markets.....	72
2. Predictive vs. Traditional Cost-Benefit Analysis.....	77
C. Comparative Benefit Analysis	86
IV. CONCLUSION.....	87

INTRODUCTION

Administrative law innovations are not ordinarily announced on the front page of the *New York Times*. A Defense Department experiment in government decisionmaking, however, made the front page of the *Times* two days in a row. The first article announced the plan for the program,¹ and the second announced that officials had responded to criticisms of it by withdrawing it even before it was to be launched later that week.² The program, sponsored by the Defense Advanced Research Projects Agency, was to have been called FutureMAP, and its goal was to develop “market-based techniques for avoiding surprise and predicting future events.”³ The immediate application that generated controversy was the use of an information market to predict developments in the Middle East, including terrorist acts. One senator denounced the program as a “federal betting parlor on atrocities.”⁴ Participants in information markets in effect place bets on future events, so the description is not an inaccurate one.

¹ See Carl Hulse, *Pentagon Prepares a Futures Market on Terror Attacks*, N.Y. TIMES, July 28, 2003, at A1.

² See Carl Hulse, *Swiftly, Plan for Terrorism Futures Market Slips into Dustbin of Ideas Without a Future*, N.Y. TIMES, July 29, 2003, at A1 (A10 in later editions).

³ <http://www.darpa.mil/iao/FutureMAP.htm> (last visited Apr. 21, 2003) (web site no longer available). FutureMAP is an acronym for “Futures Markets Applied to Prediction.”

⁴ Ken Guggenheim, *Pentagon Says Threat-Bet Program To Be Canceled*, AP, July 29, 2003 (quoting Sen. Ron Wyden).

The firestorm over FutureMAP mostly reflected its implications for anti-terrorism efforts. Concerns included that FutureMAP actually might encourage terrorism,⁵ and politicians reasonably worried that federally sponsored wagering on whether foreign leaders might be assassinated could damage foreign relations.⁶ Yet there was virtually no discussion in the media of the relevance of information markets themselves for governmental decisionmaking.⁷ Policymaking by governmental agencies often depends, either implicitly or explicitly, on predictions about the future. Economic policy depends on anticipation of economic conditions. Decisions whether to build prisons depend on projections of the number of prisoners. Highway appropriations depend on anticipated traffic. And on and on. Even where prediction is not central to determining whether there is a need for government action, an agency decision to enact any regulation presumably depends at least on a prediction that the regulation will have certain consequences. Predictions both about conditions that might affect government policy and about the consequences of such policy are thus central to regulation.

Despite the central of prediction to policymaking, there is no provision in the Administrative Procedure Act specifying how agencies should go about making predictions, and so predictions in effect are treated like any other agency findings.⁸ If FutureMAP had any promise, it was not so much as an anti-terrorism tool, but as a method for systematizing administrative agency predictions, a method that might be useful in other regulatory areas even if not for terrorism. The FutureMAP debacle has provided a setback, maybe a permanent one, to anyone who might have hoped to use information markets in administrative decisionmaking. The academic question remains, however, of whether information markets would be a useful tool for government agencies to employ. That is the question this Article considers, and the answer it provides is mixed. Information markets could help constrain administrative decisionmaking and

⁵ *Id.* (quoting Sen. Thomas Daschle as providing “an incentive actually to commit acts of terrorism”).

⁶ Hulse, *supra* note 1 (quoting Sen. Byron Dorgan as rhetorically asking, “Can you imagine if another country set up a betting parlor so that people could go in . . . and bet on the assassination of an American political figure?”).

⁷ For two exceptions, see Daniel Gross, *Book-Makers for the Bomb-Makers*, SLATE, July 29, 2003, available at <http://slate.msn.com/id/2086315/>; and Jeremy Kahn, *Is A Futures Market on Terror Outlandish?*, FORTUNE, July 30, 2003, available at <http://www.fortune.com/fortune/investing/articles/0,15114,471785,00.html>. These articles, however, focus on the mechanics of information markets, not on whether information markets might be useful in government decisionmaking generally.

⁸ There is some uncertainty in the case law about whether predictions should be treated as fact findings, and thus subject to the substantial evidence test, or as policy decisions, and thus subject to hard look review. Compare, e.g., *American Textile Mfrs. Institute, Inc. v. Donovan*, 452 U.S. 490 (1981) (Stewart, J., dissenting) (“OSHA’s prediction of the cost of the Cotton Dust Standard lacks a basis in substantial evidence.”), with *Marsh v. Oregon Natural Resources Council*, 490 U.S. 360, 373-74 (1989) (noting that the agency is required to take a “hard look” at the environmental effects of its action).

limit ideological decisionmaking, but their usefulness depends on some difficult empirical questions that the economics literature on information markets has not yet adequately answered.

The phrase “information market” evokes the mechanics of the approach and highlights the intuition underlying it. An information market, as traditionally constructed, is a stock market created for the purpose of extrapolating information from share prices. The securities in such a market do not serve as claims to corporate ownership, but rather offer payoffs contingent on some future contingency specified by the market’s sponsor. A security in an antiterrorism program, for example, might be issued with the proviso that its face value will be paid if and only if a cyberterrorism attack succeeds in shutting down the New York Stock Exchange in the next year. By examining the price at which the security is traded, the Defense Department would be able to obtain an estimate of this risk.⁹ Although some firms might use such a market as a means of hedging risk,¹⁰ its primary purpose is to harness the power of securities markets to aggregate information. An information market is potentially useful whenever an agency decision depends in part on information about the future and a security can be constructed whose price would provide a relevant prediction.

The word “market,” however, can be misleading. Information markets need not depend on either trading or the issuance of securities at all, and indeed this Article will endorse an approach that does not rely on buying or selling.¹¹ Broadly conceived, an information market is any device that gives third parties financial incentives to make predictions or to improve upon the others’ predictions and that combines the predictions into a single consensus value. The literature suggesting that conventional stock markets exhibit at least a weak form of market efficiency¹² does not itself guarantee comparable performance levels in all information markets,

⁹ Presumably, such information would have been useful so that the Department of Defense might take action to prevent an attack. The example, however, immediately presents an obvious problem with this particular implementation of information markets. A trader with information that a cyberterrorism attack is likely would have no incentive to trade on the information if such trading would succeed in thwarting the attack. *See infra* Part II.B.1.

¹⁰ The securities in information markets are derivatives, but because securities will exist only for a small subset of all risks, the introduction of these derivatives may not necessarily enhance welfare, placing aside the use of these securities for informational purposes. *See* Peter H. Huang, *Securities Price Risks and Financial Derivative Markets*, 21 NW. J. INT’L L. & BUS. 589 (2001) (explaining why the introduction of a new derivative market has indeterminate consequences for consumer and investor welfare).

¹¹ *See infra* Part I.C.2 (discussing the “market scoring rule”).

¹² For a discussion of the three possible levels of market efficiency, see Eugene F. Fama, *Efficient Capital Markets: A Review of Theory and Empirical Work*, 25 J. FIN. 383, 383 (1970). Some finance scholars have found deviations from semi-strong form efficiency. *See* Ronald J. Gilson & Reinier H. Kraakman, *The Mechanisms of Market Efficiency*, 70 VA. L. REV. 549, 626 n.205 (1984). There has long been a consensus, however, that markets exhibit at least weak-form efficiency, meaning that future price movements cannot be predicted solely on the basis of past prices. *See* Jonathan R. Macey & Geoffrey P. Miller, *Good Finance, Bad Economics: An Analysis of the Fraud-on-the-Market Theory*, 42 STAN. L. REV. 1059, 1081 (1990) (noting “overwhelming

and similarly the behavioral finance literature identifying irrationalities in aggregate stock market investment is of little applicability, at least to the information market form endorsed here.¹³ A small economics literature on information markets, however, has offered preliminary assessments of the accuracy of information markets,¹⁴ considering such existing information markets as the Iowa Electronic Markets,¹⁵ which primarily allow for trading on the elections' outcomes. These studies indicate that information markets are generally superior to other forecasting tools, such as polls, because information markets aggregate a variety of types of information and a range of individuals' predictions.¹⁶

There has been no legal scholarship on information markets' potential uses to aggregate information for administrative decisionmaking.¹⁷ This aggregation function alone could make information markets a modestly useful tool in administrative decisionmaking, perhaps providing a small advantage over relying on a single expert to combine various information sources. There is, however, an additional attribute of information markets, receiving little attention in the economics literature,¹⁸ that is of far greater significance for governmental decisionmaking.

empirical support for weak-form efficiency *See generally* Michael L. Wachter, *Takeover Defense When Financial Markets Are (Only) Relatively Efficient*, 151 U. PA. L. REV. 787, 801-04 (2003) (providing a recent review of the market efficiency literature).

¹³ For a review of the literature, see Andrei Shleifer, *INEFFICIENT MARKETS: AN INTRODUCTION TO BEHAVIORAL FINANCE* (2000). Behavioral finance scholars worry less about the relative valuations of securities than about aggregate investment. *See, e.g.*, Jeeman Jung & Robert Shiller, *ONE SIMPLE TEST OF SAMUELSON'S DICTUM FOR THE U.S. STOCK MARKET* (NBER Working Paper No. 9348, Nov. 2002) (providing a test confirming the prediction that capital markets are more efficient for individual stocks than for the aggregate stock market). Paul Samuelson made this point as follows:

Modern markets show considerable *micro* efficiency (for the reason that the minority who spot aberrations from micro efficiency can make money from those occurrences and, in doing so, they tend to wipe out any persistent inefficiencies). In no contradiction to the previous sentence, I had hypothesized considerable *macro* inefficiency, in the sense of long waves in the time series of aggregate indexes of security prices below and above various definitions of fundamental values.

Letter from Paul Samuelson to John Campbell and Robert Shiller, quoted in ROBERT J. SHILLER, *IRRATIONAL EXUBERANCE* 243 (2d ed. 2001). Aggregate investment is not a concern with information markets constructed in the manner that this Article recommends. Participants' sole incentive with this approach is to offer the best predictions possible, so while increased participation in a market should enhance accuracy, unlike conventional markets, there are no securities whose prices will be bid up as a result of increased investment.

¹⁴ *See infra* Part I.B. The literature is small, but the Defense Department's program has triggered some interest among economists, many of whom gave presentations at a DARPA-sponsored June 2002 workshop on information markets. *See* <http://marteksys.com/martek/DARPAConference.html> (last visited June 17, 2003) (conference announcement).

¹⁵ <http://www.biz.uiowa.edu/iem/> (last visited June 17, 2003).

¹⁶ For an article describing how information markets aggregate information, see David M. Pennock & Michael P. Wellman, *A Market Framework for Pooling Opinions* (2001) (unpublished manuscript, available at <http://citeseer.nj.nec.com/399182.html>).

¹⁷ FutureMAP has received far less media attention than the Information Awareness Office's data mining program. Media references include *DARPA Continues Efforts to Develop Future-Forecasting Markets*, *INSIDE AIR FORCE*, Jan. 6, 2003, available at 2003 WL 7601951; Shane Kite, *Project Seeks Terror Clues in Marts*, *SEC. INDUS. NEWS*, Mar. 17, 2003, available at 2003 WL 7547217; and James Surowiecki, *Decisions, Decisions*, *NEW YORKER*, Mar. 24, 2003 (mentioning FutureMAP in a brief article on information markets).

¹⁸ At least one commentator has mentioned potential usefulness in controlling bias. *See* George R. Neumann, *Using Markets to Make Decisions* at slide 2 (2002) (unpublished presentation, on file with author) (noting the possibility of "'yes'-man effect in organizations" or "institutional biases"). also underlie one of the first papers to consider information markets, which considered

Information markets provide *objective* predictions, though only if concerns such as the possibility of market manipulation can be overcome. Even if information markets offered little in the way of improved accuracy relative to a single, well-motivated decisionmaker's prediction, the objectivity of well-functioning information markets makes them a promising potential administrative decisionmaking tool. Agencies that regularly utilized information markets might limit the influences of interest group pressure, availability cascades,¹⁹ and related pathologies of bureaucratic life. Information markets thus help discipline agency predictions in much the same way as cost-benefit analysis disciplines normative agency decisionmaking.²⁰ The results of information markets could be useful, though not dispositive, in judicial review of agency action, either as an agency uses the information market's result to defend a challenged decision or as an opponent of the agency action challenges the decision in court.

The objectivity of information markets, assuming that they could be made to function as proposed, might appear to be both their greatest limitation as well as their greatest virtue in governmental decisionmaking. The most difficult challenge for administrative agencies, after all, is generally not the prediction of the future, but the evaluation of what should be done given such predictions. As described in the existing literature, information markets serve to provide data inputs into policy analysis, not to resolve a normative inquiry that may depend on a number of variables. Information markets emerge in the literature as predictive tools analyzing objectively verifiable facts about the world.²¹ The identity of the decisionmaker who announces the final value from which traders' profits are calculated thus has little or no relevance. An information

possible use as an alternative to peer review in science. See Robin D. Hanson, *Could Gambling Save Science? Encouraging an Honest Consensus*, 9 SOC. EPISTEMOLOGY 3 (1995). No commentator to my knowledge, however, has considered the potential for decision markets to overcome the problems frequently identified with governmental decisionmaking.

¹⁹ See Timur Kuran & Cass R. Sunstein, *Availability Cascades and Risk Regulation*, 51 STAN. L. REV. 683, 683 (1999) (defining an "availability cascade" as "a self-reinforcing process of collective belief formation by which an expressed perception triggers a chain reaction that gives the perception increasing plausibility through its rising availability in public discourse").

²⁰ See *infra* Part II.A.1.

²¹ Actual markets recognize the possibility that there might be debate about the value of a number being predicted. See, e.g., <http://www.ideosphere.com/fx-bin/Claim?claim=Clone> (last visited June 17, 2003) (recognizing the possibility of ambiguity in a claim being used to predict the possibility that a human clone will exist by 2005). I have found no existing information market that attempts to predict or largely subjective decision. Nor have I found any articles suggesting that information markets be used in this way. To the contrary, the literature on information markets seems to emphasize that the variable being predicted be capable of objective measurement. Robin Hanson, for example, discusses an information market assessing a criminal justice policy, "[Y]ou must state your claim clearly.... You might decide to focus on your state's murder rate, using some standard government statistic M as your official measure of it." Robin Hanson, *Decision Markets*, IEEE INTELLIGENT SYS., May/June 1999, at 16, 17. In a discussion, however, Hanson readily acknowledged that as long as there was a clear means of identifying the eventual decision, a market could evaluate the eventual subjective decision and that any bias in decision would not affect the market prediction to the extent that it was known *ex ante*.

market, however, could be easily used to make predictions about future normative assessments of legal policy, even if these normative assessments are wholly subjective. There must be some procedure for determining who will make the future subjective assessment, so in a sense any normative assessment that resolves an information market is in itself an objectively verifiable fact. Such a market, however, would predict what someone will say about the relevant issue rather than directly predicting some number on which all decisionmakers should be expected, more or less, to agree.

This Article's principal contribution is to defend the proposition that markets of this type, which I will call *normative information markets*, may be particularly useful, and indeed that they have significant advantages in legal decisionmaking over information market types that the literature previously envisioned, which I will label *positive information markets*. The literature's failure to discuss normative information markets is understandable. Opinions, after all, are inherently subjective, and it might appear that there is no point in making objective predictions of subjective assessments. This response, however, ignores that whether policies are ultimately enacted depends on subjective assessments of whether those policies are desirable, and normative information markets can reveal how decisionmakers on average, rather than just particular agency bureaucrats, would assess policies. While positive information markets can provide data that may serve as useful inputs into a policy decision, a normative information market can provide a bottom-line assessment of the policy itself. Normative information markets are thus easier to interpret and avoid some of the technical problems that may beset positive information markets.²²

For an example of how a prediction of a subjective assessment could be useful, imagine a very primitive normative information market used to predict whether ten years from now an agency official will conclude that a set of safety regulations being considered today are "good" or "bad." If each market security would pay \$1 if the final evaluation is that the regulations are judged as good, then a trading price of thirty cents would reflect an objective market prediction that there is a thirty-percent chance that the eventual decisionmaker will also conclude that the decision was good. Such a prediction, if made before the decision,²³ would serve as a reality

²² See *infra* Part III.B.1.

²³ A complication here is how to liquidate the market if the agency in fact decides not to make the decision. There are at least two possibilities. First, the rules might provide that all investments will be refunded at cost if the decision is not made. This approach

check that might help avoid decisions that seem superficially attractive but that on closer analysis seem likely to prove misguided over time. At the same time, this information market type might assist in exposing an agency, or agency decisionmaker, acting on impermissible or socially inefficient motivations,²⁴ in a way that a market reporting only one input into an agency decision cannot, given the agency's ability in many cases to justify decisions on the basis of other, not easily measured factors. Of equal significance, a market prediction indicating that there is a seventy-percent possibility of a "good" assessment might help an agency to defend farsighted actions against superficial attacks.

The prediction of a normative information market, however, may not be all that meaningful if the decisionmaker who will perform the eventual policy assessment is known in advance. The market would be considerably more useful if the eventual decisionmaker were unknown, because the market prediction would then reflect an expectation of what an *average decisionmaker* would decide. An information market that predicts a future retrospective evaluation of a decision allows for a normative policy analysis that is not influenced by the party affiliation or ideology of current agency decisionmakers. Thus, even if the eventual evaluation is partly or completely subjective, an information market can furnish a relatively objective datum on how the average person would be expected to make that subjective, retrospective assessment. Although a hypothetical average decisionmaker is not identical to the heuristic median voter commonly identified in political science literature,²⁵ the concepts are similar. An agency's announced intent to enact regulations does not, by itself, indicate whether policymakers would generally share the agency's assessment of the issue or whether the agency is driven by idiosyncratic factors. A normative information market can provide an objective assessment of whether policymakers in general would reach the same conclusion as the agency.

leads to a potential selection bias problem, however. *See infra* paragraph accompanying note 75. Second, the rules might provide for a retrospective assessment regardless of whether the safety regulations are in fact enacted. *See infra* Part III.A (describing a similar approach).

²⁴ For example, an agency might be motivated by a desire to help a particular well-connected group or corporation but deny such a motivation. If traders assume that the future retrospective decisionmakers allegiance is unknown, then they may predict a negative evaluation if the factors that the agency uses to justify its policy decision do not themselves seem persuasive.

²⁵ Whether a political system should aspire to enshrine the preferences of the median voter is more complicated. As Robert Cooter notes, there may be a tradeoff between institutions that produce policies appealing to the median voter and those that allow bargaining to reflect intensity of preferences. *See* Robert Cooter, *Constitutional Consequentialism: Bargain Democracy Versus Median Democracy*, 3 THEORETICAL INQ. IN LAW 1 (2002). Because the primitive information market described here assesses whether a majority of decisionmakers would approve of regulations, it will tend to produce policies of which the median voter would approve, assuming that the decisionmakers are representative. Predictive cost-benefit analysis, described below, does not allow for explicit bargaining, but it does register the intensity of preferences.

The primitive normative market described so far, while producing useful data, is flawed, because it does not take into account predictions about either how good or bad a decision is likely to be. If, for example, there were a forty percent chance that a decision would have disastrous consequences and a sixty percent chance that a decision would have modest benefits, the information market described so far would produce an incomplete and misleading assessment. One way to take into account the magnitude of effects, as well as the intensity of preferences, would be to combine an information market with the tool that agencies often use to perform normative policy evaluations: cost-benefit analysis. This Article thus imagines *predictive cost-benefit analysis*, an information market used to predict a cost-benefit analysis that would be performed some time after a decision whether to enact a policy. Such a market would produce an objective prediction of an eventual subjective evaluation. The purpose of the eventual cost-benefit analysis that the information market predicts would not be to make a decision, but rather solely to discipline the cost-benefit analysis produced by the information market.

As long as the actual cost-benefit analysis is performed sufficiently in the future that traders would be unable to predict the political parties and ideologies of those who will perform it, predictive cost-benefit analysis is not vulnerable to the principal critiques made of the more familiar form of cost-benefit analysis. First, because information markets are objective, an agency cannot manipulate predictive cost-benefit analysis as an ad hoc rationale for decisions that the administrative agency would have made anyway.²⁶ Second, the cost-benefit analyses that the information market predicts could be relatively flexible. Some proponents of cost-benefit analysis have argued that such analysis must proceed according to relatively clear rules to prevent manipulation.²⁷ The problem with rigid cost-benefit analysis is that it necessarily makes value choices, whether by setting particular values for variables or by declaring certain types of

²⁶ See, e.g., Staff of Subcomm. on Oversight and Investigations of the House Comm. on Interstate and Foreign Commerce, 96th Cong., 2d Sess., *Cost-Benefit Analysis: Wonder Tool or Mirage?* 5 (Comm. Print 1980) (arguing that agency officials manipulate cost-benefit analyses to suit their policy preferences); Thomas O. McGarity, *Professor Sunstein's Fuzzy Math*, 90 GEO. L.J. 2341, 2366 (2002) (arguing that cost-benefit analysis entails “frequently preposterous and always manipulable number spinning”); Christopher H. Schroeder, *In the Regulation of Manmade Carcinogens, if Feasibility Analysis Is the Answer, What Is the Question?*, 88 MICH. L. REV. 1483, 1494 (1990) (“Cost-benefit analysis has been frequently faulted for requiring data that is, practically speaking, often unavailable and that, when available at all, is subject to manipulation by industry interests.”); Steve Bennett, Note, *Cost-Benefit Analysis and the Feasibility Requirement of the Occupational Noise Regulation*, 55 GEO. WASH. L. REV. 123, 146 (1986) (rejecting applicability of cost-benefit analysis to a particular regulation on manipulability grounds).

²⁷ See, e.g., Steve P. Calandrillo, *Responsible Regulation: A Sensible Cost-Benefit, Risk Versus Risk Approach to Federal Health and Safety Regulation*, 81 B.U. L. REV. 957, 998-1000 (2001) (urging use of objective data and providing an anecdotal example of how subjectivity may adversely affect risk calculations); see also *infra* notes 290-298 and accompanying text.

noneconomic considerations irrelevant.²⁸ The result is that such cost-benefit analysis can be justified only to those who agree with the rules governing it. If agencies can no longer use cost-benefit analysis to rationalize decisions, then manipulation is no longer a danger, and rigid rules are no longer necessary. Predictive cost-benefit analysis would reflect the values that average decisionmakers would set for variables about which there is substantial controversy, including discount rates,²⁹ the value of a statistical life,³⁰ and the benefit that citizens receive from the existence of species, clean air, and the like.³¹ Similarly, it would allow incorporation of effects, such as distributive consequences, that are typically ignored. In sum, predictive cost-benefit analysis overcomes the tradeoff generally associated with conventional cost-benefit analysis between the dual goals of constraining agencies and allowing cost-benefit analysis to be relatively inclusive.

This Article seeks both to identify information markets' potential technical problems with, and to consider information markets' potential usefulness assuming that these potential problems either do not exist or can be overcome. Part I offers a critical introduction to information markets. In addition to describing existing implementations and proposals for information market design, this Part considers information markets' accuracy. One of the most intriguing implementations of information markets is a conditional market, which is used to predict how a decision might affect some variable of interest. This Part, however, explains that the results of conditional markets must be cautiously interpreted, because variance attributable to slight imperfections in market design might appear to represent market predictions of the effect of the decisions at issue. The part also describes how problems associated with thin information markets can be overcome, and how a sponsor might subsidize an information market to improve participants' incentives to engage in research and analysis.

²⁸ Frank Ackerman & Lisa Heinzerling, *Pricing the Priceless: Cost-Benefit Analysis of Environmental Protection*, 150 U. PA. L. REV. 1553, 1576 (2002) ("Because value-laden premises permeate cost-benefit analysis, the claim that cost-benefit analysis offers an 'objective' way to make government decisions is simply bogus.").

²⁹ See, e.g., John J. Donohue III, *Why We Should Discount the Views of Those Who Discount Discounting*, 108 YALE L.J. 1901 (1999); Daniel A. Farber & Paul A. Hemmersbaugh, *The Shadow of the Future: Discount Rates, Later Generations, and the Environment*, 46 VAND. L. REV. 267 (1993); Lisa Heinzerling, *Discounting Our Future*, 34 LAND & WATER L. REV. 39 (1999); Richard L. Revesz, *Environmental Regulation, Cost-Benefit Analysis, and the Discounting of Human Lives*, 99 COLUM. L. REV. 941 (1999).

³⁰ For a recent work on improving valuations of statistical lives by the most influential commentator in this area, see W. Kip Viscusi, *The Value of Life: Estimates with Risks by Occupation and Industry* (Harvard John M. Olin Ctr. For Law, Econ., and Bus. Discussion Paper No. 422, May 2003).

³¹ See generally Donald J. Boudreaux et al., *Talk Is Cheap: The Existence Value Fallacy*, 29 ENVTL. L. 765 (1999) (providing an overview of methods for calculating existence value and arguing that the notion of existence value is conceptually flawed).

Part II discusses information markets' uses in administrative decisionmaking. The administrative markets' principal virtue is their objectivity, and this part begins by explaining how an objective prediction tool could help overcome various problems of administrative decisionmaking. It then presents both the affirmative case that information markets are relatively objective and two concerns regarding their objectivity. The concerns are that interested parties might manipulate them, and that the market participants' unrepresentativeness might affect market results. There are some theoretical reasons to think that these dangers are small, but attempts at manipulation might foster suspicions and decrease market accuracy. This Part describes how a novel two-phase information market design might alleviate this concern and encourage information sharing among market participants. With any market design, however, the impact of manipulation and unrepresentativeness are ultimately empirical, and this Part also explains how experiments might assess the extent of these problems. Part II then describes various possible uses of information markets—some modest, others less so—and evaluates how the various technical challenges in information market design might affect these uses.

Finally, Part III describes and defends predictive cost-benefit analysis. Because predictive cost-benefit analysis can result in a policy assessment regardless of the original decision, it does not suffer from selection bias and other technical problems that make conditional markets difficult to interpret. This part also explains that although predictive cost-benefit analysis itself may be more costly than traditional cost-benefit analysis, it offers significant benefits over the traditional approach. Predictive cost-benefit analysis is more objective than traditional cost-benefit analysis, and it thus provides a more reliable signal about the policies that an agency wishes to undertake. Indeed, this form of cost-benefit analysis potentially could satisfy both camps in various debates on regulatory theory. Part III also considers an alternative means of structuring information markets making predictions of subjective evaluations for situations in which it is particularly difficult or problematic to reduce cost or benefit assessments to dollar terms.

I. THE ECONOMICS OF INFORMATION MARKETS

Part I.A explains the mechanics of simple information markets, and Part I.B reviews the limited literature assessing the accuracy of existing markets. Part I.C considers alternative means of structuring information markets to address particular challenges, identifying problems with

“conditional markets” used to assess the effects of a possible policy and explaining how information markets can be constructed when thin trading is anticipated and when the government wishes to subsidize participation to ensure greater accuracy.

A. *The Mechanics*

A basic information market’s operation is simple. The information market’s sponsor issues one or more securities and provides some form of prospectus specifying how each security will eventually be redeemed. Each security’s payout will be some function of a number or numbers that will become objectively verifiable by the time of redemption. The sponsor sells securities to those who wish to participate in the market, and participants subsequently can trade securities with one another. Typically, a market maker will match those wishing to buy and sell a particular security, for example facilitating a transaction when someone is willing to buy at a price greater than or equal to the price at which someone else is willing to sell, which is when the largest bid price is greater than or equal to the largest ask price.³² The prices at which these transactions occur, as well as the bid and ask prices, reflect market predictions of the eventual payout, and thus of the number of numbers on which that payout is based.

The Iowa Electronic Markets, the only legal, continuously operating information markets using real money,³³ offer numerous examples. Most of the information markets operated by the Iowa Electronic Markets involve predictions of the results of political elections,³⁴ although a few nonpolitical markets also exist.³⁵ The election markets are of two types: vote-share markets and winner-take-all markets. In a vote-share market, a security corresponds to a particular candidate or political party and pays off the number of cents equal to the percentage of votes obtained. For example, each security corresponding to a candidate who wins thirty percent of the total votes

³² Bid-ask spreads will be relatively large when there is a relatively large possibility that traders will have asymmetric information, such as in markets in which insider trading is possible. See James Harlan Koenig, *The Basics of Disclosure: The Market for Information in the Market for Corporate Control*, 43 U. MIAMI L. REV. 1021, 1030 n.47 (1989).

³³ Certain Internet gambling sites have created illegal information markets based on the results of sporting events. See, e.g., <http://www.tradesports.com> (last visited June 17, 2003) (providing information market services from Ireland). The Iowa Electronic Markets, by contrast, have received regulatory clearance for their activities. See *infra* note 230 and accompanying text. Other real-money information markets are planned. See, e.g., <http://www.simonmarket.org> (last visited June 17, 2003) (planning to offer an information market for claims about science).

³⁴ See <http://www.biz.uiowa.edu/iem/archive/> (last visited Apr. 24, 2003) (providing an archive of closed markets and historical data).

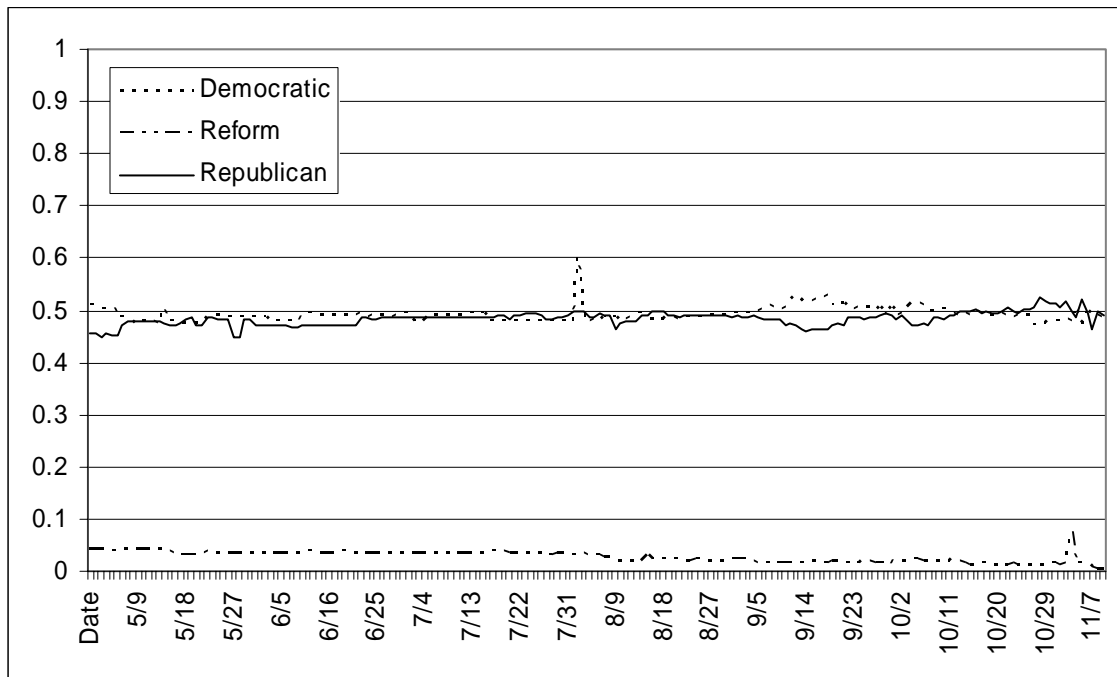
³⁵ For example, a current market predicts interest rates set by the Federal Reserve Board. See <http://www.biz.uiowa.edu/iem/markets/fedpolicyb.html> (last visited Apr. 24, 2003).

would pay \$0.30. In a winner-take-all market, a security corresponding to a candidate or party pays off if and only if that candidate or party receives the most votes in the election.

Consider, for example, Figure 1, which reports the last share price for securities in the vote-share market for the 2000 Presidential election. The graph is relatively stable, with both the Democratic and Republican expected vote shares hovering near fifty percent. The data, however, shows some variation. For example, as Democratic nominee Al Gore gained in the polls in early September his vote share also rose, but by late October, Republican nominee George W. Bush had a slight edge. The contracts eventually paid off at virtually identical prices, with Democratic securities paying at \$0.499, Republican securities paying at \$0.497, and Reform securities paying at \$0.004, all calculated based on the vote totals and shares reporting on November 10, 2000, in the *New York Times* and the *Washington Post*.³⁶

Figure 1: The 2000 Presidential Election Vote Shares Market

This graph indicates the last price at which the securities in the Iowa Electronic Markets' 2002 Presidential election vote shares market were traded on each day from May 1 to November 10.



³⁶ See Dan Balz, *Resolution Days Away as Bush's Lead Shrinks in Fla.*, WASH. POST, Nov. 10, 2000, at A1 (reporting that Gore led in the popular vote by 49,059,936 to 48,858,335). The market defined vote shares as the vote shares among the three parties, so any votes received by other parties did not count. An alternative approach would have been to have an "other" security, representing the vote share of all parties other than those explicitly represented by securities.

For the most part, this chart provides a tentative indication that vote shares markets provide at least roughly rational predictions. Changes in the candidates' fortunes, as indicated by news reports and polls, seem to be reflected in the trading prices. Market prices, moreover, are far less volatile than conventional stock markets.³⁷ Shares of companies sometimes seem to move dramatically without explanation or at least disproportionately to the apparent significance of new information,³⁸ but this market's predictions seem almost always to be at least close to what common sense would suggest. At least one anomaly is visible, however; the Democratic share price has a one-time dramatic spike to 0.60 on August 3.³⁹ This spike has no apparent reflection in campaign information, and it tellingly is not mirrored by a concurrent fall in the Republican share price. This price presumably only reflects the final transaction on August 3, and it may be the product of a foolish or manipulative trader. The lesson is that in a relatively thin market, the last price traded may reflect an anomaly, at least in the absence of a sufficiently sophisticated automated market maker.⁴⁰ If one were relying on the market predictions for practical purposes, it might thus be advisable to ignore price spikes of very short duration and to consider inframarginal bid and ask prices, or average transaction prices over a series of transactions, as potentially better predictors than the share price of the last transaction.

Figure 2 presents the 2000 Presidential election winner-take-all market, and with it an unexpected surprise: Gore wins. This victory in the Iowa Electronic Markets, however, conveys no governmental power, as the prospectus for the election clearly defines the winner of the winner-take-all market as the winner of the popular vote, as reported on November 10.⁴¹ That

³⁷ Many of the causes of stock market volatility, such as economic uncertainty, seem unlikely to affect information markets. Cf. Clifford W. Smith, Jr., *Market Volatility: Causes and Consequences*, 74 CORNELL L. REV. 953, 953-54 (1989) (discussing the causes of stock market volatility).

³⁸ The 1987 stock market crash is a common example of volatility that did not appear to reflect any underlying changes in market fundamentals. See *id.* at 954-56.

³⁹ Another anomaly is less visible on the graph. The final trading prices of the securities on November 10 before the payout were 0.481 for Democratic, 0.004 for Reform, and 0.491 for Republican. This is slightly different from the ultimate payout, even though it appeared on November 9 that Gore would win the popular vote. See, e.g., Akhil Reed Amar, *The Electoral College, Unfair from Day One*, N.Y. TIMES, Nov. 9, 2000, at A23. This anomaly reinforces that in a market without much liquidity, very small price differences based on actual sales may not mean anything. The bid-ask spread at the end of November 10 may have been more sensible, but that information is unavailable.

⁴⁰ See, e.g., Morris Mendelson & Junius W. Peake, *Intermediaries' or Investors': Whose Market Is It Anyway?*, 19 J. CORP. L. 443, 481-82 (1994) (discussing automated market makers). The mechanism by which trades are executed in the Iowa Electronic Markets is described by Robert Forsythe et al. See *Anatomy of an Experimental Political Stock Market*, 82 AM. ECON. REV. 1142, 1145 (1992).

⁴¹ See http://www.biz.uiowa.edu/iem/closed/pr_Pres00_WTA.html (last visited Apr. 24, 2003). The prospectus makes explicitly clear, "Payoffs are NOT affected by ... the outcome of the electoral college or any vote taken by the House of Representatives should such vote be necessary." *Id.*

definition is perhaps unfortunate, as a market dependent on the political party of the individual actually sworn in as President would have revealed the extent to which events like the Supreme Court's decision in *Bush v. Gore*⁴² surprised the market.⁴³ The trends in this graph in any event correspond to those in Figure 1, but the effects are magnified. This is precisely what one would expect when polling numbers are relatively stable but the election is close. The slight lead that Gore had in the polls in September corresponds to a relatively large disparity in the probability that each candidate would win the popular vote, and similarly, although Bush had only a slight lead in expected vote share, that lead seemed sufficiently stable that his chance of winning the popular vote was above 70% shortly before the election. Of course, Bush ended up losing the popular vote, but the fact that what is sometimes unlikely turns out to happen does not by itself mean that the market's prediction was wrong. At least, the market's predictions seem to correspond with third parties' assessments of the likely result of the popular vote,⁴⁴ and once it became clear that Gore would win the popular vote, the price line did an about face.⁴⁵

⁴² 531 U.S. 98 (2000).

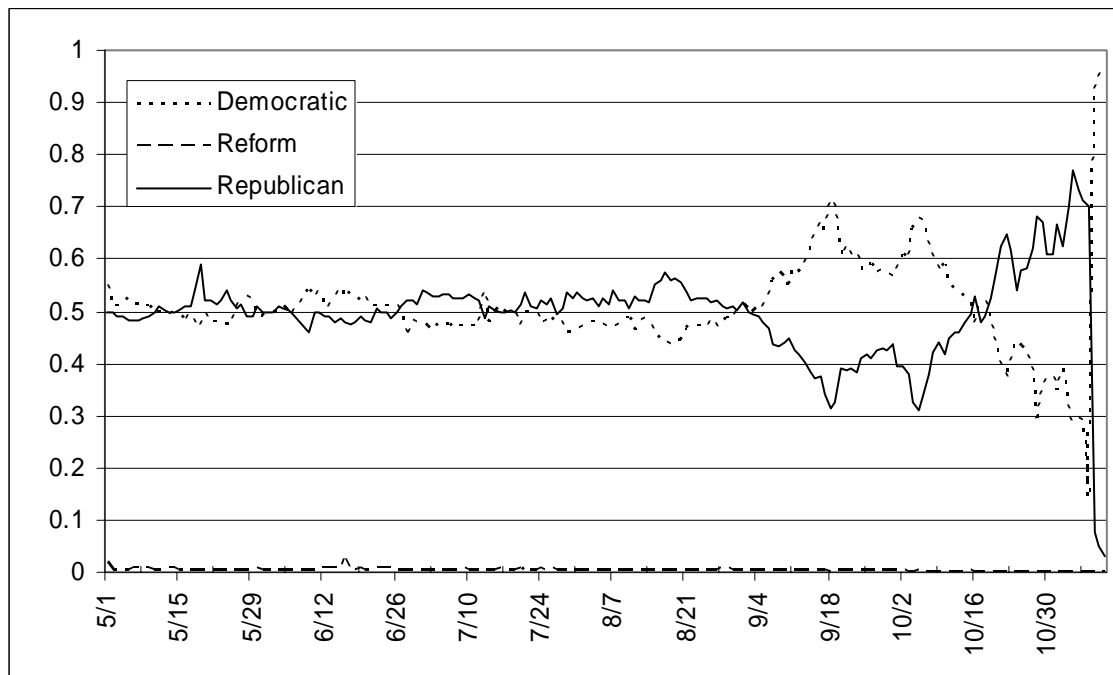
⁴³ The decision may have been motivated by concern that the market produce an unambiguous winner. Suppose, for example, that the House of Representatives had initially deadlocked and then as a compromise chosen an independent as the next President or decided on joint rule by Bush and Gore. These possibilities are fanciful, but the market designers' possible desire to avoid these contingencies may reflect the notion that the markets will be most meaningful when they are entirely objective. *See supra* note 21 and accompanying text.

⁴⁴ *See, e.g.*, Ronald Brownstein, *Polls Show Victory Could Come Without Winning*, L.A. TIMES, Nov. 3, 2000, at A1 (noting that although Bush was winning in the polls, he might win the popular vote and yet lose the electoral college).

⁴⁵ There is another anomaly in this graph, however. The last trading price of the Democratic share was slightly below 1.0, and the last trading price of the Republican share was slightly above 0. At least by the end of November 10, the outcome should have been entirely clear, since the prospectus defined the winner relative to the popular vote shares reported on November 10. This again reinforces that in a market with relatively low liquidity, very slight deviations in the prices at which shares are traded may not be meaningful. *See supra* note 39. Note that bid-ask prices presumably would be more so, given that any rational trader would trade only at \$0 or \$1.

Figure 2: The 2000 Presidential Election Winner-Take-All Market

This graph indicates the last price at which the securities in the Iowa Electronic Markets' 2002 Presidential election winner-take-all market were traded on each day from May 1 to November 10.



One distinguishing feature of the Iowa Electronic Markets is that the values of all the securities in any given market collectively add up to \$1. This is convenient, because it allows for an easy way to sell the securities. For each dollar that a participant pays, the participant initially receives one of each security.⁴⁶ This system ensures that the amount that the Iowa Electronic Markets pays at liquidation is equal to the amount contributed. As a result, the Iowa Electronic Markets avoids both the possibility of losing money and of gaining money, the latter of which might make the markets appear to be more akin to a casino than a stock market. There is nothing inherent about information markets, however, requiring security values to add up to a constant. For example, one can imagine a market in which a security would be deemed worth one cent for every million votes received by a particular candidate. Such a market would provide more information than the vote shares market, predicting total turnout as well as the share of the vote received by each security.⁴⁷

⁴⁶ See, e.g., http://www.biz.uiowa.edu/iem/markets/pr_Pres04_VS.html (last visited June 17, 2003) (“Fixed-price bundles consisting of one share of each of the contracts in this market can be purchased from or sold to the IEM system at any time. The price of each bundle is \$1.00.”).

⁴⁷ A caveat is that prices in such a market might be affected by interest rates. If, for example, traders expected turnout of 50

The Hollywood Stock Exchange is an information market structured along these lines, although it uses only “fake money.”⁴⁸ The Exchange allows for trading on securities corresponding to movies, including both those in production and those in theaters, with each movie security liquidated four weeks after the movie release for \$1 (well, one fake dollar) per \$1 million in box office gross.⁴⁹ Figure 3 provides an example. The Exchange also provides for options trading, thus providing some additional information, although that information can be difficult to decode. For example, if a call option⁵⁰ to purchase a security at a strike price of \$50 trades for \$10, that price is roughly equal to the expected probability that the underlying security will be worth at least \$50 multiplied by the average expected price assuming it is worth at least \$50. That may give some information about the film’s upside potential, but the information is not easy to digest. Only with a large number of options trading for different strike prices could one obtain a full assessment of the probabilities corresponding to each possible level of the film’s success.

million in an election a year away, they would only be willing to pay less than 50 cents for the security. There are several potential solutions to this problem, however. For example, all transactions could be cleared at the time the market is liquidated, regardless of when those transactions actually took place. Thus, if a participant made a purchase for fifty cents five months before the market close, the fifty cents would be paid at the conclusion of the market.

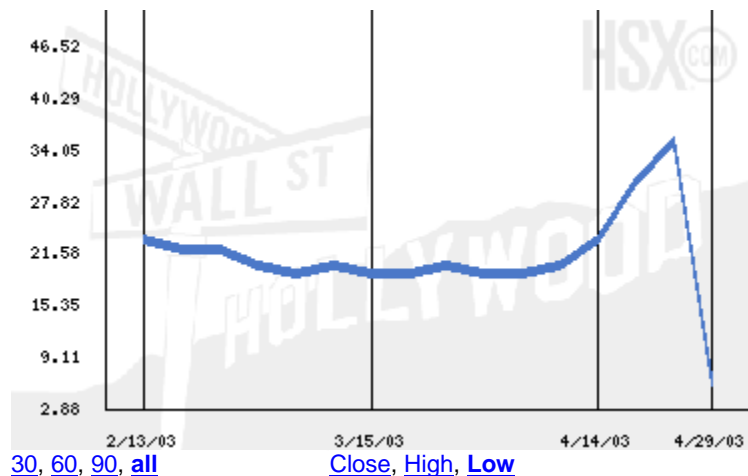
⁴⁸ See <http://www.hsx.com> (last visited Apr. 23, 2003).

⁴⁹ See <http://www.hsx.com/help/topics/whatcan.htm#1> (last visited Apr. 23, 2003). Similarly, the Exchange offers StarBonds, corresponding to individual movie stars. These bonds are liquidated when a star’s career ends, whether by unnatural causes, natural causes, or retirement, based on the average gross of the star’s previous five pictures. *See id.*

⁵⁰ A call option is an option to purchase at a set price. *See generally* Kenneth A. Froot et al., *A Framework for Risk Management*, 72 HARV. BUS. REV. 91, 99 (1994) (providing an introduction to options).

Figure 3: A Stock on the Hollywood Stock Exchange

This screen shot from the Hollywood Stock Exchange's Internet site traces the price line of the stock corresponding to the "Real Cancun." The stock fell precipitously after an unexpectedly poor opening weekend.⁵¹



B. Accuracy

The 2000 Presidential election information markets alone provide some anecdotal evidence about the accuracy of information markets. The winner-take-all market's prediction on election eve that Bush had a 70% chance of winning the popular vote, for example, seems to roughly correspond to what we might have expected based on the evidence at the time. The price seems to reflect a reasonable consensus, even though it would have been more impressive if the market had predicted, against the apparent pre-election consensus, that Gore would win the popular vote. Forsythe et al. provide a more complete analysis of the 1988 Iowa Electronic Markets, showing how the market prediction of the candidates' vote was more stable over the course of the election than poll results.⁵² Perhaps more interesting, Forsythe et al. argue that poll results did not drive market prices; that is, traders anticipated shifts in candidates' fortunes as reflected in polls before those polls actually occurred.⁵³ The study thus suggests that information markets are not simply crude aggregators of other predictors, but manage to effectively incorporate difficult-to-interpret data. This conclusion should be treated with caution, however,

⁵¹ See, e.g., Leanne Potts, 'Cancun' Loses Its Shirt, ALBUQUERQUE J., May 2, 2003, at D1.

⁵² Forsythe et al., *supra* note 40, at 1150.

⁵³ *Id.* at 1153 ("Evidently, traders were able to find out about the mood of the electorate without relying on opinion polls. In this sense, polls are not 'news' to traders who have an incentive to seek out information from other sources.")

as a study of a similar market used to predict an election in the Netherlands reached a different conclusion.⁵⁴

It is only over a number of markets that information markets' accuracy can be assessed. Because information markets purport only to give the best guesses possible based on available evidence, any single success or failure may reflect luck. Over a number of markets, however, it is possible to measure information markets' average performance.⁵⁵ Berg et al., for example, have considered all of the Iowa Electronic Markets' vote share markets.⁵⁶ In the five vote-share markets related to presidential elections, involving a total of twelve contracts, the markets had an average absolute error of 1.37%, meaning that on average each election eve prediction was off by 1.37%.⁵⁷ In other U.S. elections, the average absolute error was 3.43%, largely because of two primary elections in which the markets exhibited unjustified confidence in Paul Tsongas, and in non-U.S. elections, the markets had an average absolute error of 2.12%.⁵⁸ Collectively, these data seem sufficient to establish that the markets' predictions were neither haphazard nor perfectly omniscient.⁵⁹

The problem in assessing these data is the lack of a control group. A decisionmaker without any alternative estimate of the relevant variable would benefit from considering an information market prediction, but in real decision contexts, the choice is between relying on an information market and relying on one or more experts. One experiment, involving Hewlett Packard printer sales' predictions, provides slight evidence supporting the hypothesis that markets are better than experts.⁶⁰ The market beat the official expert forecast six of eight times,⁶¹

⁵⁴ Ben Jacobsen et al., *(In)accuracy of a European Political Stock Market: The Influence of Common Value Structures*, 44 EUR. ECON. REV. 205, 216 (2000).

⁵⁵ Formulas exist that make it possible to rate the accuracy of those who make repeated probability estimates, even though each individual probability estimate admits some uncertainty. *See, e.g.*, G.W. Brier, *Verification of Forecasts Expressed in Terms of Probability*, 78 MONTHLY WEATHER REV. 1 (1950) (offering an early example of such a formula).

⁵⁶ Joyce Berg et al., *Results from a Dozen Years of Election Futures Markets Research* (Nov. 2000) (unpublished manuscript, on file with author).

⁵⁷ *Id.* app. at 1.

⁵⁸ *Id.*

⁵⁹ The mean average error for the securities in the Netherlands study was 3.1%, which was greater than the mean absolute errors that had been reported up to that time, but still seems roughly consistent with this conclusion. *See* Jacobsen et al., *supra* note 54, at 211. Jacobsen et al. also argue that a form of the winner's curse led to overpricing of securities corresponding to parties with relatively small shares of the vote. *See id.* at 224-27. If this explanation is accurate, it seems likely to endure, given the opportunity that others would have to exploit this tendency once it is recognized.

⁶⁰ *See* KAY-YUT CHEN & CHARLES R. PLOTT, *INFORMATION AGGREGATION MECHANISMS: CONCEPT, DESIGN AND IMPLEMENTATION FOR A SALES FORECASTING PROBLEM* (Cal. Inst. of Tech. Social Science Working Paper No. 1131, 2002).

⁶¹ *Id.* at 20.

even though the experts made their predictions after the markets closed. An analysis of the Hollywood Stock Exchange, however, shows that an expert predictor of movie returns, Brandon Gray of Box Office Mojo,⁶² slightly outperformed the market.⁶³ The average stock exchange had an average percent error of 31.5%, while Box Office Mojo exhibited an error of only 27.5%.⁶⁴

Even placing aside their conflicting results, these studies would be relatively uninformative because the outcomes may depend on the setup of the particular experiment. Perhaps Hewlett-Packard hired unusually bad forecasters to make the official prediction. Perhaps Gray got lucky, or perhaps the Hollywood Stock Exchange's imperfections can be traced to the use of fake rather than real money. The ultimate question is whether experts or markets are likely to outperform the other on average assuming that equal resources are provided for each task. It would be possible to allocate some money amount either to hiring an expert or subsidizing an information market.⁶⁵ No set of even hypothetical experiments seems sufficient to provide a definitive answer, or at least one in favor of information markets, given the impossibility of finding an objectively superior process for hiring experts.

Perhaps the most that can be said on the basis of such experimental data is that information markets and well-motivated experts are roughly comparable. In my judgment, this is probably sufficient to justify further corporate as well as governmental experimentation with information markets, but any ultimate benefit attributable to markets' information aggregation powers alone is likely to be relatively small. The more significant potential payoff from information markets comes if there is reason to believe that some experts make systematic errors or are not well-motivated. This is possible in some corporate contexts; perhaps an internal market used to predict quarterly earnings would be less susceptible to optimism biases than more traditional approaches.⁶⁶ It is in governmental decisionmaking, however, where there is the greatest reason to be suspicious of experts, either because of external influence or because of

⁶² See <http://boxofficemojo.com> (last visited Apr. 28, 2003).

⁶³ David M. Pennock et al., *Extracting Collective Probabilistic Forecasts from Web Games*, in PROCEEDINGS OF THE SEVENTH ACM SIGKDD INTERNATIONAL CONFERENCE ON KNOWLEDGE DISCOVERY AND DATA MINING (2001).

⁶⁴ *Id.* at 4 tbl. 1.

⁶⁵ See *infra* Part I.C.3 (discussing subsidizing markets).

⁶⁶ Optimism biases sometimes lead individuals to believe that they are more likely than most to avoid risks such as that of business failure. See, e.g., David A. Dana, *A Behavioral Economic Defense of the Precautionary Principle*, 97 NW. L. REV. 1315, 1325 (2003) (“[P]eople may believe, even in the absence of any factual basis, that with time they will find a costless means to avoid future risks.”); Barton H. Thompson, *Tragically Difficult: The Obstacles to Governing the Commons*, 30 ENVTL. L. 241, 264 (2000) (“[W]hen confronted by an uncertain future, most people assume that they will be able to avoid, reduce, ameliorate future risks.”).

ideological agendas. Information markets seem to have information aggregation capabilities that are at least good enough for government work. The question, to which I shall return below,⁶⁷ is whether information markets are indeed objective, and whether a tool that allows the government to make objective forecasts would be useful in governmental decisionmaking.

C. Market Design Challenges

1. Conditional Markets

Perhaps the most ambitious information market form proposed by economists to date is a conditional market,⁶⁸ which assesses the results of different choices that a decisionmaker might make. Such a market harnesses information about the anticipated effects of a decision on some number of interest to the market participants. Berg and Rietz have used a version of the Iowa Electronic Markets in the 1996 Presidential election campaign to demonstrate the power of information markets to assess conditional probabilities.⁶⁹ The securities in the vote share market included two securities for each potential Republican nominee conditional on that nominee's receiving the nomination. For example, one security represented the percentage of votes that Robert Dole would receive in the general election, while another represented the percentage of votes that the Democratic candidate would receive against Dole. At any time before the nomination, the sum of these two securities would have reflected the market's estimate of the probability that Dole would win the nomination. The relative values of these shares would indicate Dole's projected performance relative to Clinton assuming that Dole was the nominee. Berg and Rietz note that Republicans could have used the market prices to recognize that Dole was not the strongest possible candidate against Clinton.⁷⁰ The same conditional market approach is being used for the 2004 Presidential election, and Table 1 shows a snapshot of market prices and what those prices indicated about the market's evaluation of candidates' relative prospects.

⁶⁷ See *infra* Part II.A.2.

⁶⁸ See, e.g., Robin Hanson, *Conditional Markets* (2002) (unpublished presentation, on file with author) (providing a careful analysis of such markets).

⁶⁹ See Joyce A. Berg & Thomas A. Rietz, *Prediction Markets as Decision Support Systems*, 5 INFO. SYS. FRONTIERS 79 (2003).

⁷⁰ This particular result may not have been all that surprising. Other commentators recognized Dole's weakness at the time, pointing to polling results involving hypothetical matchups between President Clinton and various Republican candidates. See David S. Broder, *Many Still Making Up Their Minds; Shifting Loyalties Mark Final Days*, WASH. POST, Feb. 12, 1996, at A1. Of course, some Republican voters might have preferred Dole over a candidate with a better chance of winning because they preferred Dole enough to make it seem worth the risk.

Table 4. The 2004 Presidential Election Conditional Vote Share Market

The table provides a snapshot of the 2004 Presidential Election market, as of April 28, 2003 at noon. The first column lists each potential nominee, including one undeclared candidate (Hillary Clinton) and an “other” category, including such potential nominees as John Edwards. Two securities correspond to each potential nominee. The two securities corresponding to the eventual nominee will be liquidated based on the eventual two-party vote share in the election, with all other contracts liquidating at zero. The third, fourth, and fifth columns report the bid and ask prices, followed by the midpoint of these. The last two columns calculate the market’s estimate of the chance of the candidate’s winning the election and the vote share assuming the candidate is nominated based on these midpoints. For example, Hillary Clinton’s 0.058 chance of winning the nomination is the sum of the bid-ask midpoints (0.018 and 0.04) of the securities corresponding to her nomination. (Note that this measures her absolute chance of winning the nomination, not her chance of winning conditional on deciding to run.) Clinton’s conditional vote share of 0.333 is equal to her bid-ask midpoint divided by her overall chance of winning the election.

Potential nominee	Conditional security for	Bid	Ask	Midpoint	Chance of winning nomination	Conditional vote share
H. Clinton	Democrat	0.016	0.02	0.018	0.058	0.333
	Republican	0.032	0.048	0.04		
R. Gephardt	Democrat	0.036	0.043	0.0395	0.096	0.391
	Republican	0.056	0.057	0.0565		
J. Kerry	Democrat	0.103	0.153	0.128	0.29	0.393
	Republican	0.159	0.165	0.162		
J. Lieberman	Democrat	0.042	0.045	0.0435	0.098	0.438
	Republican	0.054	0.055	0.0545		
Other	Democrat	0.206	0.209	0.2075	0.451	0.462
	Republican	0.24	0.247	0.2435		

The calculations in Table 1 are straightforward largely because the securities in the conditional market are normalized so that the total payoffs are known, in this case one dollar. This approach can be used to assess the extent to which a particular decision will affect some other probability. Robin Hanson, for example, offers analysis of a hypothetical conditional market used to determine how a decision whether to move troops will affect the probability that a war will occur.⁷¹ The market involves four securities. The first will pay off \$1 if troops are moved and there is a war; the second, if troops are moved and there is no war; the third, if the troops are not moved and there is no war; the fourth, if troops are not moved and there is a war. From the securities’ prices one could calculate both the market’s estimate of the chance that a decision to move troops will be made and the market’s estimates of the correlation between a decision to move troops and the probability that war will result from such a decision.

Hanson, however, also notes that conditional markets might be used to assess how specified decisions might affect non-binary variables, such as stock price, GDP per capita, or

⁷¹ Hanson, *supra* note 68, at slides 2-3.

unemployment rates.⁷² The mechanics here become somewhat more complicated. Suppose, for example, that a conditional market is to be used to assess the effect of proposed airline safety regulations on the number of people who will choose to fly on airplanes in a particular year, because one anticipated effect of the regulations is increased consumer confidence.⁷³ The government could issue a security that will pay off one cent per a specified number of passenger miles if the regulations are adopted and another that will pay off at the same rate if the regulations are not adopted. The values of these securities, however, would depend also on the market's prediction about whether the regulations are adopted. A simple solution would be to create an additional information market estimating the probability that the agency will indeed adopt the regulations, for example by issuing two securities, with one selected to pay off \$1 depending on whether the regulations are adopted.⁷⁴ The price of each security divided by the probability of the corresponding decision would then provide a prediction of the passenger miles in each scenario.

Information markets may produce misleading predictions, however, if there is a possibility that the eventual decisionmaker might have information unavailable to the market. Suppose that an airline safety market participant estimates, based on her own information, that there will be 500 billion passenger miles if the regulations are adopted and 490 billion otherwise, but the participant also knows that the decisionmaker has additional information that could allow for refinement of these estimates and therefore might affect the eventual decision. The decision itself thus might provide some indication of the content of this information. The market participant should reason that if the regulations are enacted, then the probability that the information is favorable to the regulations is higher than the participant's initial estimate, and vice versa if the regulations are rejected. As a result, the participant should price the security

⁷² *Id.* at slide 4.

⁷³ Of course, a market equally could be used to measure the variable that is more obviously of interest, the number of anticipated deaths in airplane crashes. I use the passenger miles to place aside questions about the appropriateness of using an information market to predict deaths and about whether information markets might create a moral hazard problem by encouraging people to predict high death totals and then cause those deaths. *See infra* note 184 and accompanying text.

⁷⁴ An alternative solution would be for the government to provide that the transactions on the irrelevant market would be cancelled, with all investors in that market receiving refunds. For example, if the regulations were adopted, then the information market corresponding to nonadoption would be cancelled. Such a market would give no indication of the probability of adoption but would indicate the effect of adoption. A drawback of this approach is that if there is a very high probability of adoption or nonadoption, there might be very little incentive to trade on the market that will be cancelled, and its results may thus be unreliable. This may not be a large concern, however, since the very high probability would presumably indicate that the government is not expected to take into account the market prediction in any event.

conditional on enactment as above the 500 billion level, and the security conditional on rejection as below that corresponding to 490 billion. Robin Hanson, who has recognized this potential problem, suggests as a possible solution: to allow decisionmakers to trade, so that there will be no information unavailable to the market.⁷⁵ This solution is problematic, however, because decisionmakers might then be tempted to make decisions that would maximize their trading gains rather than social welfare.

There are two additional reasons, besides this selection bias problem, suggesting decisionmakers should hesitate before blithely accepting information market predictions calculated on the basis of differences in security values. First, conditional markets may lead to expenditure of effort on issues of little interest to the decisionmaker, as market participants may factor in a number of variables besides whether a particular decision is made. In the market used to predict the effect of proposed regulations on airline safety, for example, market participants might devote considerable resources to the overall task of modeling passenger miles. If such a market were subsidized,⁷⁶ only a portion of the subsidy would go toward efforts focused on the variable of interest, the effect of the regulations. If the regulations are expected to have a very small effect on consumer behavior, few if any resources would be spent on this portion of the analysis. There is no obvious way to determine how much effort market participants expended at calculating the effect of the regulations. Any incomplete evaluation of relevant information in effect adds noise to the estimate of interest.

Second, slight differences in security prices could be a result of a different type of noise. If there is some noise in the data used to generate predictions, such as the last transaction price or the midpoint of bid-ask spreads, then that noise may overwhelm the variable of interest. If the increase in passenger miles caused by the regulation corresponds to a price effect that is smaller than the bid-ask spread, for example, traders might have insufficient incentive to arbitrage away any difference in price associated with the two securities. The greater the liquidity of the market, the more closely market variables such as transaction price are likely to reflect the actual consensus of market participants regarding the variables in question. As discussed below in more

⁷⁵ Hanson, *supra* note 68, at slide 12.

⁷⁶ See *infra* Part I.C.3.

detail,⁷⁷ in a relatively thin market, there could be significant deviation between price and consensus.

Each of these problems associated with conditional markets could have significant effects on the calculation of the effect of the relevant decision on the variable of interest. The bias attributable to this selection effect, for example, might be slight relative to the overall prediction of passenger miles, but large relative to the predicted difference in passenger miles attributable to a particular decision. If the bias leads to security prices corresponding to 485 and 505 billion miles, for example, then the bias in effect doubles the conditional market's estimate of the effect of airline safety regulations. The resulting data might still be useful by showing that the regulations will be expected to have only small effects. In some cases, this might itself be sufficient the regulations should not be enacted, assuming that consumer confidence rather than actual safety is the regulatory goal. But often in regulation, even relatively small effects matter, and indeed an agency might consider effects on consumer confidence only because the issue is close with respect to safety fundamentals. In that case, the government might want to know the market's exact prediction of number of lives saved by enacting the regulations.

In sum, while information markets in general can make predictions that may be used as inputs in subsequent decisionmaking, conditional markets' usefulness particularly stems from their ability to predict the effect of hypothetical decisions on variables of interest. Slight imperfections in information markets, however, are magnified when the predicted decision's effect can only be discerned by analyzing multiple market prices. Conditional markets are thus most likely to be useful where the decision is expected to have a substantial effect on the variables of interest. Often, though, conditional markets will not be useful in such a situation, because when a decision is expected to have a significant effect, the direction of that effect is likely to be known. These problems with conditional markets might thus suggest that it is not possible to construct an information market that produces a useful evaluation of a potential decision. I will argue in Part III, however, against this conclusion; predictive cost-benefit analysis, while having some weaknesses relative to conditional markets, also succeeds in avoiding many of their problems.

⁷⁷ See *infra* Part I.C.2.

2. *Thin Markets*

Information markets are a tool for generating consensus predictions from a number of participants, but they may be less effective when markets are thin. In thin markets, trades occur relatively rarely, and there is a danger that the most recent transaction will not represent the market consensus. For example, even if all traders recognize that the value of the securities has changed, they may have no reason to buy or sell these securities. The problem may be particularly severe in an information market in which some traders may obtain significant information that others lack. The greater the degree of information asymmetry, the larger bid-ask spreads will be,⁷⁸ and transactions will occur less frequently. The Iowa Electronic Markets seems to have sufficient liquidity to enable frequent transactions, but even there, as indicated above, the relative thinness of the market can make it hazardous to rely on the last transaction price.⁷⁹ When dealing with an information market in which there are fewer likely traders and asymmetric information is more of a concern, however, the standard mechanism may prove more problematic.

The central problem of thin information markets is that information markets rely on voluntary transactions to assess market consensus, and under certain market conditions voluntary transactions will not sufficiently occur. One can imagine many nonmarket mechanisms, however, that allow compulsory transactions. Consider, for example, the following simple mechanism: An initial predictor makes a prediction of the variable of interest.⁸⁰ Anyone can challenge the initial predictor to a bet, which the initial predictor is required to take. The challenger then announces a new prediction and can be challenged in turn. So, if *A* predicts 5, *B* could bet *A* that the amount will be higher and announce a new prediction of 8. Then, *C* could bet *B* that the actual amount is 7. Suppose the market then closes and the amount turns out to be 6. If each unit of prediction corresponds to \$1, then *B* would win \$1 against *A*, but *C* would win \$1 against *B*. Such a simple betting scheme performs the same information aggregation function as a securities market, but it encourages the predictors to update their predictions to reflect new

⁷⁸ See *supra* note 32.

⁷⁹ See *supra* text accompanying note 40.

⁸⁰ Because the initial predictor is required to take all bets, some incentive would need to be given for someone to become the first predictor.

information lest they be challenged, and it allows those with new information to trade on that information even if others would prefer not to trade because of the informational asymmetry.

There are ways to integrate compulsory transactions into the market metaphor, for example, enacting a rule that requires a security holder to announce a price which then serves as an offer for anyone else to either buy the security at that price or sell an identical security to the holder at that price.⁸¹ The central point, however, is that the market metaphor is unnecessary. Information markets work because they allow market participants to profit when they correctly identify that the consensus prediction is inaccurate. The market mechanism is both useful and familiar, because we know that capital markets tend to aggregate predictions. The above compulsory betting procedure, however, accomplishes much the same thing, and at least in one respect, is better than the traditional capital market, because it allows participants to make positive expected-value bets based on even small amounts of information. An additional virtue of this procedure is that it helps distinguish information markets from more conventional securities markets and the pathologies that behavioral finance scholars have identified within them.⁸² No one would casually use this form of the information market as an investment vehicle, given the risk of being subject to a bet.

Robin Hanson offers a similar proposal to overcome market thinness.⁸³ Hanson builds on the literature on “scoring rules,” which are rules that can be used to compensate individual experts to induce honest valuations from them. For example, suppose that I wanted to give an expert an incentive to estimate for me the average temperature this coming winter. I could ask for an estimate and agree to pay the expert based on the difference between the prediction and the eventual observed value, with a higher payment amount the closer the expert is. Modestly more complicated scoring rules, the focus of the scoring rules literature, can be used to generate predictions as to the probability of events.⁸⁴ A scoring rule is similar to an information market in that participants have an incentive to predict the events of interest to the sponsor. It is, in effect,

⁸¹ See Michael Abramowicz, *Trial by Market* (2003) (unpublished manuscript, on file with author) (describing such a mechanism).

⁸² See *supra* note 13.

⁸³ Robin Hanson, *Combinatorial Information Market Design*, 5 INFO. SYS. FRONTIERS 103 (2003).

⁸⁴ See, e.g., Robert T. Clemen, *Incentive Contracts and Strictly Proper Scoring Rules*, 11 TEST 167 (2002); Robert L. Winkler, *Scoring Rules and the Evaluation of Probability Assessors*, 64 J. AM. STAT. ASS'N 1073 (1969).

an information market that could be used in the thinnest of all possible markets, a market with just one participant.

Hanson suggests what he calls a “market scoring rule” as a way to provide a mechanism that works like a scoring rule when there are only a small number of participants and like a more traditionally structured information market when there are more participants. Under the market scoring rule, after the initial prediction, anyone else can make a subsequent prediction, as long as the subsequent predictor in effect agrees to pay off the current predictor when the market closes. The subsequent predictor thus receives a payment equal to the difference in the payments that the scoring rule would provide to the two predictors; if this number is negative, the subsequent predictor would pay money.⁸⁵ This system is identical to the betting scheme discussed above, except that each subsequent predictor’s bet is with the house rather than with the previous predictor. The advantage of this approach relative to the betting scheme described above is that a potential predictor need only worry about coming up with a prediction that is likely to be better than that of the current predictor, but need not worry that a subsequent predictor will further refine his or her prediction. Given the risk, individuals might still choose not to act on very minor pieces of information, but there will be more of an incentive to do so than with the betting scheme.

3. *Subsidized Markets*

While the Iowa Electronic Markets are not subsidized, subsidies will improve market participants’ incentives to engage in research and analysis, and further subsidies may be necessary for markets that are less intrinsically interesting than the Iowa Electronic Markets. Hanson’s market scoring rule provides an easy mechanism for subsidizing information markets. Recall that after the initial predictor, each subsequent predictor receives a sum equal to the amount by which the predictor achieved a better payout than her predecessor based on the scoring rule. The sponsor of the market therefore will end up paying the difference between the

⁸⁵ Suppose a market scoring rule is to be used to predict interest rates a year hence. A scoring rule for a market predicting interest rates might provide that the predictor will receive \$12,000 minus \$1,000 for each quarter point by which the estimate turns out to be inaccurate (and thus will have to pay money if the prediction turns out to be wrong by more than three percentage points). Suppose that Predictor 1 picks 2.5%; Predictor 2, 4.0%; Predictor 3, 3.25%; and the actual interest rate turns out to be 3.75%. The scoring rule would assign Predictor 1 \$7,000; Predictor 2, \$11,000; and Predictor 3, \$10,000. With the market scoring rule, Predictor 1 receives \$7,000, Predictor 2 receives \$4,000 (\$11,000 - \$7,000), and Predictor 3 pays \$1,000 (\$10,000 - \$11,000). The example shows the possibility that there might be a windfall for the first predictor (or depending on the scoring rule, an expected liability, in which case no one would want to participate). See *infra* note 89 and accompanying text.

amount that the scoring rule would indicate that the last predictor should receive and the amount that the scoring rule would indicate that the first predictor should receive, and this difference is the effective subsidy. There are two issues that the literature does not explicitly consider, however. The first is how the first predictor would be selected. Being the first predictor would provide a windfall, at least with any scoring rule that always produces a positive reward.⁸⁶ The second issue is how large a subsidy should be provided.

One approach would be for the sponsor to make an initial prediction of the variable of interest. It could then announce a scoring rule that would produce a reward based on the amount by which the eventual prediction improved upon this additional one (or a penalty if the eventual prediction were further away). For example, in the airline safety regulation market described above,⁸⁷ the government might announce a prediction of 500 billion miles and then announce that for each one billion improvement in predicting miles, the government would pay a total of \$1000. If the eventual prediction were 505 and the total number of miles turned out to be 510, the government would pay a total of \$5000.⁸⁸

This does not identify a complete solution to the first problem identified above, the selection of the first predictor. Rather, there still might be a windfall to the second predictor, and thus a race to be that second predictor, if the initial market consensus is that the government's initial prediction was poor. A simple solution to this is to auction off the right to be the second predictor.⁸⁹ If the government announced a prediction of 490 billion miles, but it were widely recognized even before research that 500 billion would be a more sensible initial estimate, then the government would receive nearly \$10,000 in auction revenues for the right to be the second predictor,⁹⁰ and then would end up paying \$10,000 back upon the close of the market. The

⁸⁶ Some subsequent predictors might lose money with such a scoring rule, as in the example of Predictor 3 in footnote 85, but that is only because they did not predict as accurately as their predecessors.

⁸⁷ See *supra* note 74 and accompanying text.

⁸⁸ Suppose, for example, after the government's initial prediction of 500 billion miles, *A* predicts 503, *B* predicts 502, *C* predicts 518, *D* predicts 522, and *E* predicts 505, with the actual total turning out to be 510. The scoring rule would assign \$1000 for each unit improvement in the prediction. Thus, the initial assignments of the scoring rule (before subtractions) would assign rewards of \$3000 for *A*, \$2000 for *B*, \$2000 for *C*, -\$2000 for *D*, and \$5000 for *E*. The amount that each predictor actually receives, however, is the difference between the scoring rule amount for that predictor and the previous one. So, the final payments as determined by the market scoring rule would be \$3000 for *A*, -\$1000 for *B*, 0 for *C*, -\$4000 for *D*, and \$7000 for *E*.

⁸⁹ Auctions are a commonly proposed approach to eliminating windfalls that would exist merely as a result of being first. See, e.g., Julie Rubin, *Auctioning Class Actions: Turning the Tables on Plaintiffs' Lawyers*, 57 BUS. LAW. 1441, 1443-1445 (1997); Robert G. Hansen & Randall S. Thomas, *Auctioning Class Action and Derivative Lawsuits; A Critical Analysis*, 87 NW. U. L. REV. 423, 428-429 (1993) (finding that in most class action lawsuits the court appoints as lead counsel the first attorney to file the lawsuit).

⁹⁰ It would not receive quite this much back, because of the risk associated with being the second predictor, given the inherent

incentives to improve on the prediction of 500 billion miles would exist as before, but the windfall would be eliminated.

Regardless of whether an auction is used to eliminate windfalls, an advantage of a subsidization scheme that provides a set reward for specified improvements in predictions is that the sponsor of the market can set the subsidy based on its assessment of how valuable improvements in the information would be. In this case, for example, potential participants will invest their time and resources into market participation whenever they believe that an investment of \$1000 (including the risk associated with participation) will be expected to produce a refinement in the estimate by one billion passenger miles. It is difficult to know how valuable more accurate information is. Presumably, the subsidy should reflect the reliance the government is likely to place on the information market; the higher the subsidy, the greater the reliance the government is likely to place on the information market. This system has the virtue of allowing the market sponsor to consider the information's value directly, whereas other means of providing subsidies may obscure the question of how much research the sponsor is encouraging individuals to conduct. Despite this benefit, there may be occasions in which the sponsor wishes to provide a set subsidy for a particular market, for example, because it needs to be able to anticipate the cost of running the market. This would be a straightforward calculation, as the subsidy, plus any auction revenues,⁹¹ could be divided in proportion to the reward indicated by the market scoring rule.⁹²

II. INFORMATION MARKETS AND ADMINISTRATIVE DECISIONMAKING

Information markets in theory might be used in any decisionmaking environment, and corporations have experimented to determine how effectively information markets aggregate

uncertainty about the prediction being made. The predictor in effect would demand some compensation for this risk by bidding less than \$10,000 even if the expected receipts from participating were \$10,000 with some variance.

⁹¹ See *supra* text accompanying notes 89-90.

⁹² In the example of footnote 88, suppose the preset subsidy were \$10,000. Then, the scoring rule would assign, instead of \$1000 per unit improvement, one point per unit improvement. The market scoring rule would subtract each participant's points from the previous participant's and then distribute the subsidy in proportion to these points. The final payments would end up being exactly twice those listed in the hypothetical: \$6000 for A, -\$2000 for B, 0 for C, -\$8000 for D, and \$14000 for E.

The only complication would occur if the final prediction were *worse* than the initial prediction. In this case, the market scoring rule could be adjusted so that all positive rewards (i.e., rewards corresponding to predictions that did move in the right direction) were multiplied by the constant that would result in the government's paying on net the amount of the subsidy. In the unlikely event that there were *no* predictions moving the market in the proper direction, then the government could distribute the subsidy so that those who caused the least damage received the most. For example, if A's market score were -1 and B's market score were -2, then A would receive two-thirds of the subsidy.

information.⁹³ If the information aggregation benefits are relatively modest, however, then it might not seem worthwhile to experiment with information markets in governmental decisionmaking. For reasons of tradition and continuity, governments are hesitant to innovate with respect to decisionmaking structures,⁹⁴ and the innovations that administrative agencies have undertaken are generally responsive to some government-specific need.⁹⁵ If information markets are to become a particularly useful tool in governmental decisionmaking, it will not be because governmental decisionmakers themselves want to use the tool to improve their own decisions. Rather, information markets are best justified as a means for disciplining governmental decisionmaking. Part II.A explains why objective prediction tools might be particularly useful in an area as rife with agency problems as governmental decisionmaking, and it assesses the extent to which concerns about manipulation and unrepresentative traders undercut the possibility of objectivity. Part II.B then considers a few specific examples of how information markets might assist governmental decisionmaking, selecting examples that illustrate various design problems with information markets.

A. *The Objectivity of Information Markets*

1. *Why Objectivity Matters*

Let us suppose that government researchers developed a crystal ball that allowed officials to see what the future will be like under any given set of policies. Would the government use the crystal ball? For some purposes, perhaps not; we might not want to know our individual destinies

⁹³ See *supra* note 60 and accompanying text. A company called Intellimarket offers services in establishing internal information markets for companies seeking to perform estimations. See <http://www.wrsasc.com/default.cfm?fuseaction=tbAboutintellimarket> (last visited June 17, 2003).

⁹⁴ Even in areas in which competitive federalism provides an incentive to improve regulatory structures, states still may hesitate to innovate. See Michael Abramowicz, *Speeding up the Crawl to the Top*, 20 YALE J. ON REG. 139 (2003) (explaining why there might be underinnovation in corporate law).

⁹⁵ A useful example of such an innovation is negotiated rulemaking, which provides a means of passing regulations when all principal constituencies can agree on what the rules should provide. See generally 5 U.S.C. § 561 (2000) (providing the statutory foundation for regulatory rulemaking); Matthew J. McKinney, *Negotiated Rulemaking: Involving Citizens in Public Decisions*, 60 MONT. L. REV. 499 (1999) (arguing that negotiated rulemaking can solve a number of problems currently associated with traditional notice-and-comment rulemaking by increasing citizen participation during both the drafting and initial implementation stages of the proposed regulation, thus increasing the level of compliance with the rule, reducing litigation over the implementation of the regulation, and shortening the total time required to pass the regulation); Jody Freeman, *Collaborative Governance in the Administrative State*, 45 UCLA L. REV. 1 (1997) (concluding that negotiated rulemaking and other forms of collaborative government, while falling short of the ideal form of collaborative government, have already proved advantageous in the areas of health and safety and environmental regulation). Negotiated rulemaking would have limited applicability in corporate contexts, where hierarchical structures minimize the need to satisfy a broad range of constituencies.

or those of our favorite sports teams. But it is hard to imagine an area of administrative law for which the crystal ball would not be useful. The Department of Homeland Security could anticipate terrorist attacks, and a Department of Precrime could intervene to stop murders before they occur.⁹⁶ The Environmental Protection Agency could assess the effects of both global warming and the policies that might limit it. The Securities and Exchange Commission could consider which market rules would lead to the highest future stock prices. The Immigration and Naturalization Service could assess the effects of immigration on economic growth and crime. And on and on. None of these analyses would be sufficient to determine the course of policy, for in all of these areas, few if any policies are Pareto efficient.⁹⁷ But a crystal ball surely would at least be useful by narrowing administrative decisionmaking to the selection of which world we would prefer.

Information markets are not a crystal ball. They cannot predict what *will* happen; rather, they can only give us probabilistic predictions. But let us suppose that information markets are accurate, in the sense that they aggregate information as well as alternative approaches such as hiring experts, and that they are objective, in that they have no ideological or other biases and cannot be manipulated by government or outside agents. Suppose further that they are accepted as providing a best guess, so that an argument against an information market prediction would receive little weight in policy discourse. While information markets cannot allow us to select a single future, with these assumptions they can permit us to select from a probability distribution of possible futures, at least based on the variables that we select for the markets to predict. This still narrows agency decisionmaking down to normative assessments rather than to the task of making scientific and other predictions. Of course,⁹⁸ these assumptions are not trivial, and

⁹⁶ See *MINORITY REPORT* (20th Century Fox 2002) (exploring such a scenario). The movie questions whether we would want to rely on such a system if it were the least bit uncertain, as represented by the occasional “minority report,” or dissenting opinion as to the future. The concern is that we might end up arresting people for crimes that in fact they would not have committed, and an additional concern is that it might be morally problematic to arrest people for crimes that they *would* have committed but did not actually commit. Somewhat puzzlingly, however, the movie never considers the possibility that the system could be used to thwart crime but not to incarcerate the putative perpetrators. Perhaps the filmmakers felt an intrinsic discomfort with predictive legal tools. Or maybe they worried that acknowledging the distinction would have made the movie less interesting.

⁹⁷ Pareto efficient policies are those that make everyone better off. See generally Guido Calabresi, *The Pointlessness of Pareto: Carrying Coase Further*, 100 *YALE L.J.* 1211 (1991) (noting that few policies are Pareto efficiency, and considering bases for evaluating policies that are not Pareto efficient).

⁹⁸ See *supra* Part I.B (considering the accuracy assumption); *infra* Part II.A.2 (considering the manipulation and ideological bias assumptions).

ultimately information markets' usefulness depends on difficult evaluations of the extent to which they are and might be true.

The immediate question is why narrowing agency decisionmaking to normative questions would be useful. It might seem that if we set aside the possibility that information markets achieve better information aggregation than any comparably costly individual decisionmaker, it should not matter whether the government uses them. On this argument, information markets represent a privatization of government functions without resource savings and should thus be a matter of indifference. The problem with this argument is that it focuses exclusively on different information aggregation technologies rather than on how agents might use or misuse such technologies. The danger is that administrative agencies might make factual assessments in order to support decisions that they would like to make for normative reasons.⁹⁹ Deprived of the ability to announce misleading predictions, agency officials would need to defend their decisions with normative arguments alone. To the extent that these normative arguments are unpersuasive, agency officials might suffer the reputational costs that a decisionmaker bears when observers view a decision as nothing more than an *ipse dixit*,¹⁰⁰ as well as face an increased possibility of skeptical judicial review.¹⁰¹

Although information markets that generate predictions of objective variables are only precursors to subsequent normative agency decisions, they can help discipline the process of making those decisions. Another analytical tool, cost-benefit analysis, has often been identified as achieving a similar type of discipline. Commentators have emphasized that cost-benefit analysis helps to reduce cognitive errors and the effect of political factors, such as interest groups

⁹⁹ Joseph Smith and Emerson Tiller have shown empirically that judges who wish to achieve results in particular cases may act strategically in choosing an instrument that will achieve the policy goal, for example in deciding whether to reject a regulation on procedural or substantive grounds. See Joseph L. Smith & Emerson H. Tiller, *The Strategy of Judging: Evidence from Administrative Law*, 31 J. LEGAL STUD. 61 (2002). Smith and Tiller note: "Strategic judges anticipate the likely responses of other players and make choices that will ultimately maximize their own policy preferences under the given constraints. To the extent that the choice of a judicial instrument affects the ability of others to reverse the court, it becomes a strategic variable in the court's decision." *Id.* at 63. The point applies as much to agencies as to judges. Agency officials who wish to ensure that a policy is enacted may justify their decisions on predictions that are difficult for courts to dispute if they are concerned that courts might be skeptical of the normative considerations that in fact motivated the decision.

¹⁰⁰ Some commentators doubt that reputation figures prominently in the judicial utility function. See Richard A. Posner, *What Do Judges and Justices Maximize? (The Same Thing Everybody Else Does)*, 3 SUP. CT. ECON. REV. 1 (1993). To the extent that judges do care about their reputations, however, judges presumably do not wish to appear results-oriented. See, e.g., Richard A. Posner, *The Meaning of Judicial Self-Restraint*, 59 IND. L.J. 1, 8 (1983/1984) (recognizing that "result-oriented" is used as a "pejorative").

¹⁰¹ See *infra* notes 238-242 and accompanying text (considering the possible use of information market predictions in judicial review).

or ideology, on legal policy. The parallel is incomplete, as agency officials may have some ability to control the numbers produced by cost-benefit analysis that they would not have with information markets, but I will return to that problem later by showing how information markets might be combined with cost-benefit analysis to reduce such control.¹⁰² The comparison, however, is useful, as information markets can increase predictive assessments objectivity in much the same way that cost-benefit analysis can increase normative assessments objectivity, as the following subsections explain.

a. Cognitive Errors

The simplest defense of cost-benefit analysis is that it prevents bad policies, which are policies whose costs, if enacted, would exceed their benefits.¹⁰³ Such a defense invites the question of *why* governmental officials would choose bad policies in the absence of cost-benefit analysis, a question to which Cass Sunstein poses one possible answer in *Cognition and Cost-Benefit Analysis*.¹⁰⁴ Sunstein argues “that cost-benefit analysis is best defended as a means of overcoming predictable problems in individual and social cognition.”¹⁰⁵ By “predictable problems,” Sunstein means the heuristics that cognitive psychologists have identified as producing systematic biases in human decisionmaking,¹⁰⁶ as well as the social dynamics that can cause group decisionmaking to err.¹⁰⁷ Cost-benefit analysis’s methodology enforces a rigor that can help decisionmakers overcome such pitfalls, pointing them to what is really at stake in a decision.

As the most significant example, Sunstein cites the availability heuristic,¹⁰⁸ which refers to the tendency of people to think that events are more likely if they can recall past examples of such events. Writing with Timur Kuran, Cass Sunstein has previously emphasized the danger

¹⁰² See *infra* Part III.

¹⁰³ For a sophisticated exposition of what they term the “conventional view,” see Matthew D. Adler & Eric A. Posner, *Rethinking Cost-Benefit Analysis*, 109 YALE L.J. 165, 176-94 (1999) (offering a more sophisticated version of such a defense).

¹⁰⁴ Cass R. Sunstein, *Cognition and Cost-Benefit Analysis*, 29 J. LEGAL STUD. 1059 (2000).

¹⁰⁵ *Id.* at 1060.

¹⁰⁶ See generally JUDGMENT UNDER UNCERTAINTY: HEURISTICS AND BIASES (Daniel Kahneman et al. eds, 1982) (collecting many of the seminal cognitive psychology articles identifying such biases).

¹⁰⁷ Sunstein himself has done considerable work on the dynamics of group decisionmaking. See, e.g., Cass R. Sunstein, *Deliberative Trouble? Why Groups Go to Extremes*, 110 YALE L.J. 71 (2000); Cass R. Sunstein et al., *Predictably Incoherent Judgments*, 54 STAN. L. REV. 1153 (2002); Cass R. Sunstein, *What’s Available? Social Influences and Behavioral Economics*, 97 NW. U. L. REV. 1295 (2003).

¹⁰⁸ Sunstein, *supra* note 104, at 1065-66.

that this heuristic will lead to “availability cascades,”¹⁰⁹ a vicious cycle in which an event leads individuals to overestimate a risk, in turn affecting public discourse, which then makes exacerbates the initial overestimation.¹¹⁰ As an example, Sunstein cites the Love Canal episode, in which residents’ concerns about environmental contamination from a toxic waste dump snowballed, eventually leading to mass relocations,¹¹¹ even though there was no valid scientific evidence validating their concerns.¹¹² The ultimate effect, Sunstein suggests, was even more dramatic, leading to passage of the Superfund statute,¹¹³ which critics have suggested is one of the most expensive and least effective environmental statutes,¹¹⁴ given the relatively small risk posed by toxic waste dumps.¹¹⁵

Cost-benefit analysis might have thwarted the Love Canal availability cascade by forcing governmental officials to make explicit estimates of risk based on scientific data.¹¹⁶ An information market might have a similar effect. For example, an information market might have been invoked to predict cancer rates in various communities. A prediction that the cancer rates in communities near dumps would be comparable to the rates in communities lacking dumps would provide an objective datum indicating that there was an overreaction. Information markets might also be used for other availability cascades that Sunstein has perceived, such as the public overreaction that he believes occurred to the September 11, 2001, terrorist attacks.¹¹⁷ In that case, for example, an information market might have been used to predict future deaths from terrorism, possibly conditional on different governmental policies. Such a market might have produced an objective datum that the risk was small, or it might have suggested that the risk was quite large, especially once the danger of nuclear terrorism is taken into account.¹¹⁸ At the same

¹⁰⁹ Kuran & Sunstein, *supra* note 19.

¹¹⁰ *Id.* at 716 (“Public discourse shapes individual risk judgments, risk preferences, and policy preferences; and the reshaped personal variables then transform the public discourse that contributed to their own transformations.”).

¹¹¹ *Id.* at 695.

¹¹² *Id.* For a different view of the episode, see ADELIN GORDON LEVINE, LOVE CANAL: SCIENCE, POLITICS, AND PEOPLE (1982).

¹¹³ 42 U.S.C. §§ 9601-9675 (2000).

¹¹⁴ STEPHEN BREYER, BREAKING THE VICIOUS CYCLE: TOWARD EFFECTIVE RISK REGULATION 38 (1993).

¹¹⁵ For an analysis of the Superfund program, see ANALYZING SUPERFUND: ECONOMICS, SCIENCE, AND LAW (Richard L. Revesz & Richard B. Stewart eds., 1995).

¹¹⁶ See Sunstein, *supra* note 104, at 1065-66 (“Cost-benefit analysis is a natural corrective [to availability cascades], above all because it focuses attention on the actual effects of regulation, including, in some cases, the existence of surprisingly small benefits from regulatory controls.”).

¹¹⁷ See CASS R. SUNSTEIN, THE COST-BENEFIT STATE: THE FUTURE OF REGULATORY PROTECTION 26 (2002).

¹¹⁸ See, e.g., MARTIN J. REES, OUR FINAL HOUR: A SCIENTIST’S WARNING (2003) (arguing that nuclear war and other potential disasters present a strong possibility of destroying civilization).

time, an information market preceding September 11 might have suggested that insufficient attention was being paid to the dangers of terrorism.

A significant caution is in order here. Just because an information market exists does not mean that it will stop an availability cascade, even assuming that it makes an accurate prediction. Public officials might well ignore the results of an information market, downplaying any data that seem inconsistent with public concerns. This is, however, a limitation of cost-benefit analysis and other objective guides to policy as well.¹¹⁹ Any methodology for improving governmental decisionmaking can be successful only to the extent that it is followed. Administrative agency officials might seek to institute information markets as a way of preventing availability cascades from emerging, but officials might also like to play on public fears, either to seem responsive or because they might prefer a greater amount of risk regulation than a well-informed public would choose.

There are at least two possible routes that could lead the government to make decisions consistent with approaches like cost-benefit analysis and information markets. One is general acceptance. Such acceptance has eluded cost-benefit analysis,¹²⁰ perhaps in part because of legitimate claims that it favors particular values over others.¹²¹ Even if information markets do prove to be accurate and objective, the public may be skeptical of their results. Their relative lack of prominence makes acceptance of their predictions particularly unlikely, although in the long term, their focus on predictions rather than on values may encourage their acceptance, especially if they prove prescient and their accuracy is publicized. The second route is to enact a set of rules requiring that administrative agency decisionmakers engage these methodologies and factor them into their decisions.¹²² Attempts to institutionalize cost-benefit analysis have been at least moderately successful,¹²³ though the courts have held that some agencies cannot use cost-benefit

¹¹⁹ As Kuran and Sunstein emphasize, a blue-ribbon panel on Love Canal “endorsed none of the reports of serious health effects,” but “its evaluation had no appreciable influence on subsequent events.” Kuran & Sunstein, *supra* note 19, at 695.

¹²⁰ Sunstein himself has documented a recent case in which a cost-benefit analysis recommendation was controversial. See Cass R. Sunstein, *The Arithmetic of Arsenic*, 90 GEO. L.J. 2255 (2002).

¹²¹ See *infra* notes 288-289 and accompanying text.

¹²² See, e.g., *id.* at 746-61 (considering a variety of correctives to prevent availability cascades).

¹²³ The Office of Management and Budget has standardized cost-benefit analysis in Circular A-94. See OFFICE OF MGMT. & BUDGET, CIRCULAR A-94 (Oct. 29, 1992), *available at* Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs, 57 FED. REG. 53519-02 (Nov. 10, 1992) [hereinafter CIRCULAR A-94].

analysis.¹²⁴ I will consider the role that information markets might play in judicial review below.¹²⁵

b. Politics

A separate virtue of cost-benefit analysis is that it may allow for easier monitoring of agencies by their principals, the President and Congress. Eric Posner has argued that cost-benefit analysis may help to overcome informational advantages that agencies have over the President and Congress by forcing an agency to convey information about projects.¹²⁶ Cost-benefit thus produces not only better projects from the perspective of social welfare, but also better projects from the perspective of the President,¹²⁷ who Posner assumes can veto agency plans.¹²⁸ Perhaps surprisingly, Posner's model suggests that cost-benefit analysis will lead to enactment of more regulations, as the President will trust the agency more once it must back up its recommendations with cost-benefit analysis and the agency in turn will propose more projects that the President otherwise would have rejected.¹²⁹ Posner extends the analysis to situations in which Congress is the principal or Congress and the President are both principals, concluding once again that cost-benefit analysis helps the agency transmit information.¹³⁰

Critical to Posner's model is that cost-benefit analysis provides an *objective* datum about the efficiency of projects.¹³¹ Posner recognizes the possibility of agencies "bias[ing] the outcomes in their own favor" to move policy closer to their ideal level of regulation,¹³² but his model depends at least on the possibility that cost-benefit analysis might provide a means of

¹²⁴ See *Whitman v. Am. Trucking Assocs., Inc.*, 121 S. Ct. 903, 909 (2001) (finding that the Environmental Protection Agency cannot consider costs under certain sections of the Clean Air Act).

¹²⁵ See *infra* Part II.C.

¹²⁶ Eric A. Posner, *Controlling Agencies with Cost-Benefit Analysis: A Positive Political Theory Perspective*, 68 U. CHI. L. REV. 1137 (2001). Posner explains that "Agency's informational advantage is due to its institutional expertise." *Id.* at 1149.

¹²⁷ *Id.* at 1156-57.

¹²⁸ *Id.* at 1148.

¹²⁹ *Id.* at 1157.

¹³⁰ *Id.* at 1171.

¹³¹ Posner complicates this assumption by suggesting that the agency can "fake" a plausible cost-benefit analysis for some additional cost. *Id.* at 1161. Although Posner does not elaborate the nature of this cost, it might be thought of as the reputational hit that the agency may suffer as a result of faked analyses. Posner does note that if this cost is relatively low, the President might not believe the agency's cost-benefit analysis. *Id.*; see also Eric A. Posner, *Cost-Benefit Analysis as a Solution to a Principal-Agent Problem*, 53 ADMIN. L. REV. 289, 292 (2001) ("If the agency's cost-benefit analysis is based on implausible data, or uses flawed calculations, then the President has the necessary information for disciplining the agency as well as for rejecting the project.").

¹³² Posner, *supra* note 126, at 1178-79.

overcoming an informational asymmetry and allowing a principal to monitor an agent. Whether cost-benefit analysis in fact provides such a function,¹³³ information markets would, assuming that they are generally objective.¹³⁴ Just as cost-benefit allows an agency to reassure the principal that a particular project moves regulation in the principal's preferred direction, so too could information markets assure the principal that the regulation will be an improvement based on the principal's policy preferences. For example, a cost-benefit analysis showing that a regulation is expected to save many lives could help assure a President who is more skeptical of interventionist government than agency officials that there is indeed reason to believe that the regulation will indeed produce considerable bang for whatever bucks are being expended.¹³⁵

Posner's model helps to underscore a straightforward point that is as applicable to information markets as it is to cost-benefit analysis: Any tool that provides an objective reflection on an agency's decision makes it easier to separate the claimed and real justifications for agency decisions. As a result, agencies are less likely to pursue policies when the real justifications would be insufficient to persuade those who have some power over the agencies. Yet agencies will be more able to pursue policies in which the claimed justification would in fact be sufficiently persuasive, because information markets can reduce the informational asymmetry and thus help to reduce distrust.

Perhaps the most significant implication is that information markets may help limit the role of interest groups. Posner notes that "interest groups sometimes help government principals control agencies by disclosing information to the principals,"¹³⁶ specifically by announcing endorsement of or opposition to a particular regulation. Although such signals may have an effect similar to cost-benefit analysis in reducing informational asymmetry, a consequence is that interest groups may bias agencies' project choices.¹³⁷ By neutralizing the informational advantage, cost-benefit analysis eliminates the potential for such bias. Information markets can do much the same thing.¹³⁸ But they can also do something that cost-benefit analysis cannot do.

¹³³ See *infra* text accompanying note 277.

¹³⁴ See *infra* Part II.A.2.

¹³⁵ As conceived so far, the information market measures only the bang and does not consider the buck. *But see infra* Part III (explaining how predictive cost-benefit analysis would allow for an overall assessment of legal policy).

¹³⁶ Posner, *supra* note 126, at 1174.

¹³⁷ *Id.*

¹³⁸ A variety of reforms of administrative agencies might limit outside influence on agency decisionmakers. One danger, however, is that reforms that seek to insulate decisionmakers might effectively cut out only public participation, leaving industry

As Posner acknowledges, “interest groups retain a hidden influence on the location of the other players’ ideal points,” but cost-benefit analysis does not change this.¹³⁹ Information markets may frustrate agency officials’ ability to credibly claim illusory benefits from policies that are in reality special interest giveaways. Information markets can thus lead agency officials concerned with their credibility to act in a manner more consistent with the ideal points that they would prefer to project to the public, who constitute the principal that they ultimately may care most about.

2. *Are Information Markets Objective?*

The affirmative case for information markets’ objectivity is that they provide financial incentives for honest predictions. When individuals are asked to make predictions in the absence of a financial incentive, they might announce what they wish would happen or what they would like to be seen as believing rather than what they truly believe.¹⁴⁰ Financial incentives aimed at accurate prediction may be sufficient to override these factors. Even if these incentives are inadequate for some traders, other traders will have an incentive to trade against these financially unmotivated traders, whose willingness to lose money for expressive purposes will ordinarily be limited. Depending on the amount of money at stake, traders may even have an incentive to gather information not previously publicly available to better inform their trading decisions.

There are, however, two significant caveats to this proposition, to be considered in each of the following two subsections. The first danger is that some market participants might affirmatively attempt to manipulate an information market because they have extrinsic reasons, including possibly financial ones, to care about the results of that market. If attempts at manipulation are successful, then the information markets may be biased, unless attempts by opposing factions happen to cancel each other out. The second danger is that the demographics of the trading population may influence results. For example, if traders happen to be on average conservative, they may have honestly different beliefs, and thus predictions, with respect to certain questions than those of a more demographically diverse group. It may not be enough for information markets to produce honest predictions if the honest predictions are those of a

and other special-interest groups in a relatively stronger position. See, e.g., David Dana, *Setting Environmental Priorities: The Promise of a Bureaucratic Solution*, 74 B.U. L. REV. 365, 379 (1994) (noting the possibility of “uneven depoliticization”).

¹³⁹ Posner, *supra* note 126, at 1175.

¹⁴⁰ See *infra* note 259 and accompanying text (discussing “position taking”).

homogeneous group. Although there are theoretical reasons to believe that both these dangers will be small, ultimately empirical evidence and analysis is necessary.

a. Market Manipulation

The Iowa Electronic Markets have appeared relatively immune to attempts at market manipulation, though an occasional unexplained price spike suggests some reason for concern.¹⁴¹ A possible explanation for this, however, is that the Markets' only effects are on the participants themselves. Imagine, however, a world in which a large number of voters based their actual election votes on the outcome of the Iowa Electronic Markets. This is somewhat fanciful,¹⁴² but the hypothetical allows for consideration of what might occur if the government were to use information markets as factors in their own decisionmaking. In the election hypothetical, there would then be many individuals, from the candidates themselves to special interest groups, that might be willing to accept some trading losses if their trading activity had an effect on the market price and thus on the election. Similarly, in an information market used as a basis for agency decisionmaking, anyone who might have an interest in the ultimate decision might transact with the primary goal of manipulating market prices, even if such manipulation resulted in financial loss for the trading party.¹⁴³

If such manipulation were possible, particularly if it could have more than a trivial effect on market prices, then the case for the objectivity of information markets is reduced. Market manipulation is likely to be impossible, however, whenever market participants can detect that the manipulation has changed asset prices in an objectively unjustifiable way. Suppose, for example, that Bill Gates seeks to pour money into the securities corresponding to the Green Party in a hypothetical version of the Iowa Electronic Markets, attempting to drive the price of such securities up. This would create an immediate opportunity even for arbitrageurs with no knowledge of the security's true value, because the price of the various securities would add up to more than \$1. By investing \$1, an arbitrageur could receive one of each security, selling them

¹⁴¹ See *supra* note 39 and accompanying text..

¹⁴² Some voters prefer to vote for the candidate whom they expect to win, a phenomenon called a "bandwagon effect." See, e.g., Roger B. Myerson & Robert J. Weber, *A Theory of Voting Equilibria*, 87 AM. POL. SCI. REV. 102, 102 (1993). If the Iowa Electronic Markets became widely recognized as accurate predictors of election results, supplanting polls in press assessments of elections, then market prices might affect some voters' decisions.

¹⁴³ Hanson characterizes the incentive by noting that interested parties might seek to trade according to criteria other than fundamental security values because "[d]ecision gains might outweigh trade losses." Hanson, *supra* note 68, at slide 13.

all at market prices, and this activity would drive down the price of the Republican and Democratic securities. Fundamental value traders, however, would recognize that the Republican and Democratic securities were undervalued, and they would thus purchase these securities, driving their prices back up.¹⁴⁴ This market activity would attract other arbitrageurs and fundamental value traders, and eventually even Gates would run out of money.

Gates, in this hypothetical, is effectively offering to make obviously losing bets. If Gates were to announce that he would give away money to anyone who wanted it, then his money would disappear quickly, and if he announced that he would bet anyone and everyone \$1 that the Green Party would win the Presidency, it would disappear with almost equal speed. An information market structured like the Iowa Electronic Markets ensures the same dynamic. Some additional protection against manipulation would be needed in an information market in which the number of securities is constant. If, for example, there were a fixed number of Green Party securities, then Gates could buy all of these securities, sell one of them to a friend for \$1, and then refuse to enter into any additional transactions. There are a number of possible solutions to this problem, however. The market scoring rules approach described above offers a simple one.¹⁴⁵ By allowing anyone to displace the current predictor with a new prediction, this approach would in effect allow market participants to take bets against anyone repeatedly entering unsupportable predictions. It should thus similarly allow the market to counter manipulation fairly quickly.

Quickly, however, may not be fast enough in certain cases. We have seen at least one blip in the Iowa Electronic Markets data. What if such a blip occurred in an information market used to support a decision at exactly the moment that the market terminated? That could happen for reasons other than pure coincidence. If there is a fixed point in time at which a snapshot of market prices is to be taken for decision purposes, then potential manipulators would have an incentive to wait until just before this point in time to act. There is, however, a potential antidote to this problem. Instead of setting a fixed point at which the market would end, the sponsors of an information market might leave the exact end time ambiguous, even if there is some period in

¹⁴⁴ If it were possible to sell short on the Iowa Electronic Markets, the fundamental value traders also might sell the Green Party short. Current rules on the Iowa Electronic Markets prohibit short sales. See <http://www.biz.uiowa.edu/iem/faq.html#shortsales> (last visited June 17, 2003).

¹⁴⁵ See *supra* Part I.C.2.

which the market is assured to be open.¹⁴⁶ Moreover, the market sponsor might provide that the market will close only after some period of time during which there has been relatively little market activity potentially consistent with market manipulation. Such activity might be defined as an unusually high volume of trading or unusually large shifts in market price.

This analysis provides strong reason to believe that as long as it is clear that a trader's activity has moved a security price away from its fundamental value, market forces should respond. Robin Hanson reaches this conclusion as well,¹⁴⁷ but he properly recognizes that this will not always be the case. In information markets for which asymmetric information is relevant, traders will be uncertain whether trades reflect genuine insider information or a trader's attempts at market manipulation. If the traders underestimate the extent to which a trade is based on an attempt at manipulation, then that attempt may partly succeed.¹⁴⁸ On the other hand, Hanson argues, if traders overestimate the extent to which a trade is based on a desire to bias, then the attempt may backfire.¹⁴⁹ Allowing interested parties to trade, Hanson argues, will not increase the danger of successful manipulation, as long as market participants know which parties are potentially interested, because other market participants remain free to conclude that an interested trader's activity reflects nothing about the security's underlying asset value.¹⁵⁰ Moreover, he notes,¹⁵¹ interested traders provide a benefit to the market in the form of increased liquidity, giving a greater incentive for fundamental value traders to engage in research because they can now use such research to bet against those attempting to manipulate the market.

¹⁴⁶ For this to work, the market sponsors would not be able to set a final time by which the market would necessarily close; otherwise, just before that moment, manipulation might be possible. Instead, the sponsors might assign a "half life" to the market beginning at some point in time, so that in each succeeding period, there is only some relatively small probability that the market will close regardless of how many such periods have occurred. See Michael Abramowicz, *Cyberadjudication*, 86 IOWA L. REV. 533, 568-69 n.93 (2001) (describing such an approach in an analogous context). Even with such a system, if the half life is sufficiently short, there is a very high probability that the relevant market will close quickly.

¹⁴⁷ See Hanson, *supra* note 68, at slide 18.

¹⁴⁸ See *id.* at slides 16-21.

¹⁴⁹ *Id.* This point is not immediately clear from the slides themselves, but Hanson's verbal explanation of the slides make clear that this is his view. Although Hanson's view of this seems analytically sound, this scenario seems unlikely to be a significant concern in reality. In Hanson's model, each trader has some private information on both the underlying asset value and the desire to bias, so if the traders overestimate the desire to bias, then the traders will misestimate the manipulator's private information on the underlying asset value. This effect is substantial, however, only if traders frequently have mixed motives in participating in information markets, both to achieve trading gains and to affect decisions. In the typical case, if traders concluded that someone was trying to manipulate the market, they would draw no inference at all about the underlying asset value from the attempt.

¹⁵⁰ See *id.* at slide 22.

¹⁵¹ *Id.*

If market participants could identify interested traders, Hanson is probably correct that allowing trading will be harmless and even potentially beneficial. The real danger, though, is that market participants will not be able to tell whether a certain trader is in fact an interested party. Market participants would seek to estimate what portion of trading is likely to reflect that of interested parties, and thus on average interested traders should not be able to change the market price, but in any given case, market participants might underestimate or overestimate the portion of trading attributable to that group. Indeed, the trick for interested parties would be to trade more than market participants expect, producing a vicious expectations cycle leading to ever more trading by interested parties.¹⁵² In sum, the possible existence of interested parties does not seem likely to produce *systematic* bias, but there is concern that trading by such parties could add considerable noise to market predictions.

This is the gloomy view, and the absence of expected systematic bias in favor of interested parties makes it not so gloomy after all. There is, in any event, a more optimistic view, which suggests that information markets will be relatively efficient at distinguishing between trading based on information and trading based on desires to bias. Some traders will presumably earn reputations for acting reliably on the basis of private information, which may include private information that they have paid other parties to access. In a world with many information markets, the market will trade against traders lacking such reputations, and they will thus be unable to influence the market. Perhaps initial experiments with information markets will not be sufficient to allow such reputations to build, but there would be strong long term incentives for traders to build and monitor reputations. Traders who fail to produce reputable information will often be unsuccessful in convincing others.

Whether the gloomy or the optimistic view is the appropriate one may well depend on the context. For example, one goal of FutureMAP was to consider the extent to which information markets may be able effectively to combine both public and top secret information.¹⁵³ If all

¹⁵² A more familiar example of expectations leading to a vicious cycle occurs with inflation policy. Inflation depends in part on expectations of inflation, yet policymakers may wish to generate surprise inflations as short term spurs to economic growth. A central bank's desire to produce just a little bit more inflation than is expected can lead quickly to very high levels of inflation. See, e.g., Robert Barro & David Gordon, *Rules, Discretion and Reputation in a Model of Monetary Policy*, 12 J. MONETARY ECON. 101 (1983) (solving for the inflation rate that results in a rational expectations equilibrium); Finn E. Kydland & Edward C. Prescott, *Rules Rather than Discretion: The Inconsistency of Optimal Plans*, 85 J. POL. ECON. 473 (1977) (explaining how binding rules can allow an escape from this expectations problem).

¹⁵³ See <http://www.darpa.mil/iao/FutureMAP.htm> (last visited June 17, 2003) (web site no longer available) ("A market that addresses defense-related events may potentially aggregate information from both classified and unclassified sources.").

participants in a market have access to the same top secret information, or even if there is known to be a critical mass of traders who have such access as well as information about who else has such access, such markets may work well. But in a market in which individual traders may have access to top secret information that they can not convey to other traders, market participants will give some credence to any trades by someone who may have such information. Market participants also would try to estimate the extent to which that individual may be biased because of policy outcomes the individual hopes to influence, but once again this could reduce the accuracy of market predictions.

By contrast, in a market that depends primarily on public information, or even on private information that market participants could share with reputable parties who could verify it, traders seem likely to place only minimal weight on the views of traders without reputations for objectivity if they do not produce persuasive analyses in support of their positions. The only danger here is that some traders without reputations might determine that they are better off trading repeatedly on any information they obtain rather than trading on that information once and then revealing the information later. That is, a trader perversely could benefit from others' skepticism about whether the trade is really supported by information, because the skepticism allows the trader to repeatedly trade against the skeptics. This may be unlikely, because traders would have some incentive to improve their reputations by showing that their trades are consistently information-based.¹⁵⁴ If this strategy did emerge, however, the market would grant at least some credence to traders without reputations, while still trying to estimate the total amount of bias, because any trade might reflect this strategy. Thus, if traders might have an incentive to suppress the information justifying their trades, market accuracy once again will be reduced, and traders who manipulate the market more than they are expected to will benefit.

There is, however, a relatively straightforward, if counterintuitive, approach to encouraging traders to reveal the information on which they base their decisions. The market could be divided into two phases, a first phase and a very short second phase, long enough for

¹⁵⁴ The advantage of gaining a reputation is that a trader with a reputation can move the market, since that trader's actions will not lead to responses by other skeptical market participants. Thus, the trader can trade on a relatively small amount of information and then exit the market quickly by selling shares. A trader without a reputation must wait until the close of the market to profit, and that is a risky strategy. Relying on reputation to produce accuracy, however, is itself a risky public policy, for example because it sometimes may be in traders' interests to deplete their reputational assets in exchange for one-time trading gains. Cf. Frank Partnoy, *Barbarians at the Gatekeepers?: A Proposal for a Modified Strict Liability Regime*, 79 WASH. U. L.Q. 491, 498 (2001) (considering the incentives of corporate gatekeepers, such as attorneys and auditors, to deplete reputation).

market participants to counter any sudden price swings but not long enough for any substantial analysis. The closing price of the second phase, rather than the event ultimately at issue, would determine the liquidation values for all trades made during the first phase.¹⁵⁵ In effect, this is a double information market, using one information market to predict the results of a second, which in turn predicts the result of interest. While the goal of a trader in the first phase is only to predict the outcome of the second, the knowledge that traders in the second will be seeking to predict the ultimate outcome means that traders in the first phase will seek to predict the ultimate event. The market sponsor could thus provide that the closing price of the securities in the first market would be the basis for any subsequent decisions. While the second phase may appear redundant, the second phase performs a useful function in disciplining pricing and behavior in the first phase. Just the threat of discipline in the second phase should be sufficient to curtail manipulation efforts, and minimal trading would be expected to occur during this phase.

This two phase approach admittedly would make the market a bit less intuitive to participants, and it might well be unnecessary. The approach does, however, give market participants incentives to reveal information supporting their trades, and eliminates the strategy of withholding information to increase trading profits.¹⁵⁶ Those who fail to convince other market participants that their first phase trades are justified will be unable to profit from the information on which they have traded, because the second phase, rather than the actual event, determines first-phase profits. A trader, meanwhile, would not be able to withhold information for strategic purposes in the first phase, because the second phase is too short for information to be processed. Meanwhile, there would be no reason for someone with actual information to wait until the second phase to trade on that information, because it would be impossible to trade repeatedly on information in the second phase. The strategy of taking advantage of skepticism to increase

¹⁵⁵ Consider a market to predict the population of New York in 2010. Suppose that the last prediction at the end of the first phase was 9.0 million, the last prediction at the end of the second phase was 9.1 million, and the eventual population turned out to be 10.0 million. The predictors in the second phase would be rewarded or penalized in the usual way, with the 10.0 million number used in the scoring rule. The predictors in the first phase would be rewarded or penalized, however, as if the population turned out to be 9.1 million. These rewards or penalties could be imposed immediately after the end of the second phase.

¹⁵⁶ The two-phase approach also has a side benefit, reducing the risk of participating in the first phase of the market, because there will no longer be risk associated with real world randomness. A trader who makes a probability estimate of an event's occurrence in a conventional information market bears a risk because the event will either occur or it will not. But a participant in the first phase is simply predicting the probability estimate of the participants in the second phase. Any participants in the second phase would still bear some risk, but there likely would be little trading in this phase. The purpose of the second phase is to discipline the first, but given the existence of this disciplining mechanism, there is no incentive to wait to trade until the second phase.

trading profits will not work in the second phase, for that strategy depends on market participants thinking that there is some chance that a trade is really an attempt to manipulate the market. No one would have an incentive to manipulate the market in the second phase, because the ultimate decisionmaker considers only the closing prices at the end of the first phase.

By promoting sharing of information, the two-phase approach seems likely to increase information markets' accuracy both directly and indirectly, because the sharing will make it easier to identify traders who are trading without a sufficient reason for doing so. Regardless of whether this method is used, it is worth emphasizing that information markets will not on average be biased in favor of those who have incentives to manipulate the market, because other participants are just as likely to overestimate as to underestimate the extent to which they try to engage in manipulation. Ultimately, the extent to which the possibility of manipulation reduces the accuracy of markets in which people have incentives to manipulate is an empirical question. An experiment involving a large number of securities could be used to address this question; if some individuals are given incentives to care about the prices of some randomly selected securities. The experiment could also examine whether it matters if an incentive to manipulate is conveyed only privately to one trader or announced publicly. Until such experiments verify that there is some means of structuring information markets to minimize problems associated with manipulation, government agencies will need to be quite cautious in relying on the results of such markets. The analysis in this section, however, suggests that there is considerable reason for optimism.

b. Unrepresentative Decisionmakers

An additional potential objection to relying on the results of information markets for governmental decisionmaking in particular is that such reliance means that an unrepresentative group of traders is in effect making decisions for society as a whole. This objection may come in two distinct flavors. The first is that even if traders were excellent predictors, relying on traders is democratically illegitimate because the traders are not democratically elected. This objection identifies an instinctive discomfort that democratic theorists have to private entities governing the public, but the discomfort comes from the concern that private entities may act on behalf of

their own interests rather than those of society as a whole.¹⁵⁷ The discomfort thus seems misplaced if indeed market predictions are unbiased; if democratically selected officials place weight in their decisionmaking on unbiased information, regardless of whether the producers of the unbiased information engage in voting behavior that suggests that they wish the unbiased information they produce were different, it seems unproblematic. At least this is true as long as attempts to manipulate the market fail.¹⁵⁸

The second flavor of the objection concedes this counterargument but worries that the results of information markets might include an unintentional bias. The argument is that even if a market is not manipulated by any individual participant or set of participants, perhaps the identity of the participants as a whole might affect the market outcome. Information market participants are not likely to be representative of the population as a whole. Aside from being more expert in the relevant area, they might be wealthier or have systematically different political views. Suppose, for example, that the participants in the hypothetical airline safety regulations information market tend to be individually more sympathetic to market-oriented solutions than airline safety experts in general or than the population as a whole. One might then worry that predictions by this market might be systematically different from predictions in a hypothetically more representative government because the decisionmakers might generally predict a smaller number of **lives saved** than a more representative group.

The limited experimental evidence produced by the Iowa Electronic Markets suggests that this is not likely to be a problem. Bias in an election market would be consistent with what is referred to by cognitive psychologists as the “false consensus effect,” in which individuals overestimate the number of others who share their views.¹⁵⁹ Robert Forsythe et al. studied the participants in the 1988 Presidential vote share market and collected, through polls, information on traders’ political preferences.¹⁶⁰ The market performed quite accurately, with election eve

¹⁵⁷ See, e.g., Jody Freeman, *The Private Role in Public Governance*, 75 N.Y.U. L. REV. 543 (2000) (arguing that while privatization increasingly involves private firms in making public decisions, “alternative accountability mechanisms . . . can allay our concerns about the particular risks posed by arrangements of public and private actors, while capitalizing on their capacities”).

¹⁵⁸ See *supra* Part II.A.2.a.

¹⁵⁹ See, e.g., Clifford E. Brown, *A False Consensus Bias in 1980 Presidential Preferences*, 118 J. SOC. PSYCHOL. 137 (1982); Thomas Gilovich, *Differential Construal and the False Consensus Effect*, 59 J. PERS. & SOC. PSYCH. 623 (1990); Lee Ross et al., *The False Consensus Effect: An Egocentric Bias in Social Perception and Attribution Processes*, 13 J. EXPERIMENTAL SOC. PSYCHOL. 279 (1977).

¹⁶⁰ Forsythe et al., *supra* note 40.

predictions coming within about one percentage point,¹⁶¹ even though the traders consisted entirely of individuals affiliated with the University of Iowa, including a disproportionate number of business students, producing a trading population that identified as Republican and favored George Bush somewhat more often than the population as a whole.¹⁶² On the whole, traders tended to buy shares associated with their preferred candidates.¹⁶³ The authors, however, isolated one group of traders, whom they called marginal traders, who exhibited little to no judgment bias. These traders were identified by the fact that they placed limit orders at prices slightly different from the market price, while other traders placed either market orders or limit orders at prices far from the market price.¹⁶⁴ In effect, these marginal traders determined the market price and ensured that judgmental bias did not affect the market,¹⁶⁵ even though these marginal traders were even less demographically representative than traders as a whole.¹⁶⁶

Though cause for optimism, there are at least two reasons that this study is not a conclusive repudiation of the possibility that traders' demographic characteristics might affect market prices. First, the closeness of election-eve predictions to the actual result does not disprove the possibility of bias; perhaps the market prediction was slightly biased but by chance the bias coincided with an election day surprise.¹⁶⁷ An experimental design that might more clearly control the effect of traders' demographic composition would involve implementation of a number of simultaneous markets, for example predicting vote shares in large numbers of congressional races, with the composition of those permitted to trade in these markets selected varied randomly across markets.¹⁶⁸ A demonstration that groups of Democratic traders were no more likely to be biased toward Democrats than other groups would bolster the Forsythe findings. If, however, it turns out that group composition matters, then the further question, also potentially testable, is to what extent group composition matters in open markets in which

¹⁶¹ *Id.* at 1148.

¹⁶² *Id.* at 1146 (noting, for example, that forty-six percent of traders identified themselves as Republican, compared with thirty-one percent in a national poll).

¹⁶³ *Id.* at 1155-56.

¹⁶⁴ A market order is an order to buy or sell a security at the best available price, while a limit order is an offer to buy or sell only if the price reaches a certain level. *See generally* Puneet Handa & Robert A. Schwartz, *Limit Order Trading*, 51 J. FIN. 1835 (1996) (explaining the rationale for limit orders).

¹⁶⁵ Forsythe et al., *supra* note 40, at 1156-57.

¹⁶⁶ *Id.* at 1158 (noting that all marginal traders were male and that 60 percent favored Bush after the third debate). The marginal traders tended to have higher investments than other traders.

¹⁶⁷ Forsythe et al. recognize this possibility, noting that they might have just been "lucky." *Id.* at 1156-57.

¹⁶⁸ With a large number of markets, thin trading is potentially a concern. *See supra* Part I.A.2.

anyone may participate.¹⁶⁹ One possibility, supported indirectly by the experiments indicating that a small percentage of informed traders is sufficient,¹⁷⁰ is that as long as there is some minimal participation by a range of individuals, the results do not depend on the traders' exact composition.

Second, the conclusion that some trader diversity will ensure objectivity is not necessarily generalizable to other forms of information markets, particularly those involving complex models that individuals with different political views might view with different levels of acceptance. Traders may recognize that others will have different political views but not give credence to alternative models relevant to the immediate prediction. It is one thing for traders to recognize that their preferred candidate may not win the election, but quite another for them to grant equal weight to others' views about economic policy. Conservative traders skeptical of command-and-control regulation, for example, might apply a discount to lives saved by an air safety regulation that liberal traders who believe government is relatively efficient might not apply. A market with only conservative traders might reach a different result from one with only liberal traders, even if the two groups would reach similar results in a market that, for example, predicted whether a candidate who supported the command-and-control approach would win an election.

This too is potentially empirically testable with an experimental design using a large number of information markets operating simultaneously. For example, a number of markets might be used to predict economic performance in a variety of different states or localities contingent on passage of some Presidential initiative, such as a tax-cut proposal,¹⁷¹ with markets again randomly assigned to allow only those with Democratic preferences, only those with Republican preferences, and/or a mix to participate. Perhaps each individual group will contain a sufficient number of marginal traders who are able to put their own views aside, and thus traders' demographic composition will not matter, but the experiment may show that because Democratic

¹⁶⁹ A possible experimental methodology would examine predictions in sports information markets. *See supra* note 33 and accompanying text (discussing such markets). If group composition matters, one might expect the markets to predict stronger performance from relatively popular teams. Such a finding would not necessarily be dispositive to other contexts, because participants might derive more utility from betting on the home team than from making a prediction consistent with a particular political orientation. It is possible, though, that such preferences can be arbitrated away, providing a strong indication that information markets are uninfluenced by what traders on average would like to see happen.

¹⁷⁰ *See supra* note 164 and accompanying text.

¹⁷¹ A potential problem here is the determination of whether the tax-cut proposal in fact passed. *See infra* note 213.

traders simply think differently than Republican traders, the markets with more Democratic traders would be more skeptical of the benefits of a Republican tax cut proposal.

While evidence indicating that an information market's demographic composition can affect the predictions these markets reach may invoke concern that militates toward placing less weight on these markets in some contexts, it does not necessarily mean that these markets should be ignored altogether. Even an assessment by a biased group may prove useful in governmental decisionmaking, especially when the group's prediction runs opposite its identified bias. The effects of any bias may be small, and the magnitude of the effects may depend on the type of information that the market is predicting. Demographic and ideological characteristics seem less likely to matter, for example, in an information market predicting election returns than in an information market predicting the economic consequences of the adoption of a particular proposal, even assuming comparable analytical complexity in these two predictive problems. Evidence that a trader personally believes that a particular proposal will be good for the economy does not necessarily mean that the trader would then assume that others will recognize the wisdom of the proposal and vote for the candidate associated with it.¹⁷²

The more complex and controversial the relevant models are, and the greater the disparity between the models' evaluations and the individuals' political beliefs, the greater the danger that the traders composition could matter. Even a controversial market, such as a market to estimate the effects of global warming,¹⁷³ might be more objective relative to alternatives than would be an information market making a noncontroversial prediction.¹⁷⁴ The situations in which information markets may be least objective are the same as those in which governmental decisionmaking is likely to be least objective, with the inherent difficulty of the projection task making it more likely that agencies will make projections in accordance with their normative views. An information market at least eliminates cheap talk,¹⁷⁵ forcing those who take positions

¹⁷² Party officials, of course, will offer positive spin on how election news is good for their preferred candidate, but few of us believe that the officials actually believe what they are saying. See Michael Kinsley, *True Lies*, NEW YORKER, Sept. 26, 1994, at 48, 50 (reviewing JAMES CARVILLE & MARY MATALIN, *ALL'S FAIR: LOVE, WAR AND RUNNING FOR PRESIDENT* (1994)) (pointing out how the insistence by the two prominent political "spin control" artists that they never lie looks hollow in the face of the stories they tell describing their great "spin control" victories). But we are less likely to doubt party officials' claims that they in fact prefer the candidate they are working for to the alternatives.

¹⁷³ Compare, e.g., BJORN LOMBORG, *THE SKEPTICAL ENVIRONMENTALIST: MEASURING THE REAL STATE OF THE WORLD* 258-322 (2001) (expressing skepticism about models predicting global warming), with LAURENCE PRINGLE, *GLOBAL WARMING: THE THREAT OF THE EARTH'S CHANGING CLIMATE* (2003) (arguing that global warming does present a threat).

¹⁷⁴ Such a market, however, is unlikely to have any immediate political consequences. See *infra* text accompanying note 260.

¹⁷⁵ Cf. Jason Scott Johnston, *Communication and Courtship: Cheap Talk Economics and the Law of Contract Formation*, 85 VA.

even on controversial matters to back up their predictions with money.¹⁷⁶ The *relative* advantage of information markets is thus likely to be greatest where talk is cheapest. Any lack of objectivity from information markets may be of greater concern than a comparable lack of objectivity by agency officials, because such officials are at least selected by democratic processes. But almost any tool that agencies use in their decisionmaking—including more traditional means of consulting experts¹⁷⁷—can introduce subjectivity without a solid democratic pedigree into the administrative decisionmaking process. Just because information markets seek to increase decisionmaking’s objectivity does not render fatal any failure to obtain a perfectly objective decisionmaking process.

B. Some Potential Governmental Uses for Information Markets

This section describes potential governmental uses for information markets. This is a very partial list.¹⁷⁸ Every administrative agency makes policy decisions, and such decisions are almost always based implicitly on predictions about what the future will bring or how different decisions will affect the future. The examples here are not information markets’ most straightforward applications; to the contrary, they illustrate a range of design challenges. Nor are they chosen for political enactability, though given the FutureMAP controversy, no information market proposal is likely to be enactable at least in the near future. Part II.B.1 addresses the direct goal of the FutureMAP program, considering whether information markets might help identify security vulnerabilities. Though this section explains how objective assessments might produce useful political feedback, it also recognizes problems with the particular application that

L. REV. 385, 412 (1999) (“Formally, cheap talk is defined as a message that does not directly affect the payoff of either the message’s sender or receiver.”).

¹⁷⁶ Without an information market, the government can take a position on global warming to support other goals. For example, if the Bush Administration is concerned that environmental regulation may be too expensive, it might claim skepticism on global warming. Cf. Eric Pianin, *Group Meets on Global Warming*, WASH. POST, Dec. 4, 2002, at A8 (noting Bush Administration skepticism about the causes and effects of global warming, along with environmentalists’ response that adequate evidence exists). Of course, the same possibility could occur in reverse; an administration sympathetic to environmental concerns could discount any doubts about global warming because it would like to minimize air pollution in any event. Information markets will not reflect information that participants wish were true will reflect information that participants *believe* to be true. Markets are thus potentially not objective only when people have trouble separating the two.

¹⁷⁷ Groups of experts themselves may not be random samples of the population. Those who decide to obtain training in environmental policy, for example, may be on average more liberal than the population at large and thus produce different policy recommendations than environmental experts would in a hypothetical world where educational choices were uncorrelated with political beliefs.

¹⁷⁸ As the *New York Times* reported in an article on information markets that appeared two months before the recent controversy, “[o]nce you get the idea, the possibilities are endless.” Hal R. Varian, *A Market Approach to Politics*, N.Y. TIMES, May 8, 2003, at C2.

reinforce the proposition that information markets must be carefully designed. Part II.B.2 provides an example involving information markets with very large numbers of securities, assessing the extent to which such markets might help regulators assess the solvency of individual financial institutions. Part II.B.3 shows how information markets might help provide objective forecasts, in this case with respect to budget deficits, and overcome partisan disputes about methodology. Finally, Part II.B.4 shows how an information market might be used to predict a future governmental decision and help administrative agencies enact interim regulations while they consider long term solutions.

1. Homeland Security Vulnerabilities

One of the goals of the FutureMAP program was to use information markets to improve homeland security. An information market could be adapted toward identifying degrees of vulnerability. An initial test market was used to predict the number of days in a month corresponding to each threat level—“red,” “orange,” and so on—providing the public with a prediction of the average threat level expected over the month.¹⁷⁹ This market received no public attention, but the example illustrates a modestly useful potential application of information markets more generally. While the government raises and lowers threat levels depending on its assessment of intelligence, it does not make predictions of where the threat level is likely to be. Yet Americans might want to factor the anticipated threat level in making travel plans, and a summary could be useful for corporate and local governmental planning as well. An information market provides a simple way for the government to provide some guidance without releasing any supporting information. Of course, the market would incorporate only information known to its participants, so it might not have turned out to have much predictive power. Even if employees with security clearances were prohibited from participating, however, the market at least could have provided a summary of what experts without clearances anticipated about threat levels.

This simple example, however, presents some of the technical and other challenges of the approach that the more controversial later experiment, called the Policy Analysis Market, was to have taken. That market would have been used to predict the probability of specific events and

¹⁷⁹ See http://www.marteksys.com/markets/ThreatLevel_home.html (last visited June 17, 2003).

specific types of attacks. Although the precise nature of the market's securities was not released,¹⁸⁰ it is easy to imagine possibilities. For example, an information market might be used to predict the total amount of property damage from cyberterrorism attacks in a particular year. At the end of that year, the government would estimate the amount of property damage, and that estimate would be used to liquidate the market. Similarly, such markets might be used to predict the number of people killed from attacks on the water supply. Markets could even be used to assess the vulnerability of particular targets, such as the Sears Tower¹⁸¹ or the Palo Verde Nuclear Plant.¹⁸² A comprehensive program might allow trading on all conceived possible means of attack and all prominent individual targets, using markets predicting damage from "other means" and to "other targets" to develop a comprehensive picture of terrorist threats.

At least from a political perspective, the principal problem with such a market appears to be that it seems distasteful. Sen. Hillary Clinton, for example, criticized the program as "a futures market in death."¹⁸³ Perhaps this criticism merely reflects intuitive revulsion, or perhaps it is based on a theory that futures markets are commodificationist and destroy personhood.¹⁸⁴ A full consideration of the implications of commodification theory is beyond the scope of this article, but unless many of our fundamental economic practices are in need of reform, such considerations cannot be dispositive. We tolerate similar distastefulness in other contexts, after all, where it is particularly useful. Life insurance companies set prices based on assessments of when people will die, and providers of annuities and viatical companies in fact benefit when people die.¹⁸⁵ The economic value of a reversionary interest in a life estate depends on an individual's age.

The promise of such a market is that it might help expose imbalances in resource allocations, especially those associated with salient threats. Suppose, for example, that the United States has paid excessive attention to safeguarding air travel and insufficient attention to certain

¹⁸⁰ News reports indicated that the claims would include "whether Palestinian leader Yasser Arafat would be assassinated, or Jordan's King Abdullah II would be overthrown." See Kahn, *supra* note 7.

¹⁸¹ See generally Robert L. Kaiser, *Taking America's Pulse Inside the Sears Tower*, CHI. TRIBUNE, Sept. 8, 2002, available at 2002 WL 26639737 (reporting on concerns about the Sears Tower's vulnerability to terrorist attacks).

¹⁸² See Bill Gertz & Jerry Seper, *Nuclear Power Plant a Potential Target*, WASH. TIMES, Mar. 20, 2003, at A13.

¹⁸³ See, e.g., Tom Detzel, *Pentagon Kills Plan for Terror Futures Market*, OREGONIAN, July 30, 2003, available at http://www.oregonlive.com/news/oregonian/index.ssf?/base/front_page/1059567193164150.xml.

¹⁸⁴ See generally MARGARET JANE RADIN, *CONTESTED COMMODITIES* (1996) (developing commodification theory).

¹⁸⁵ See generally Russell J. Herron, Note, *Regulating Viatical Settlements: Is the Invisible Hand Picking the Pockets of the Terminally Ill?*, 28 U. MICH. J.L. REFORM 931 (1995).

other threats. An information market prediction indicating greater risk from other threats would be consistent with that claim. A limitation of such markets is that data on threats alone is insufficient to justify a policy conclusion. In theory, the marginal decrease in risk associated with the last dollar of spending to combat each of a variety of threats should be equal,¹⁸⁶ but the information markets measure only the risk side of the equation, not the extent to which spending can address that risk. Conditional markets are a possible solution here, though the caveats about such markets discussed above¹⁸⁷ are compounded by the non-binary nature of the decision of how much money to allocate to each risk. Given the wide range of policy responses to potential terrorist risks, it might be difficult to create information markets that would sufficiently assess the full range of policy options.¹⁸⁸

In the absence of an adequate conditional market, an additional problem is the danger of circularity.¹⁸⁹ Suppose, for example, that a trader identifies a new security vulnerability that terrorists might exploit, but one that the government could address and eliminate if the market price rose to reflect the risk. The information suggests that the security is undervalued, but if the trader drives up the price of the security, it will then be overvalued as the government response reduces the risk back to its initial level. It is difficult to predict how such a market will equilibrate. Perhaps the optimal strategy might be to trade on securities only up to a point at which the increased price would not trigger a government reaction. Information markets would then reflect only information about which nothing can or would be expected to be done, considerably reducing their usefulness. Maybe the optimal way to avoid this problem would be for information markets to have a relatively short time horizon, perhaps estimating different risks over a period of about a month.¹⁹⁰ Such an approach would allow estimates of current risk levels while still providing data to facilitate long-term responses.

¹⁸⁶ A caveat is that it may make sense to spend extra on risks that cause the most anxiety. See Frank B. Cross, *The Public Role in Risk Control*, 24 ENVTL. L. 887, 924-26 (1994) (considering “dread risks”).

¹⁸⁷ See *supra* Part I.C.1.

¹⁸⁸ It might be possible to have conditional markets for each of a number of different possible levels of spending. The greater the segmentation of conditional markets, however, the less the subsidy is for each information market, and the greater the noise in the market price, making differences between prices of conditional securities less meaningful. See *supra* text accompanying note 77.

¹⁸⁹ Gross offers the following explanation of the circularity:

[T]he market might defeat itself. The Pentagon wanted to create the PAM in order to gather information it could use to stop terrorism and reduce instability. If it saw, say, that people were betting heavily on the assassination of Iraq's interim president, the Defense Department would start searching for some assassination plot in the hopes of rooting it out. But preventing the assassination would cause all the people who bet on it to lose their money.

Gross, *supra* note 7.

¹⁹⁰ A complication here is that the probability estimates would likely be very small. This presents two potential problems, the

The problem upon which many critics of FutureMap focused is the danger that insider trading profits might encourage terrorist acts. Insider trading in information markets generally should be encouraged, because the purpose of such markets is to encourage consideration of all potentially relevant information.¹⁹¹ Indeed, information markets in theory could even provide a means of implementing anti-terrorism bounty programs.¹⁹² The is, however, a concern that terrorists might trade on their acts just before implementing them and thus profit from terrorism. Worse, information markets could create a moral hazard problem,¹⁹³ encouraging individuals who otherwise would not have been terrorists to commit terrorist acts for profit.

The amount of money that could be made in this way, however, would be relatively small, unless the government poured millions of dollars into such a market. Moreover, trading activity might leave a paper trail that could lead to the terrorists' apprehension and that might prevent terrorists from cashing in on the appreciation in share value. The government could require that traders be identified, perhaps even subject to background checks, and that government payouts be made and transactions cleared only after it is determined that traders did

second more serious than the first. First, cognitive psychology experiments indicate that people generally are not as good at estimating small probabilities. *See, e.g.,* Daniel Kahneman & Amos Tversky, *Prospect Theory: An Analysis of Decision Under Risk*, 47 *ECONOMETRICA* 263, 281 (1979). At least some market participants, however, would likely be sophisticated enough to rely on formal models rather than intuitive probability estimations to guide trading in such a market. Second, what may be most important in such a market is the relative prices of securities, reflecting the relative risk of different types of attacks. To make comparisons of security prices meaningful, the government would need to promise a relatively large payment in the event of an actual terrorist attack, so that traders would have incentives to investigate the levels of different risks and factor them into trading decisions.

¹⁹¹ The same argument can be made for traditional securities markets, and indeed Henry Manne has argued that insider trading should generally be permitted, as insider trading will promote market efficiency. *See* HENRY G. MANNE, *INSIDER TRADING AND THE STOCK MARKET* (1966); *see also* Dennis W. Carlton & Daniel R. Fischel, *The Regulation of Insider Trading*, 35 *STAN. L. REV.* 857 (1983) (building on Manne's general thesis); David D. Haddock, *Academic Hostility and SEC Acquiescence: Henry Manne's Insider Trading*, 50 *CASE W. RES. L. REV.* 313 (1999) (commenting on Manne's work and its influence). Whatever the merits of Manne's claim as applied to the stock market, it seems powerful in the context of an information market. An information market's sole purpose is to encourage the release of information, and the welfare of market participants is not a significant concern. Thus, the strongest counterarguments to Manne's theory do not apply. For example, Michael Manove has argued that insider trading may discourage corporate investment, but corporate investment is not an issue here. *See* Michael Manove, *The Harm from Insider Trading and Informed Speculation*, 104 *Q.J. ECON.* 823 (1989).

¹⁹² The Justice Department currently sponsors a "Rewards for Justice" program that promises to pay individuals who provide information that thwarts terrorist acts. *See* <http://www.rewardsforjustice.net/> (last visited June 19, 2003); *see also* 22 U.S.C.A. § 4821 (West Supp. 2001) (providing the statutory basis for the program). It would be straightforward to use an information market as a means of providing the reward. An information market might be used, for example, to determine whether the government uncovers a terrorist plot of a particular type during a given time period. By trading on such a market and then releasing information to the government, an informant could assure a trading profit. Note that the market could not be used to predict successful terrorist attacks, however, for then trading that succeeds in alerting the government to a terrorist attack would not be profitable.

¹⁹³ In the insurance context, "[m]oral hazard is the . . . tendency of an insured to underallocate to loss prevention after purchasing insurance." KENNETH S. ABRAHAM, *DISTRIBUTING RISK: INSURANCE, LEGAL THEORY, AND PUBLIC POLICY* 14 (1986). This context provides a more direct and troubling form of moral hazard that insurance law generally worries about only in the life insurance context.

not themselves participate in any terrorist acts. With appropriate regulation, an information market would become an extremely unlikely venue for this type of scheme. It would be far easier to make money by trading on conventional securities markets,¹⁹⁴ or even by bribing a basketball player to shave a few points in a game. Nonetheless, even if the possibility of such side effects are remote, the worst case scenario may be so serious that even the probability-adjusted costs might outweigh any benefits of such a program. In retrospect, the moral hazard danger made this application of information markets a poor choice for preliminary experimentation, although the project might not have been funded at all in the absence of the terrorism hook.

2. *Solvency Regulation*

The primary virtue of terrorist vulnerability information markets is that they help expose resource misallocations, possibly providing feedback to the governmental decisionmakers, and ultimately leading to the correction of such misallocations. Information markets might prove more directly useful, however, at governmental tasks that are explicitly predictive and involve processing of massive amounts of information. A useful example, though in the end perhaps not a practicable one, is solvency regulation. The federal agencies that share responsibilities for regulating depository institutions regularly monitor those institutions' financial performance to prevent insolvency.¹⁹⁵ A principal justification for such monitoring is that with federal deposit insurance, institutions in danger of insolvency have an incentive to take excessively risky bets in a last desperate effort to avoid losing their investments.¹⁹⁶ States engage in a similar form of monitoring of insurance companies, seeking to ensure that insurance companies will be able to meet their obligations even in the event of unexpectedly high liabilities.¹⁹⁷ In both cases, the regulatory regime is complicated, as regulators focus on a variety of indicia of safety, such as reserve ratios and other capital standards.

¹⁹⁴ Indeed, there were reports after September 11 that the terrorists might have profited by trading on options markets before the attacks, though subsequent investigations indicated that they did not do so. See Erin E. Arvedlund, *Chance, Not Terror, Was Behind Suspicious Deals*, ASIAN WALL ST. J., Feb. 1, 2002, at W2, available at 2002 WL-WSJA 3343602.

¹⁹⁵ For an overview of bank solvency regulation, see GEORGE J. BENSTON & GEORGE G. KAUFMAN, RISK AND SOLVENCY REGULATION OF DEPOSITORY INSTITUTIONS: PAST POLICIES AND CURRENT OPTIONS (Fed. Res. Bank of Chicago Staff Mem. No. 88-1, 1988); MICHAEL P. MALLOY, BANK REGULATION 229-62 (1999).

¹⁹⁶ See, e.g., R. Mark Williamson, *Regulatory Theory and Deposit Insurance Reform*, 42 CLEV. ST. L. REV. 105, 120-21 (1994) ("Moral hazard is present with deposit insurance because banks which increase the riskiness of their portfolio will capture all upside profits while bearing only part of the downside losses . . .").

¹⁹⁷ See, e.g., KENNETH S. ABRAHAM, INSURANCE LAW AND REGULATION: CASES AND MATERIALS 98-104 (3d ed. 2000).

Simple information markets could be used to provide some form of political feedback. For example, a market might be used to predict the number of depository institutions that will be expected to go insolvent in a particular year or the total dollar amount of the federal government bailout corresponding to such failures. Such a market might have led to earlier recognition of the savings and loan crisis in the 1980s.¹⁹⁸ Although such recognition presumably would not have averted closures, it might have prevented some of the high-risk activity that exacerbated the problem.¹⁹⁹ The more challenging task for information markets, however, would be to perform the task of monitoring individual institutions. The goal, after all, of these regulatory regimes is to make individualized predictions about the likelihood of insolvency and the corresponding likelihood of either consumer or governmental losses, or both, in the event of insolvency. Such markets could foster more objective decisionmaking. Given the criticism that special interests have affected monitoring in the past,²⁰⁰ objective predictions might be a considerable improvement.

The large number of institutions at issue should not be a significant problem for such a market, as long as a different security is issued for each one. Just as traders in the Hollywood Stock Exchange, or any other information market, have incentives to consider each security individually,²⁰¹ so too would insolvency information markets' traders, assuming that there was a sufficient financial incentive. The greater challenge is overcoming traders' limited access to relevant information. Although mandatory reporting requirements furnish some publicly available data, governmental inspections are a key aspect of such regulatory regimes.²⁰² Privacy considerations presumably would prevent traders from obtaining the full access that governmental officials enjoy. Traders would thus need to base their predictions on publicly available information and any other information that they are able privately to obtain.

¹⁹⁸ See generally Carl Felsenfeld, *The Savings and Loan Crisis*, 59 *FORDHAM L. REV.* 57 (1991) (providing an overview of the crisis).

¹⁹⁹ See Joseph A. Grundfest, *Lobbying into Limbo: The Political Ecology of the Savings and Loan Crisis*, 2 *STAN. L. & POL'Y REV.* 25 (1990) (describing how regulatory laxity produced high-risk behavior).

²⁰⁰ See, e.g., Mark Seidenfeld, *A Civil Republican Justification for the Bureaucratic State*, 105 *HARV. L. REV.* 1511, 1567 (1992) (identifying special interest lobbying of Congress as a factor in loose monitoring of endangered thrifts).

²⁰¹ Each security in the Hollywood Stock Exchange is liquidated based on the performance of a particular film. See *supra* text accompanying notes 48-51. While it might be profitable to trade on a number of securities based on information expected to affect filmgoing generally, traders generally profit by finding errors in the market's pricing of individual films.

²⁰² Critical to the programs are the randomness of inspections, since banks and insurance companies may be able to make temporary changes to conceal weaknesses. See Andrew Chin, *Spoiling the Surprise: Constraints Facing Random Regulatory Inspections in Japan and the United States*, 20 *NW. J. INT'L L. & BUS.* 99, 99-100 (1999) (discussing a scandal in Japan involving governmental officials tipping off banks as to the dates of their inspections).

That information alone, however, may be sufficient to enhance the monitoring process. Whistleblowers with particular knowledge of problems at depository institutions have an incentive to trade on the information and thus reveal the problems. Depository institutions, moreover, would have some incentives to open up their files to private parties that will independently verify their financial information, lest traders infer from a lack of cooperation that the institutions are hiding information.²⁰³ Theory alone is insufficient to predict the extent to which these dynamics will affect information markets' success at incorporating private information. Even a financially modest experiment, publicly or privately financed, could provide traders with sufficient incentives to develop models predicting the probability of insolvency based on public information,²⁰⁴ and such models might help overcome any loopholes in current governmental accounting rules that may allow institutions in danger to clear government benchmarks. In theory, if such an experiment proved successful, governmental agencies could consider as a next step incorporating market predictions into governmental formulas determining the extent of governmental intervention.

3. *Budget Forecasts*

The examples so far have involved areas in which the application of information markets may be controversial, but the policy issues corresponding to the relevant predictions tend not to be controversial. Information markets could be used in more controversial areas, but many information markets assessing controversial issues might offer little other than academic value. For example, it would be straightforward, if macabre, to use a conditional market to assess the deterrence effect of the death penalty, by creating securities whose redemption value would depend, for a particular geographic area, on both the number of executions in a specified time period and the number of murders in the next time period.²⁰⁵ Yet, for at least two reasons, such a

²⁰³ Once some institutions reveal information, an adverse selection effect increases the pressure on other institutions to do so as well. Because an institution will have an incentive to allow inspections whenever they believe that traders are overly suspicious of them, the best of the institutions that initially do not reveal information will have a strong incentive to do so. This dynamic can produce an unraveling effect, so that even institutions who fear that inspections will reveal negative information will allow such inspections lest traders infer that the situation is even worse than it actually is. *See generally* Paul Milgrom & John Roberts, *Relying on the Information of Interested Parties*, 17 RAND J. ECON. 18 (1986) (providing a model of how adverse selection may force information release).

²⁰⁴ Such an experiment might be particularly successful if regulators were allowed to trade, given that regulators have the best information. If the markets were significant enough, the chance of trading profits in effect would become part of the governmental officials' compensation and would motivate them to scrutinize institutions carefully.

²⁰⁵ A recent study considers the deterrent effect of the death penalty by examining retrospectively the crime rate shortly before and after such executions. *See* Hashem Dezbakhsh et al., *Does Capital Punishment Have a Deterrent Effect? New Evidence from*

study would be likely to play only a peripheral role in the death penalty debate. First, although deterrence figures largely in debates on the death penalty, people's views on the death penalty may depend more on the extent to which they believe in retributive justifications for punishment.²⁰⁶ Second, an information market prediction does not resolve the underlying question, but provides at most a best guess. An information market is thus akin to event studies in corporate law, serving as a sophisticated form of opinion poll that might prove incorrect.²⁰⁷ An information market thus provides the educated analyst little reason to change an opinion, and on an issue like the death penalty, there are likely to be highly knowledgeable experts on both sides.

If information markets ever were to receive acceptance as a legitimate tool, then information markets on controversial issues would most likely be relevant when competing parties have some incentive to commit to them in advance. One area in which this might be true would be in budget forecasting, for example the prediction of the size of government deficits in future years conditional on the enactment of legislation.²⁰⁸ A desire to have relatively objective data has led Congress to create the Congressional Budget Office (CBO), which generally is viewed as seeking to guard its independence and thus the integrity of its estimates.²⁰⁹ Budget issues, however, can become political despite such attempts, in part because the majority party appoints the leadership of the CBO.²¹⁰ A prominent controversy has been whether budget forecasters should employ "dynamic scoring,"²¹¹ a method that assumes tax cuts will increase economic growth, and how much dynamic scoring they should allow.²¹² The disagreements about the accuracy of dynamic scoring may imperil the legitimacy of budget forecasts, and potentially both parties could be amenable to a solution that would depoliticize forecasts.

Post-Moratorium Panel Data (2003) (unpublished manuscript, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=259538).

²⁰⁶ See Dan M. Kahan, *The Secret Ambition of Deterrence*, 113 HARV. L. REV. 413 (1999) (arguing that deterrence arguments have little effect on citizens' views, even if they are common in public discourse).

²⁰⁷ See Jill E. Fisch, *Picking a Winner*, 20 J. CORP. L. 451, 465 (1995) (book review) (noting that reliance on an event study may produce bad policy if the market has processed information irrationally).

²⁰⁸ For an example of deficit forecasts given existing policies, see <http://www.cbo.gov/showdoc.cfm?index=1944&sequence=0> (last visited June 19, 2003).

²⁰⁹ See generally Philip G. Joyce & Robert D. Reischauer, *Deficit Budgeting: The Federal Budget Process and Budget Reform*, 29 HARV. J. ON LEGIS. 429, 431-33 (1992) (discussing the origins of the Congressional Budget Office).

²¹⁰ See 2 U.S.C. § 601(A)(2) ("The Director shall be appointed by the Speaker of the House of Representatives and the President pro tempore of the Senate . . .").

²¹¹ *The Reaction—The Battle over the Tax Cut*, ECONOMIST, Jan. 18, 2003, available at 2003 WL 6244623 (discussing the attempt by Republicans to implement dynamic scoring).

²¹² See generally Alan J. Auerbach, *Dynamic Revenue Estimation*, 10 J. ECON. PERSP. 141 (1996) (considering methodological and implementation issues associated with dynamic scoring).

Information markets offer a straightforward solution, which presumably would incorporate dynamic considerations to some extent but not as much as dynamic scoring's most ardent proponents might like. A simple approach would be for securities to be tied toward ex post evaluations of the deficit, and conditional securities could be used to assess the effects of different possible policies.²¹³ A possible concern here is that there may be subjectiveness even in such retrospective assessments, for example because of accounting issues concerning investment and asset sales.²¹⁴ As long as the identity and party affiliation of the budget officials who will make the assessments is unknown, however, information markets will predict an average of what different officials might be expected to decide,²¹⁵ thus stripping political bias out of the estimate. For the officials to be unknown, assessments should be scheduled for a significant period of time, perhaps ten years, after the launch of the market, even when the year in question follows soon afterward launch. This is a critical point, to which I will return in discussing predictive cost-benefit analysis:²¹⁶ The possibility for subjectiveness in determinations of market payouts is not a problem as long as the ultimate decision is postponed and as long as a prediction of how an average decisionmaker would make the assessment is sufficient for the relevant purpose.

4. *Interim Regulation*

That predictions in a budget forecasting information market are in fact predictions of a future governmental assessment does not make that market unique. Any information market requires that some individual, whether or not governmental, make a final assessment of what the number being predicted turned out to be. Information markets, however, might be used to predict not only governmental assessments made specifically for the purpose of concluding the markets, but also governmental decisions on issues of regulation. In many administrative contexts, an agency has insufficient resources to address all of the regulatory questions facing it. For

²¹³ The concern about selection bias with conditional markets, *see supra* note 75 and accompanying text, is not salient here, because Congress is unlikely to have substantial private information that would be relevant to the market's assessment of what the deficit would be conditional on the enactment of a particular policy. The most significant challenge with such a market would be developing suitable definitions for the relevant conditions. The problem is that a particular budget proposal almost never will be adopted in full. Thus, the conditional market might estimate the deficit conditional on the specified tax rate.

²¹⁴ Thomas Seto has emphasized this point in arguing that a Balanced Budget Amendment to the Constitution would need an "independent scorekeeper" to make deficit assessments. *See* Theodore P. Seto, *Drafting a Federal Balanced Budget Amendment That Does What It Is Supposed to Do (and No More)*, 106 YALE L.J. 1449, 1511-15 (1997).

²¹⁵ *See infra* note 252 and accompanying text.

²¹⁶ *See infra* Part III.

example, the Environmental Protection Agency is notoriously slow in determining whether to classify particular chemicals as pollutants.²¹⁷ Similarly, the Food and Drug Administration has often faced criticism for the amount of time that it takes to conduct drug trials.²¹⁸ In these cases, an information market might be used to predict the decision that the agency eventually will reach, and that decision could be used to fashion an interim regulatory regime.

In the EPA context, for example, an information market might include a number of securities corresponding to different chemicals, with the security to be liquidated at a set value if the chemical eventually is classified as a carcinogen. Without conducting any scientific research on its own, the agency might enact rules concerning the regulation of such carcinogens. Conceivably, the EPA might decide to limit any expanded use of chemicals for which the information market indicated a significant probability of ultimately being labeled a carcinogen. A more elaborate information market might predict the “dose-response” curve corresponding to a particular chemical.²¹⁹ In such a market, securities would be traded predicting the effect of each chemical on humans at various levels of exposure. Although these are notoriously difficult assessments,²²⁰ the market at least would provide an objective prediction of the EPA’s later determination. The information market alone cannot determine the best interim legal regime,²²¹ but it can provide an agency with limited resources data that it can use pending its own investigation. At least, such an information market can help an agency decide which decisions to prioritize.

²¹⁷ See, e.g., John P. Dwyer, *The Pathology of Symbolic Legislation*, 17 *ECOLOGY L.Q.* 233, 259-60 (1990).

²¹⁸ For an evaluation of a federal statute that seeks to expedite the process, see Deborah G. Parver, Note, *Expediting the Drug Approval Process: An Analysis of the FDA Modernization Act of 1997*, 51 *ADMIN. L. REV.* 1249 (1999).

²¹⁹ “Dose/response assessment estimates the relationship between the amount, intensity, or duration of exposure and the risk of a particular outcome (e.g., . . . a worker’s lifetime probability of contracting leukemia).” Robert A. Pollak, *Regulating Risks*, 33 *J. ECON. LIT.* 179, 183 (1995).

²²⁰ See Sunstein, *supra* note 120, at 2279-80 (discussing dose-response curves and the difficulties of calculating them).

²²¹ In the FDA context, a relevant question would be whether drugs should be provisionally allowed when there is a prediction that they ultimately will be approved. An argument for such a rule is that many people may die while waiting for a drug to be approved, and an information market prediction at least provides some assurance that the drug company is not seeking to deceive consumers by peddling snake oil. Cf. *Investigational New Drug, Antibiotic, and Biological Drug Product Regulations*, 53 *FED. REG.* 41516-01 (Oct. 21, 1988) (noting the FDA’s goal of “speed[ing] the availability of new therapies to desperately ill patients, while preserving appropriate guarantees for safety and effectiveness”). A contrary argument is that the government might find it difficult to eliminate supplies of the drug should it ultimately decide not to approve it.

C. Institutionalizing Information Markets

An information market used for interim regulation provides a useful example of why an administrative agency might like to launch an information market. It also shows a modest way in which an information market might be helpful in judicial review. Some critics of the administrative process complain that it is “ossified,”²²² with the cumbersomeness of notice-and-comment rulemaking preventing administrative agencies from enacting useful regulations.²²³ Although the Administrative Procedure Act allows an agency to skip the notice-and-comment process for “good cause,”²²⁴ courts have been reluctant to allow agencies to count the interim nature of a rule as sufficient good cause.²²⁵ Such reluctance is understandable, given the danger that agencies might abuse such a prerogative by regulating through series of interim rules. An information market used to predict an ultimate agency decision, however, can provide the agency with some data to support the view that it will indeed eventually find sufficient evidence to act in a particular way. Perhaps this data would be sufficient to convince courts that interim regulation is appropriate.

This is a modest use of information markets, complementing rather than replacing existing decisionmaking processes. There are, however, likely to be many steps before agencies and courts would use information markets even in this way. Even the now-defunct FutureMAP program involved only experimental uses of information markets, and there was no indication that the Defense Department would take the predictions of such information markets into account in decisionmaking, let alone present them to support its actions in judicial review.²²⁶ As I have suggested, further experimentation would be necessary in any event to ensure that such markets are sufficiently unbiased and invulnerable to attempts at manipulation.²²⁷ Perhaps if

²²² See Thomas O. McGarity, *Some Thoughts on “Deossifying” the Rulemaking Process*, 41 DUKE L.J. 1385 (1992).

²²³ *Id.* at 1387-96. The Administrative Procedure Act appears to require relatively little information by agencies in the rulemaking process. See 5 U.S.C. § 553(b)(3) (“The notice shall include . . . either the terms or substance of the proposed rule or a description of the subjects and issues involved.”). The D.C. Circuit, however, has interpreted the Act aggressively, for example by requiring agencies to release background documents that might be necessary for parties to make informed comments on the rules. See *Portland Cement Ass’n v. Ruckelshaus*, 486 F.2d 375 (D.C. Cir. 1973) (rejecting a Notice of Proposed Rulemaking, even though the agency had provided the text of the proposed rule in draft form).

²²⁴ See 5 U.S.C. § 553(b)(B).

²²⁵ See, e.g., *Tenn. Gas Pipeline Co. v. FERC*, 969 F.2d 1141 (D.C. Cir. 1992) (rejecting two rationales for application of the “good cause” exception).

²²⁶ The Defense Department is often exempt from notice-and-comment rulemaking requirements anyway. See 5 U.S.C. § 553(a) (providing an exemption for “a military or foreign affairs function of the United States”).

²²⁷ See *supra* Part II.A.2.

experimentation confirms the promise of information markets, an intrepid agency might take an additional cautious step by creating some information markets that provide public information without any requirements that agency officials pay attention to such markets in making decisions.²²⁸

Even apart from negative publicity, obstacles remain even to such invocation of information markets. First, one might argue that the Commodity Futures Trading Commission has jurisdiction over information markets.²²⁹ The CFTC, however, has provided no action letters for the Iowa Electronic Markets,²³⁰ and it presumably would do so as well for governmental agencies. Second, gambling on information markets might be seen as a violation as state gambling laws,²³¹ although federally authorized trading would be immunized under the Supremacy Clause.²³² Third, there is no existing infrastructure to support the creation of information markets. In theory, a number of agencies might create their own information market programs, but there are presumably economies of scale to running information markets, and the hassle of creating the infrastructure, combined with the publicity that would accompany any agency that was the first to experiment with information markets after the FutureMAP debacle, might discourage agencies that otherwise would be willing to devote some resources to particular predictive tasks.

Perhaps a long time from now, FutureMAP will be a distant memory, and use of information markets in the private sector or in other countries would lead to a sense that further experimentation is warranted. Congressional intervention could greatly facilitate further experimentation with information markets by administrative agencies. As an initial matter, a

²²⁸ Information markets to predict terrorist activity might be a useful initial project. *See supra* Part II.B.1.

²²⁹ The question hinges in part on the definition of “commodity.” The statute provides that “it shall be unlawful for any person to offer to enter into . . . a contract for the purchase or sale of a commodity for future delivery” in the absence of regulatory approval or exemption. 7 U.S.C. § 6 (2000). “Commodity” is in turn defined broadly to include, in addition to specifically enumerated commodities like wheat, cotton, and rice, “all services, rights, and interests in which contracts for future delivery are presently or in the future dealt in.” *Id.* § 1a(4). There is a strong argument that the securities in information markets are not “contracts for future delivery,” as no “delivery” is contemplated.

²³⁰ *See* Commodity Futures Trading Commission Ltr. No. 93-66, Comm. Fut. L. Rep. ¶ 25.785, *available at* 1993 WL 595741 (providing a redacted version of the authorization).

²³¹ Participating in information markets at least superficially appears to meet some definitions of “gambling.” *See, e.g.,* MCKINNEY’S N.Y. PENAL LAW § 225.00 (2003) (defining “gambling” as “A person engages in gambling when he stakes or risks something of value upon the outcome of a contest of chance or a future contingent event not under his control or influence, upon an agreement or understanding that he will receive something of value in the event of a certain outcome.”). For a consideration of the regulatory obstacles to information markets, see Peter Locke, *Regulatory Obstacles to Information Markets* (2002) (unpublished PowerPoint presentation, on file with author).

²³² U.S. CONST. art. VI.

statute might explicitly empower agencies to create information markets, at least in the absence of moral hazard concerns such as those present in the terrorism context. Such a statute would eliminate any concerns about CFTC jurisdiction and gambling laws.²³³ In addition, Congress might designate a particular agency, perhaps the Office of Management and Budget (OMB), to coordinate a single web site hosting information markets for a variety of agencies, with other agencies providing the finances necessary for subsidizing their particular projects. Finally, Congress might appropriate funds for specific projects, although agencies also might use lump sum appropriations in the absence of specific funding.²³⁴ Perhaps concerns with appearances will serve as a permanent block to such support, but at least if information markets became more accepted among academic, there are no obvious interest groups that would be expected to oppose such moves.

Even if the regulatory climate were more favorable, information markets probably would not become pervasive in administrative decisionmaking. Cost-benefit analysis received a significant impetus when President Reagan signed an executive order,²³⁵ since renewed in varying forms,²³⁶ mandating that it be used for certain decisions. A similar blanket policy is less plausible for information markets, because discretion is needed to identify what types of information to seek and how individual information markets should be designed. Thus, information markets can grow only when individual agencies take initiative, either because of genuine interest in the information produced, or because an agency wished to provide some support for the predictive judgments influencing policy decisions. A President who wanted to rationalize agency decisionmaking with information markets might give a single agency, such as OMB, responsibility for deciding on appropriate prediction problems. That agency would then

²³³ A more aggressive approach might authorize private entities to create information markets as well, perhaps subject to some regulatory supervision, for example by the CFTC. Private implementation of information markets, for example by corporations interested in particular information, not only might be intrinsically useful, but also might provide further data on the design of such markets. Aside from the Iowa Electronic Markets, current experiments have been quite limited, as participation has been restricted to particular individuals, none of whom bears any risk of a net financial loss. *See, e.g.,* CHEN & PLOTT, *supra* note 60 (noting the closed nature of the HP market). Any statute authorizing private information markets might include specific exceptions designed to reflect concerns that information markets can serve as a form of gambling, for example by prohibiting information markets for sporting events. *See supra* note 33 (noting that existing overseas information markets allow for gambling on sports teams).

²³⁴ *See Lincoln v. Vigil*, 508 U.S. 182 (1993) (holding that agencies' decisions about how to spend lump sum appropriations are not subject to judicial review).

²³⁵ Exec. Order No. 12,291, 3 C.F.R. 127 (1982), *reprinted in* 5 U.S.C. § 601 (1988).

²³⁶ *See, e.g.,* Exec. Order No. 12,866, 3 C.F.R. 638 (1994), *reprinted in* 5 U.S.C. § 601 (2000).

have much the same role that OMB has today in enforcing regulatory requirements ranging from cost-benefit analysis to paperwork reduction.²³⁷

Should information markets exist on topics relevant to particular administrative decisions, no amendments to the Administrative Procedure Act would be necessary for courts to take them into account in checking the rationality of judicial policy decisions. Under the “hard look” doctrine, one of the means by which courts assess whether agency actions are “arbitrary” or “capricious,”²³⁸ courts closely analyze administrative agencies’ justifications for their policy choices.²³⁹ In *Motor Vehicle Manufacturers Ass’n v. State Farm Mutual Automobile Ins. Co.*,²⁴⁰ for example, the Supreme Court scrutinized the National Highway Traffic Safety Administration’s interpretation of studies on the effect of seatbelt usage, noting that the agency’s prediction of less than a five percentage point increase in safety belt usage failed to take into account distinctions between two types of belts.²⁴¹ Similarly, a court might vacate an administrative action when an agency has failed to produce an adequate explanation of why that decision is justified given a particular prediction by an information market. This is a limited power, as the court can only guarantee that the agency conducts a “reasoned analysis,”²⁴² not that the agency picks among a variety of reasonable conclusions that which the court thinks best. The more objective data that courts have to assess administrative agency reasoning, however, the more effectively courts will be able to conduct this task.

In theory, judicial review might take into account even privately run information markets. Nothing in statutory or case law requires an agency or a court to consider only information that the agency itself has created, and indeed agencies participating in notice and comment presumably consider a wide range of evidence. If some form of private information market existed and produced information relative to a particular administrative rulemaking, then both

²³⁷ See generally Robert V. Percival, *Rediscovering the Limits of the Regulatory Review Authority of the Office of Management and Budget*, 17 ENVTL. L. REP. 10017 (1987) (providing an overview of OMB’s authority and limits on that authority).

²³⁸ See 5 U.S.C. § 706(2)(A) (requiring the reviewing court to “hold unlawful and set aside agency action, findings, and conclusions found to be . . . arbitrary, capricious, an abuse of discretion, or otherwise not in accordance with law”).

²³⁹ Though firmly entrenched, the hard look doctrine remains controversial. See, e.g., Richard J. Pierce, Jr., *The Unintended Effects of Judicial Review of Agency Rules: How Federal Courts Have Contributed to the Electricity Crisis of the 1990s*, 43 ADMIN. L. REV. 7 (1991) (exploring the effect of hard look review on substantive regulatory outcomes in the electricity regulation context).

²⁴⁰ 463 U.S. 29 (1983).

²⁴¹ *Id.* at 54 (“[T]his and other statements that passive belts will not yield substantial increases in seatbelt usage apparently take no account of the critical difference between detachable automatic belts and current manual belts.”).

²⁴² *Id.* at 57 (holding that a “reasoned analysis” is sufficient even when an agency is changing its policy course).

agencies and courts presumably could take such information into account. The Iowa Electronic Markets exist despite regulatory obstacles, and it is plausible that universities or non-profit organizations might create information markets devoted to particular policy issues. Judges at first presumably would be reluctant to place any substantial weight on information market results, but acceptance could build over time, and it would take but one judicial panel to take the first critical step. My purpose, of course, is not to predict that information markets will figure in judicial review any time soon, and my own instinct is to doubt it, but the open-endedness of administrative law at least does not forbid the consideration of evidence from information markets.

III. PREDICTIVE COST-BENEFIT ANALYSIS AND BEYOND

A principal challenge in creating information markets that the examples above demonstrate is the determination of what information the markets should be constructed to predict.²⁴³ The information must be sufficiently important to bear on the decision, but even once a clearly relevant type of information is predicted, governmental decisions will often depend on a variety of additional factors. The information markets described so far can provide inputs into governmental decisions, but they cannot provide comprehensive assessments of the decisions themselves. Predictive cost-benefit analysis provides at least a solution to these dilemmas, by creating an information market to predict the outcome of a cost-benefit analysis. Because a cost-benefit analysis can incorporate, at least in theory, all of the costs and benefits of a particular decision, the methodology avoids the problem of defining the relevant contours of a decision. It thus produces bottom-line assessments of potential government decisions. While existing information markets predict subsequently objectively verifiable information, the potential for comprehensive analysis of individual policy decisions makes predictive cost-benefit analysis a useful candidate for the future.

Predictive cost-benefit analysis is attractive for reasons other than that it provides a template that can easily accommodate virtually any form of governmental action. Predictive cost-benefit analysis is responsive both to concerns about information markets and to concerns

²⁴³ In the cyberterrorism context, for example, the difficult questions are which terrorist activities to predict and over what time frame. *See supra* Part II.B.1. The problem is especially acute for conditional markets, which require also a definition of the condition upon which the markets will depend.

about traditional cost-benefit analysis. Although conditional markets can also be used to assess individual decisions, predictive cost-benefit analysis overcomes some of the technical problems that make the results of conditional markets difficult to interpret. A disadvantage of predictive cost-benefit analysis relative to other information markets is that predictive cost-benefit analysis results in predictions of what people will say their opinions are, and sometimes this might deviate from what their opinions really are. In most contexts, however, this disadvantage is minor, as a result of the most significant virtue of predictive cost-benefit analysis.

This virtue is that predictive cost-benefit analysis allows for an assessment of how an average decisionmaker, rather than any particular decisionmaker, would rate a particular policy. Unlike traditional cost-benefit analysis, predictive cost-benefit analysis will not vary based on the happenstance of which political party or which particular individuals happen to be in control of an agency at the time of the analysis. A proponent of traditional cost-benefit analysis might protest that it would be worse to have politically moderate decisionmakers all of the time than to have politically more extreme decisionmakers from opposite sides of the ideological spectrum sharing power across time. Allowing Democrats free reign for a few years and Republicans for another few years is better than having moderates always control policy, for example because experimentation with different approaches is likely to produce useful feedback into the political process.²⁴⁴ This perspective sees the abrupt shifts in power following elections as inherently desirable rather than as the mere byproduct of a political system that seeks to enshrine the preferences of a hypothetical median voter but can do so only approximately and only in fits and starts.

Though important, this objection is beyond my scope, because my argument is simply that predictive cost-benefit analysis is a more useful tool than traditional cost-benefit analysis, not that agencies necessarily should be required to follow its recommendations. The more decisionmakers can manipulate cost-benefit analysis to reflect their own preferences, the less valuable cost-benefit analysis becomes as a signal of administrative action.²⁴⁵ At an extreme of pliability, cost-benefit analysis becomes useless, providing meaningless positive evaluations of

²⁴⁴ This position dates at least to President Jackson. *See, e.g.,* John F. Duffy, *The FCC and the Patent System: Progressive Ideals, Jacksonian Realism, and the Technology of Regulation*, 71 U. COLO. L. REV. 1138 (2000) (discussing how President Jackson limited agency officials' tenure to promote rotation in office).

²⁴⁵ *See supra* text accompanying notes 131-135.

every administrative action that an agency wishes to take. The same motivations that lead agencies to employ cost-benefit analysis, such as providing clear signals to other policy actors, should thus motivate them to prefer relatively more objective forms of cost-benefit analysis, at least if the cost of increasing objectivity is not so high and if others recognize the increase in objectivity. If information markets generally are indeed relatively objective,²⁴⁶ then predictive cost-benefit analysis also will be more objective than traditional cost-benefit analysis, in the sense that it will be less sensitive to the identity of the policymaker. Perhaps there remains a second-best argument that we should not seek to improve the objectivity of tools like cost-benefit analysis for fear that they might be misused, but I will assume simply that if cost-benefit analysis does exist, it might as well be as objective as possible.

Indeed, attempts to structure cost-benefit analysis by providing rigid *ex ante* rules about measuring costs and benefits form one means by which agencies already seek to improve the objectivity of their analyses. By providing an alternative means of assuring objectivity, predictive cost-benefit analysis limits the need for such *ex ante* rules. It therefore can help overcome the complaint that *ex ante* rules merely mask subjectivity, moving discretion from the evaluators of individual policies to those who determine the form of cost-benefit analysis. This complaint is at the heart of many debates about cost-benefit analysis and about regulatory policy more generally, and Part III.B.2 will explain why predictive cost-benefit analysis should appease both sides in a variety of such debates, even if considerable debate would remain about the circumstances in which the results of predictive cost-benefit analysis should constrain or motivate agency action. Before that, Part III.A will describe the mechanics of predictive cost-benefit analysis, and Part III.B.1 will compare it to other information markets. Finally, in Part III.C, I will describe a variant of predictive cost-benefit analysis, which I call comparative benefit analysis and may be useful when it is difficult to measure aggregate benefits in dollars.

A. The Mechanics

The mechanics of predictive cost-benefit analysis are straightforward, once an infrastructure for creating information markets exists. Suppose, for example, that the government wished in 2004 to conduct a predictive cost-benefit analysis of a proposed policy to reduce

²⁴⁶ See *supra* Part II.A.2.

arsenic in the water. The information market then would be initiated at that time, with a final prediction taken at some designated time, say September 1, 2004, and the government could then consider that prediction in deciding whether to enact the proposed policy. Then, at some far later point, say September 1, 2014, the government would perform a cost-benefit analysis of the policy, regardless of whether it in fact had decided to implement it. If the policy had been adopted, that cost-benefit analysis would estimate the costs and benefits that had occurred as well as any that might still result, discounting all to 2004 dollars; if it had not been adopted, the analysis would estimate what the costs and benefits would have been, discounting once again. The 2014 cost-benefit analysis would not have any direct policy effect, but it would determine prices at which securities from the 2004 information market would be liquidated.

The case for predictive cost-benefit analysis is largely the same regardless of how the information market component itself is constructed. Two innovations developed in Part I of this Article will be particularly helpful here, however. First, auctioning off the right to be the first participant helps solve a technical complication.²⁴⁷ The result of a cost-benefit analysis might be a positive number (net benefits) or a negative one (net costs).²⁴⁸ A security liquidated at a negative number would mean that the holder of the security would pay money to the government, so there should be no difficulty allowing trading of such a security. In effect, the buyer receives money from the seller.²⁴⁹ But no one would want to be the initial holder of the security. An auction, however, could allow negative bids, and when net costs are expected, the highest bidder would be the one with the least negative bid. The government would then pay to that bidder the corresponding amount to compensate for holding a security that ultimately is expected to result in payments to the government. The risk associated with holding the security would make this amount more than the expected eventual payment back to the government, with the difference between these amounts the government's expected cost from running the information market.

²⁴⁷ See *supra* text accompanying note 38.

²⁴⁸ An alternative approach would be to have two information markets, one used to estimate benefits and one used to estimate costs. This leads to various definitional questions, however, as some considerations may be seen either as reducing benefits or as imposing costs.

²⁴⁹ The market-scoring rule eliminates the need for such transactions to occur formally. See *supra* Part I.C.2.

Second, the two-phase information market may be particularly useful.²⁵⁰ Of the benefits of this approach already discussed, perhaps the most useful in this context is that it provides incentives to share information. Information market participants would publicize any costs and benefits that they expect future decisionmakers to take into account in their calculations. Similarly, a trader who concludes that a particular methodology for estimating a cost or benefit is flawed will have an incentive to trade on that and reveal the better methodology. The agency officials eventually conducting the ex post cost-benefit analysis could then consider any discussions during the information market in reaching their own conclusions, possibly saving some expense. The two-phase approach, however, achieves an even more dramatic savings of expense if, as might often be true, no transactions actually occur in the short second phase. Recall that with the two-phase information market, trades in the first phase are liquidated based on the estimate at the end of the second phase.²⁵¹ The possibility of the second phase is necessary to discipline the first phase, but once it exists, there would be no incentive to wait until that phase to act on a prediction. The eventual measurement, here the ex post cost-benefit analysis, is needed only to discipline any trades made in the second phase, so the measurement can be skipped if there are no trades in the second phase to discipline. In such a case, trades in the first phase would be resolved simply by the prediction at the end of the first phase.

Regardless of the form that the information market takes, it will result in a prediction of what an *average* decisionmaker would be expected to decide. If, for example, the market consensus is that Republicans would assess a regulation as having net costs of \$50 million and that Democrats would assess net benefits of \$100 million, then, assuming that it is equally likely that Republicans and Democrats will be in power ten years later, the market prediction would be net benefits of \$25 million. Traditional securities markets similarly assign probabilistic weights to different possible contingencies. A stock price for a company that may be acquired, for example, reflects the expected value of the company if it is and if it is not acquired and the probability of acquisition.²⁵² Much the same result would occur if the market scoring rule

²⁵⁰ See *supra* notes 155-156 and accompanying text.

²⁵¹ See *supra* text accompanying note 155.

²⁵² Finance scholars have developed methods of inferring probability distributions of future stock prices based on the prices of options. See, e.g., J.C. Jackwerth & Mark Rubinstein, *Recovering Probability Distributions from Contemporary Security Prices*, 51 J. FIN. 1611 (1996); Stephen A. Ross, *Options and Efficiency*, 90 Q.J. ECON. 75 (1976).

described above were used.²⁵³ While the incentive of participants is to come as close as possible to the eventual prediction, risk-averse participants will seek to minimize the risk associated with deviation from predictions and thus choose intermediate predictions rather than extreme ones.²⁵⁴

Because the information market predicts what an average decisionmaker would decide, it is not quite so important here as in traditional regulatory contexts to ensure that the actual retrospective decisionmaker is highly qualified. The decisions of highly qualified decisionmakers can be expected to differ from those of less qualified decisionmakers in two ways. First, less qualified decisionmakers are likely to be less predictable, with a greater variance of cost-benefit assessments. Such variance is of immediate concern in traditional administrative decisionmaking, because anomalous decisionmaking has a policy effect. With predictive cost-benefit analysis, however, the ex post evaluation has no direct policy effect, and thus the possibility that some will err too high and others too low is of little concern.²⁵⁵ Second, less qualified decisionmakers may systematically neglect subtle but important variables, or place too much weight on variables that more qualified decisionmakers would recognize are not as important as they might appear. This second effect remains of direct concern with predictive cost-benefit analysis, although the incentives of traders to provide information defending their choices²⁵⁶ may mute the effect.

Predictive cost-benefit analysis thus requires decisionmakers who are capable of understanding all dimensions of the relevant problem but not necessarily experts in the particular area. A virtue of using generalist decisionmakers, perhaps even decisionmakers not attached to any particular agency, is that they are less likely to shade their retrospective evaluations based on prospective agendas. Biased ex post decisionmakers present no hurdle to predictive cost-benefit analysis as long as traders cannot anticipate the direction of the bias ex ante. Generalist decisionmakers, however, reduce the chance that traders might anticipate systematic bias, for

²⁵³ See *supra* note 85 and accompanying text.

²⁵⁴ If such risk aversion were insufficient, the market scoring rule could be modified to provide incrementally greater rewards the closer is the prediction. A simple scoring rule with this property would square the difference between the prediction and the actual result. For example, if the actual result turns out to be \$50, then a prediction of \$30 would count as four times further away from the result than a prediction of \$20.

²⁵⁵ It is of modest indirect concern, however, because variability in decisionmaking will make participation in information markets riskier. As a result, the government would receive less (or pay more) in auctions determining who would become the first predictor. With risky information markets, the government must spend more to compensate participants for the risk that they are undertaking.

²⁵⁶ See *supra* note 156 and accompanying text.

example because decisionmakers in a particular agency from both political parties might be expected to share some normative commitments that generalists would reject.²⁵⁷ Regardless of whether specialists or generalists are the ultimate decisionmakers, an agency must have some procedure in place for selecting the relevant decisionmaker and ensuring that the cost-benefit analysis is indeed performed. The procedure could simply be that the agency will designate the decisionmaker; once again, even if the ultimate designation reflects idiosyncratic political preferences, that does not matter so long as traders *ex ante* do not know which idiosyncratic preferences will control.

B. Comparisons

I have identified objectivity as the chief virtue of information markets, and a particular type of objectivity, independence from the current administration's political and ideological biases, as the chief virtue of predictive cost-benefit analysis. This independence provides predictive cost-benefit analysis its principal advantage both over information markets used to make predictions of objectively verifiable numbers and over traditional cost-benefit analysis. The following subsections will explain why, after first assessing the disadvantages of predictive cost-benefit analysis relative to positive information markets and traditional cost-benefit analysis.

1. Normative vs. Positive Markets

Suppose an administrative agency is considering a nationwide regulation that would ban citizens from carrying concealed handguns. Robin Hanson has suggested that a conditional market would be an ideal tool for evaluating the effect of such a policy change, allowing a prediction to emerge about how the policy would affect the murder rate or the number of deaths involving handguns.²⁵⁸ It would be equally straightforward to use a predictive cost-benefit analysis to assess how decisionmakers would evaluate the policy retrospectively. Neither approach can guarantee the "correct" answer, as both approaches reflect only traders' predictions. An advantage of the objective approach, however, is that it gives traders an incentive

²⁵⁷ This raises the familiar debate about whether the beliefs of informed generalists should be preferred to those of specialists. See, e.g., Richard A. Posner, *Will the Federal Courts of Appeals Survive Until 1984?: An Article on Delegation and Specialization of the Judicial Function*, 56 S. CAL. L. REV. 761, 775-91 (1983) (providing a skeptical view of the need for specialized courts).

²⁵⁸ See Hanson, *supra* note 21, at 16.

to predict the actual effects of the policy rather than what decisionmakers will claim they believe the effects of those policy to be. If retrospective cost-benefit analysis is merely a form of “position taking,”²⁵⁹ then predictive cost-benefit analysis will reflect anticipated positions on policy effects rather than the policy effects themselves.

This objection suffices to justify a conclusion that predictive cost-benefit analysis may fail to uncover universal truths on the nation’s most controversial policy issues. But an information market predicting an objectively verifiable number seems equally irrelevant to government policy on such issues,²⁶⁰ because political actors are unlikely to yield to a prediction on a controversial issue, even an objective one. A positive information market might influence public opinion and thus indirectly political parties’ positions, but an agency is unlikely to ignore the administration’s positions on a fundamental issue like gun control solely because of an information market prediction. Predictive cost-benefit analysis may reflect average political preferences as well as hard-nosed analysis, but this is more a virtue than a vice in the political process. Predictive cost-benefit analysis will recommend against policies that all retrospective decisionmakers are expected to conclude had net costs, even if such conclusions are expected to be politically motivated rather than genuine.²⁶¹ Such policies would elude enactment in any event, and a tool that recommended them would lack credibility in the political process. In contrast, predictive cost-benefit analysis can verify whether an agency’s support for a policy is idiosyncratic, providing a politically useful signal to the President, Congress and the courts, rather than just providing objective analyses that in the long run might affect voters.

²⁵⁹ “Position taking” is defined in the political science literature as the “public enunciation of a judgmental statement on anything likely to be of interest to political actors.” DAVID MAYHEW, CONGRESS: THE ELECTORAL CONNECTION 64 (1974). When engaging in “position taking,” legislators are motivated by a desire to appeal to constituents or others, rather than by their underlying views on the issues. Whether legislative behavior constitutes position taking is often controversial. *See, e.g.,* Keith & Krhbiel & Douglas Rivers, *Sophisticated Voting in Congress: A Reconsideration*, 52 J. POLITICS 548, 564 (1990) (considering a particular example). With retrospective cost-benefit analysis, the more decisionmakers are expected to engage in “position taking,” the less the underlying merits will matter.

²⁶⁰ *See supra* text accompanying note 207.

²⁶¹ When the subjective assessments predicted by predictive cost-benefit analysis point in a direction opposite from what an positive information market would recommend, additional values besides whatever the objective market directly measures may be at stake. Suppose, for example, that an positive information market would predict that a ban on concealed weapons would reduce murders, *cf.* JOHN R. LOTT, JR., MORE GUNS, LESS CRIME: UNDERSTANDING CRIME AND GUN-CONTROL LAWS (2000) (arguing that laws permitting concealed weapons save lives), but a predictive cost-benefit analysis would anticipate greater costs than benefits. The discrepancy might reflect concerns other than lives saved, such as the intrinsic discomfort of knowing that many individuals have on their persons tools for ending lives. If these concerns are sufficiently powerful to affect conclusions about the net benefits of a concealed weapons ban, then perhaps they should influence policy.

Disingenuous policy evaluation should prove less dangerous with predictive cost-benefit analysis than with traditional administrative decisionmaking anyway. One reason is that disingenuousness will often roughly cancel out; an average decisionmaker is less likely than a partisan to read evidence selectively. More significantly, the retrospective assessor will be less susceptible to political and other forms of pressure than an administrative agency official making an actual policy decision. The only effect of a decision is on the traders, and, at least for most analyses, there would not be sufficient money at stake to cause traders to seek to influence the process.²⁶² Because the decision has no direct policy effect, interest groups should have far less interest in retrospective evaluations of decisions made a decade before than in the original decisions, and retrospective evaluations will receive less attention than actual decisions from other agency officials, the public, the President, and Congress. The incentives of ex post evaluators would be similar to the incentives of judges, including maintaining a reputation for high-quality work. At times, evaluators might seek to demonstrate solidarity with the present or past positions of a particular political party, or to use the retrospective valuations as a vehicle for advancing a world view. But unless evaluators from one party or associated with one view are more likely to act in this way than others, such tendencies will have no effect on predictions. Like judges, indeed perhaps even more than judges because judges have a direct influence on policy,²⁶³ retrospective evaluators may care about demonstrating careful and competent analysis.²⁶⁴

Predictive cost-benefit analysis's predictions ordinarily should not differ greatly from those of an positive information market evaluating the same effect of a policy. The predictions of information markets are themselves informed opinions, and informed opinions about future

²⁶² If a sufficiently large amount of money is at stake, then some means of preventing or minimizing trader influence of the eventual decisionmaker is necessary. A simple approach would be to ensure that the decisionmakers are independent. For example, Article III judges rather than political officials might themselves choose the decisionmakers, at least if the decisions are appealable to the courts. *See, e.g., Landry v. FDIC*, 204 F.3d 1125 (D.C. Cir. 2000) (finding that administrative law judges were inferior officers, and thus eligible for the alternative appointment approaches specified in the Appointments Clause, because their decisions are nonfinal). In addition, the decisionmakers might serve single, nonrenewable terms, and rules of procedure might prevent ex parte contacts between traders and decisionmakers.

²⁶³ A leading group of judicial politics scholars contend that judges tend to vote their "attitudes." *See* JEFFREY A. SEGAL & HAROLD J. SPAETH, *THE SUPREME COURT AND THE ATTITUDINAL MODEL REVISITED* (2002). Retrospective evaluators might suppress their attitudes more, given the inability actually to affect policy. On the other hand, the absence of policy consequences might liberate the evaluators to entertain even extreme positions.

²⁶⁴ *See* Sidney E. Shapiro & Richard E. Levy, *Judicial Incentives and Indeterminacy in Substantive Review of Administrative Decisions*, 44 DUKE L.J. 1051, 1053 (1995) (identifying "craft" and "outcome" as two goals that judges seek to achieve); *see also* Brett G. Scharffs, *Law as Craft*, 54 VAND. L. REV. 2245 (2001) (analogizing legal decisionmaking to other craft traditions).

expressed opinions of a policy's effect will rarely deviate greatly from informed opinions about the policy effect itself. This is especially true for run-of-the-mill administrative issues that elicit little passion in the present, and can be thus be expected to attract virtually no attention when retrospectively assessed in the future. There may be occasions when there is little need to undertake the cost of a retrospective cost-benefit analysis, because there is widespread agreement about what variable the agency should seek to optimize and there already exists a mechanism for measuring that variable. Most policies, however, have numerous effects, and no objective algorithm exists to convert all of them to a single scale. By averaging the costs and benefits that different decisionmakers are expected to assign to various effects of a regulation, predictive cost-benefit analysis avoids the subjectivity associated with *ex ante* determination of how to weigh different variables in assessing the regulation. In practice, a policymaking process that uses predictive cost-benefit analysis may be less prone to subjective influences than one that uses positive information markets.

The most significant advantages of predictive cost-benefit analysis over positive information markets, however, are technical. First, with predictive cost-benefit analysis, there is less of a danger that the demographics of traders will have a significant influence on the market outcome. Traders' ideologies are more likely to affect trading behavior when the proper valuation depends on ideologically contested assumptions. The task of predicting individuals' preferences, however, is not necessarily ideologically charged, even if the preferences themselves are largely ideological.²⁶⁵ Liberal and conservative traders might have different views about the magnitude of costs that a proposed regulation will impose on businesses, but they might nonetheless largely agree on how decisionmakers on average would assess those costs. Methodologies for measuring and predicting public opinion may be controversial, but such controversies tend to become ideological only in debates in which opponents each seek to claim public approval. Predictive cost-benefit analysis does not solve the problem of ideological influence entirely; traders may wrongly assume that others will agree with their own views,²⁶⁶ and they may believe that the passage of time will prove their perspective correct to the

²⁶⁵ See *supra* note 171 and accompanying text.

²⁶⁶ See *supra* note 159 and accompanying text (discussing the false consensus effect)

retrospective evaluator. It lessens the problem considerably, however, and the effect seems likely to be quite small as long as there is some ideological diversity among the traders.

Second, predictive cost-benefit analysis reduces the danger of spurious interpretations of effectively random phenomena. Conditional markets require a comparison of two closely related securities, so any noise in price determinations that randomly affects the two securities differently may be misinterpreted as reflecting the market's consensus about the effect of the relevant condition.²⁶⁷ Predictive cost-benefit analysis requires the issuance of only one security, so noise should not be a factor. Even if such analysis proceeded with two securities, one to measure costs and the other to measure benefits, there is no overlap between what the two securities measure, so the price difference is meaningful. There will be some instances, of course, in which the evaluation of a regulation is so close that whether predictive or traditional cost-benefit analysis anticipates slight net benefits or slight net costs is essentially random.²⁶⁸ With conditional markets, however, noise may be of comparable magnitude to the effect of interest, rather than just a small factor that may on occasion prove to be the decisive difference.

Third, predictive cost-benefit analysis avoids the selection bias problem, even if agency decisionmakers have information unavailable to traders. A conditional market tends to exaggerate the effect of the relevant condition because traders will reason that if the agency in fact adopts the policy, it may have information indicating that the policy will in fact be successful.²⁶⁹ This effect can occur only on a security that will have value only contingent on the government decision. With predictive cost-benefit analysis, whether the policy is adopted does not affect whether the retrospective assessment will occur, so selection bias is eliminated. Traders' assessments might still depend slightly on their estimate of the probability of adoption should they believe that retrospective cost-benefit analyses will be kinder to a regulation that is

²⁶⁷ See *supra* text accompanying note 77.

²⁶⁸ The policy decision is not necessarily unimportant in such cases. There may turn out to be very large net benefits or costs even if the average decisionmaker's evaluation would be very close to zero. When net benefits or costs are close to zero, however, there is little significance in whether there are a few dollars of net benefits or of net costs, because in such cases the probability that the regulation should be enacted will be close to 50% either way. As an analogy, many the outcome of the 2000 Presidential election might be seen as essentially random, in that very slight changes in conditions such as advertising expenditures could have made either candidate the clear winner. Though the decision itself is of obvious importance, in advance of an election, it may be far more important to ensure that election processes and machinery produce the "correct" winner when the population is split 75%-25% than when it is split 50.01%-49.99%, since in the latter case the outcome is essentially a tie. Cf. JAMES W. CEASAR & ANDREW E. BUSCH, *THE PERFECT TIE* (2001) (providing political scientists' perspectives on the 2001 election).

²⁶⁹ See *supra* notes 74-75 and accompanying text.

adopted than one that is not, all else being equal.²⁷⁰ Such an effect would make the government seem modestly wiser than the traders' views warrant and prevent predictive cost-benefit analysis from being entirely independent of the present administration. Predictions, however, would still be of retrospective cost-benefit analyses, and predictive cost-benefit analysis would still be far less dependent on the identity of current agency decisionmakers than cost-benefit analysis of the more traditional sort.

2. *Predictive vs. Traditional Cost-Benefit Analysis*

The principal disadvantage of predictive cost-benefit analysis relative to traditional cost-benefit analysis is the cost of the procedure itself. Predictive cost-benefit analysis may well be more expensive than traditional cost-benefit analysis, because the government must both subsidize the information market itself and then, at least in some cases,²⁷¹ pay for a retrospective cost-benefit analysis. Theoretically, however, this balance is indeterminate, for at least three reasons. First, the total cost will depend on how often retrospective analysis turns out to be necessary. Second, the retrospective component of predictive cost-benefit analysis might be cheaper or more expensive than prospective analysis.²⁷² Third, and most importantly, the government can choose the level of market subsidization. Should the government invest little in a predictive cost-benefit analysis, then predictive cost-benefit analysis might be cheaper than traditional cost-benefit analysis. Of course, the government can decide how much to invest in traditional cost-benefit analysis as well, with greater resources improving the report's accuracy and reliability.²⁷³

An analysis of both the cost and accuracy of the two forms of cost-benefit analysis thus depends on the respective levels of funding. Presumably, at some level of funding, predictive cost-benefit analysis will be as thoughtful as a typical traditional cost-benefit analysis. If enough

²⁷⁰ This might be true if the retrospective decisionmakers are prone to a status quo bias. See, e.g., Russell Korobkin, *The Status Quo Bias and Contract Default Rules*, 83 CORNELL L. REV. 608, 625-30 (1998) (discussing experimental evidence of such a bias). A status quo bias, of course, also affects other political decisionmaking, and information markets should not make the effects of such a bias any more severe than they already are.

²⁷¹ It would not need to do so when there are no transactions in the second phase of a two-phase information market. See *supra* text accompanying note 251.

²⁷² It might be cheaper if the policy is adopted and costs and benefits become easily ascertainable. On the other hand, it might be more expensive for two reasons. First, actual measurement may be more costly than speculation, depending on how carefully each is done. Second, when the policy is not adopted, it may be more complex to imagine a counterfactual world than it would have been to imagine a possible future world.

²⁷³ By "reliability," I mean the extent to which different practitioners of cost-benefit analysis would reach the same result.

money is at stake, each predictor will have an incentive to spend as much time considering every line item in the cost-benefit calculation as carefully as a traditional practitioner would. That level of funding may not be necessary to achieve comparable accuracy; even if each individual predictor has less knowledge than the practitioner of traditional cost-benefit analysis, the collective prediction that emerges may be equally accurate. Determination of the precise amount of funding needed to achieve comparable accuracy requires empirical evaluation, but the amount may well be greater than that for traditional cost-benefit analysis because of the redundancy associated with multiple traders' studying the same information.²⁷⁴ Concerns about the implementation cost of the analysis itself are already important to debates about traditional cost-benefit analysis, as executive orders require cost-benefit analysis only for sufficiently important regulations.²⁷⁵ Predictive cost-benefit analysis may magnify those concerns.²⁷⁶

A comprehensive assessment would provide a cost-benefit analysis of the forms of cost-benefit analysis. A numerical comparison of costs and benefits, however, is not yet possible, and the above analysis suggests that the benefits may be a function of the costs. To focus the analysis on the benefit of predictive cost-benefit analysis, assume that it is funded at a sufficient level so that it is as accurate as a typical traditional cost-benefit analysis, and retain the assumption that predictive cost-benefit analysis is objective. Predictive cost-benefit analysis's benefit with these assumptions is that the signal it provides depends less on the identity of the current agency officials. The superiority of predictive cost-benefit analysis along this dimension does not necessarily mean that this improvement is worth the expense, or that predictive cost-benefit analysis is sufficiently better than traditional cost-benefit analysis to justify greater reliance upon it in the regulatory process. Nor is this advantage necessarily sufficient to overcome what can be a healthy governmental reluctance to discard the procedures that it has traditionally employed. These are critical issues, but my ambition is only to explain the significance of predictive cost-benefit analysis's production of a signal that is relatively independent of agency officials' policy preferences.

²⁷⁴ Similar costs of redundant evaluation have long been recognized in the context of traditional securities markets. See Jack Hirshleifer, *The Private and Social Value of Information and the Reward to Inventive Activity*, 61 AM. ECON. REV. 561, 563-67 (1971).

²⁷⁵ President Reagan's original executive order required regulatory impact analyses only for "major rules." 46 Fed. Reg 13193 (1981).

²⁷⁶ Relatedly, if predictive cost-benefit analysis is expensive, it may magnify concerns about administrative ossification. See McGarity, *supra* note 222.

The defenses of cost-benefit analysis generally described above²⁷⁷ provide a straightforward preliminary explanation of why a tool that makes cost-benefit analysis relatively more ideologically neutral represents an improvement. To allow for a clean comparison, consider two extremes, which we may call *purely ideological* and *purely objective* means of conducting cost-benefit analysis. With purely ideological cost-benefit analysis, agency officials do not consider the underlying merits at all, except insofar as the merits are their actual ideological concerns. With purely objective cost-benefit analysis, agency officials' ideological preferences have no effects on the outcome of the cost-benefit analysis. Sunstein should prefer the purely objective approach, both because it thwarts heuristics to the extent that those heuristics have become impounded into agency decisionmakers' ideology,²⁷⁸ and because it eliminates the possibility that agency officials will exploit heuristics of members of the public to achieve ideological ends. Eric Posner should prefer it as well, because a purely ideological cost-benefit analysis would have no credibility with the President, Congress, and the courts, and it would thus fail to reduce the skepticism about agency motives that he identifies as limiting agency action. Of course, predictive and cost-benefit analysis do not represent these extremes, but the examples reveal that cost-benefit analysis will be more effective, the less its results confound the ideology of its practitioners.

This conclusion should be no surprise. To the extent that the results of cost-benefit analysis represent anything other than the merits or how people on average will view the merits, it becomes a noisier and less valuable measure. If a computer randomly added variance to cost-benefit measurements, the new measure would be less useful than the old one. Ideological perspectives and idiosyncratic views are not random, of course, but confounding cost-benefit measures with these factors is even worse. The regulatory system does not need an indication that agency officials favor a particular regulation; that is obvious from the agency's decision to

²⁷⁷ See *supra* Part II.A.1.

²⁷⁸ Sunstein seems to assume that heuristics affect agency decisionmakers as much as the public. See SUNSTEIN, *supra* note 117, at 9 (suggesting that cost-benefit analysis helps prevent agency decisionmakers from acting like "intuitive toxicologists"). But see Michael Abramowicz, *Toward a Jurisprudence of Cost-Benefit Analysis*, 100 MICH. L. REV. 1708, 1716 (2002) (reviewing SUNSTEIN, *supra* note 117) (arguing that well-informed decisionmakers are likely to be relatively unaffected by heuristics, but that they might seek to take advantage of the heuristics of others). To the extent that heuristics do affect such decisionmakers, they might not affect decisionmakers of all ideology equally. For example, liberals who are generally pro-environment might be more susceptible to an availability cascade resulting from an environmental catastrophe, because even an exaggerated policy response will generally be in the policy direction that they prefer. Similarly, conservatives who are generally pro-national security might be more affected by an availability cascade resulting from a terrorist incident, because the incident confirms their own fears.

adopt it. More worrisome, agency officials have the greatest incentive to shade the results of cost-benefit analysis in those cases in which others are least likely to share their views. Thus, to the extent that predictive cost-benefit analysis isolates average assessments rather than particular assessments, it improves the enterprise. To some critics of cost-benefit analysis, making cost-benefit analysis more useful should count as a deplorable development, for the same reason that an environmentalist might not welcome a new piece of heavy equipment that makes it easier to pave over wetlands.²⁷⁹ This critique is beyond my scope here, as my ambition is to show only that predictive cost-benefit analysis can be a better tool than traditional cost-benefit analysis, not to defend more broadly the rise of cost-benefit analysis.

Predictive cost-benefit analysis, however, may facilitate a reform that would mute at least some of the concerns that critics of cost-benefit analysis offer. The reform would be a loosening of the rules and norms underlying the practice of cost-benefit analysis,²⁸⁰ such as the requirements of OMB Circular A-94,²⁸¹ leaving practitioners of cost-benefit analysis free to make their own decisions about how to measure and discount costs and benefits. For example, such a reform would permit practitioners to count distributional effects of governmental regulations explicitly as benefits or costs, rather than merely discussing such effects, as Circular A-94 requires,²⁸² effectively making the controversial Kaldor-Hicks criterion²⁸³ the lodestar of

²⁷⁹ Alternatively, cost-benefit analysis could be irrelevant to agency decisions, in which case a superior approach to cost-benefit analysis accomplishes nothing. Matthew Adler and Eric Posner make this point:

Suppose, for example, that the public has no influence on political decision making and that all regulations are approved if and only if interest groups that benefit from them have more political power than interest groups that are harmed by them. Under these circumstances, it is hard to imagine a normative argument in favor of using cost-benefit analysis. The results of cost-benefit analysis performed by agencies would not influence their choice of regulations, and it is hard to see why any political actors would want agencies to use cost-benefit analysis in the first place.

Matthew D. Adler & Eric A. Posner, *Cost-Benefit Analysis: Legal, Economic, and Philosophical Perspectives—Introduction*, 29 J. LEGAL STUD. 837, 839-40 (2000).

²⁸⁰ Some agencies develop their own guidelines for cost-benefit analysis. See, e.g., U.S. EPA INNOVATIVE STRATEGIES & ECONOMICS GROUP, OFFICE OF AIR QUALITY PLANNING & STANDARDS, REGULATORY IMPACT ANALYSES FOR THE PARTICULATE MATTER AND OZONE NATIONAL AMBIENT AIR QUALITY STANDARDS AND PROPOSED REGIONAL HAZE RULE 12-43 (1997) (describing methodologies for calculating willingness to pay).

²⁸¹ See CIRCULAR A-94, *supra* note 123. The Circular could provide far more detailed guidance than it does, and much of the text consists of statements that are relatively uncontroversial. See, e.g., *id.* ¶ 7.a (“Economic analyses are often most readily accomplished using real or constant-dollar values, i.e., by measuring benefits and costs in units of stable purchasing power.”). Other requirements, however, appear to be attempts to standardize administrative practice through choices with which some reasonable people might disagree. See, e.g., Revesz, *supra* note 29, at 978 (suggesting that the Circular’s approach to discount rates “can produce perverse results”).

²⁸² The Circular provides:

The principle of maximizing net present value of benefits is based on the premise that gainers could fully compensate the losers and still be better off. The presence or absence of such compensation should be indicated in the analysis. When benefits and costs have significant distributional effects, these effects should be analyzed and discussed, along with the analysis of net present value.

CIRCULAR A-94 ¶ 10.

cost-benefit analysis if not of regulatory policy more generally.²⁸⁴ Similarly, practitioners might choose whether to measure deaths by the total number of lives lost or the total number of life-years,²⁸⁵ rather than having that choice dictated by a general policy.

It may at first seem perplexing why allowing such flexibility should be considered a reform given the preceding analysis. After all, guidelines constraining the practice of cost-benefit analysis seek to limit the danger that agency officials will view cost-benefit analysis as a hurdle to their regulatory goals and will thus choose whatever parameters justify their ideological conclusions. In addition, guidelines seek to assure that different agencies act consistently in performing cost-benefit analysis,²⁸⁶ thus seeking to create some coherence in a regulatory regime that critics have complained lacks rational inter-agency priority setting.²⁸⁷ Guidelines thus advance the same goal as predictive cost-benefit analysis, reducing the extent to which cost-benefit analysis results depend on the identity of those performing the analysis. That is precisely, however, why predictive cost-benefit analysis can allow for relaxation of standards. To the extent that predictive cost-benefit analysis succeeds in divorcing conclusions from the identity of particular agency decisionmakers, guidelines constraining the practice of cost-benefit analysis are less necessary. Flexibility in the retrospective cost-benefit analysis that would survive in a predictive regime does not present the same danger as flexibility in traditional cost-benefit analysis. Because the retrospective decisions do not directly affect policy, predictive cost-benefit

²⁸³ For a critique of the Kaldor-Hicks approach, see Jules Coleman, *The Normative Basis of Economic Analysis: A Critical Review of Richard Posner's The Economics of Justice*, 34 STAN. L. REV. 1105, 1106-12 (1982).

²⁸⁴ The existence of regulatory effects that cost-benefit analysis does not consider has led some proponents of cost-benefit analysis to acknowledge that there may be sound reasons for agencies to ignore the recommendations of cost-benefit analysis. See SUNSTEIN, *supra* note 117, at 22 (“[R]egulators might reasonably decide that the numbers are not decisive if, for example, children are mostly at risk, or if the relevant hazard is faced mostly by poor people, or if the hazard at issue is involuntarily incurred or extremely difficult to control.”).

²⁸⁵ See Richard H. Pildes & Cass R. Sunstein, *Reinventing the Regulatory State*, 62 U. CHI. L. REV. 1, 83-85 (1995) (advocating consideration of “quality-adjusted life years”); Tammy O. Tengs et al., *Five-Hundred Life-Saving Interventions and Their Cost-Effectiveness*, 15 RISK ANALYSIS 369 (1995) (recommending independently discounting each year of life saved); Richard Zeckhauser & Donald Shepard, *Where Now for Saving Lives?*, 40 LAW & CONTEMP. PROBS. 5, 11-15 (1976) (considering use of life-years rather than lives). The Environmental Protection Agency recently decided to end its practice of considering life-years rather than lives when senior citizens groups complained that this approach devalued the lives of the elderly. See John J. Fialka, *EPA to Stop ‘Death Discount’ to Value New Regulations*, WALL ST. J., May 8, 2003, at D3. That is, of course, exactly what using life-years does, but society may be justified in spending more to prevent the accidental death of a six-year old than of a ninety-six-year old. In both cases, the agency is merely delaying death, but with the six-year old, it is presumably delaying it much longer. Considering lives rather than life-years in effect values each year of elderly people’s lives at more than other people’s.

²⁸⁶ Existing guidelines, however, are in many instances not sufficiently detailed to avoid considerable inconsistencies. See, e.g., Matthew D. Adler & Eric A. Posner, *Implementing Cost-Benefit Analysis When Preferences Are Distorted*, 29 J. LEGAL STUD. 1105, 1146-47 (2000) (producing tables indicating differences in agencies’ valuation of life and selection of discount rates); cf. Edward R. Morrison, Comment, *Judicial Review of Discount Rates Used in Regulatory Cost-Benefit Analysis*, 65 U. CHI. L. REV. 1333, 1336 (1998) (“OMB’s guidelines appear to have had little effect on the discount rates that agencies actually use.”).

²⁸⁷ The critique is made most prominently by BREYER, *supra* note 114, at 21-33

analysis would result in predictions about how average decisionmakers would resolve any issues for which the retrospective decisionmakers will have discretion.

That predictive cost-benefit analysis provides a mechanism that substitutes for the effect of guidelines explains only why the loss of guidelines need not be mourned. There are sound reasons, however, that the elimination of guidelines should be celebrated once the need for constraint disappears. The guidelines are not neutral; rather, they necessarily impound value choices. The guideline against factoring distributional effects directly into cost-benefit analysis may well be sensible in existing practice, given the semantic awkwardness of calling a particular redistributive effect a “cost” or a “benefit.”²⁸⁸ Perhaps the omission of distributional effects even makes cost-benefit analysis seem more value neutral than under a hypothetical set of rules indicating that monetary effects on the wealthy should count only at some prespecified fraction as those on the poor.²⁸⁹ But any appearance of value neutrality from the omission of such rules is an illusion. More generally, whenever the rules of cost-benefit prevent an effect of governmental policy from being weighed in the balance, that decision carries a practical consequence, making the omitted variable less important in the practice of cost-benefit analysis itself.

Predictive cost-benefit analysis would not eliminate debate of controversial methodological issues, but it would allow the practice of cost-benefit analysis to continue without unanimous resolution. Consider, for example, the recent debate about relative position and cost-benefit analysis. Robert Frank and Cass Sunstein argued that by basing and costs and benefits on individuals’ willingness to pay, cost-benefit analysis substantially undervalued certain regulatory benefits.²⁹⁰ Individuals, they noted, care not only about their own economic well-being, but also their position relative to those they know.²⁹¹ Relative position matters more for some goods, like sports cars, than for others, like health care,²⁹² so an individual might be favorable to a regulation mandating a benefit for everyone even if the amount of income workers will lose as a result is greater than what individuals would be willing to pay in the absence of the

²⁸⁸ For an extended criticism of incorporating distributional concerns into cost-benefit analysis, see W. Kip Viscusi, *Risk Equity*, 29 J. LEGAL STUD. 843 (2000).

²⁸⁹ For an argument that cost-benefit analysis unjustifiably biases policy toward the rich, see David Copp, *The Theory and Rationale of Cost-Benefit Analysis*, 23 THEORY & DECISIONS 65, 74-77 (1987).

²⁹⁰ See Robert H. Frank & Cass R. Sunstein, *Cost-Benefit Analysis and Relative Position*, 68 U. CHI. L. REV. 323 (2001).

²⁹¹ *Id.* at 323-47 (reviewing survey and behavioral evidence).

²⁹² *Id.* at 351.

regulation.²⁹³ Thomas Kniesner and W. Kip Viscusi offer a detailed response,²⁹⁴ arguing in part that Frank and Sunstein overestimate the importance of relative position.²⁹⁵ Equally importantly, they argue that factoring relative position into cost-benefit analysis is not practicable, given the methodological difficulties in assessing how much relative position matters.²⁹⁶ They accordingly reject the incorporation of positional concerns and conclude that “the most important refinements one could make in the area of regulatory evaluation would be for agencies involved to adhere more to the framework of what is generally considered a carefully done cost-benefit study.”²⁹⁷

A potential counter to Kniesner and Viscusi’s complaint about the methodological difficulty of measuring the significance of positional effects is that the difficulty does not justify a conclusion that the effects are zero. The government should attempt, this counterargument suggests, to calculate the magnitude of positional effects as best it can, adjusting for different contexts and different groups of workers to the extent possible. Perhaps anticipating this counterargument, Kniesner and Viscusi argue that “the fact that the estimates are based on real market data for life and death choices rather than hypothetical thought experiments is a major contributing factor” to the acceptance of cost-benefit analysis.²⁹⁸ Although omission of a consideration from cost-benefit analysis as a practical matter constitutes a value judgment, the observation might be correct as a matter of public relations. Nuanced attempts by agencies to account for subtleties might make the value-laden nature of cost-benefit analysis more glaring.²⁹⁹

²⁹³ *Id.* at 372-73.

²⁹⁴ Thomas J. Kniesner & W. Kip Viscusi, *Why Relative Economic Position Does Not Matter: A Cost-Benefit Analysis*, 20 YALE J. ON REG. 1 (2003).

²⁹⁵ *Id.* at 12-16.

²⁹⁶ *Id.* at 12-15. For example, Kniesner and Viscusi offer the following critique:

[S]uppose we consider the effects of others' incomes on my behavior, and my true reference group is only my neighbor living in the house to the east. The researcher cannot know that only the income of one neighbor enters my decisions. Therefore, a statistical model incorrectly identifying all the houses on my block as my reference group will find that the average income on my block is statistically significant to my behavior because incomes are positively correlated across houses nearby.

Id. at 15 (footnotes omitted).

²⁹⁷ *Id.* at 2; *see also id.* at 19-22.

²⁹⁸ *Id.* at 23. They continue: “Moreover, given the sensitivity of the concerns [associated with the valuation of statistical lives], it is noteworthy that implicit value of life estimates derive from the value workers themselves place on risks of death as reflected in their labor market decisions.” *Id.* at 23-24.

²⁹⁹ Viscusi himself, however, has considered many subtleties affecting the valuation of life. *See, e.g.,* VISCUSI, *supra* note 30; W. KIP VISCUSI & JOSEPH E. ALDY, *THE VALUE OF A STATISTICAL LIFE: A CRITICAL REVIEW OF MARKET ESTIMATES THROUGHOUT THE WORLD* 15-16 (NBER Working Paper No. 9487, Feb. 2003) (considering whether a statistical method used to measure the value of life from wage data should apply the logarithm of wages). Perhaps Kniesner and Viscusi are skeptical of Frank and Sunstein’s recommendation for increasing estimating the statistical value of life not so much because they are opposed to complex methodologies that are difficult to verify, but because they believe that confidence in cost-benefit analysis will be undermined by any techniques not focusing on revealed preferences. Although this distinction may have technical merit, it seems

By omitting a variable, an agency can plausibly claim that it wishes to simplify the process of cost-benefit analysis and avoid making value judgments on a case-by-case basis.

Predictive cost-benefit analysis, however, provides a means through which a factor like positional effects can be considered without any *ex ante* agreement about the extent to which it should be a factor. Traders would seek to anticipate the importance that an average decisionmaker would attribute to positional effects, and an agency would not need to take a position on the value of positional effects in advance. The anticipated hypothetical average decisionmaker, of course, might not give positional effects precisely the weight that either camp in the debate believes to be appropriate. Predictive cost-benefit analysis, though, appears to satisfy the procedural concerns of both camps. First, it provides a mechanism by which positional factors, and more broadly any factors that some decisionmakers are likely to consider important in cost-benefit analysis, can be counted. If positional concerns receive less weight than Frank and Sunstein recommend, that would be because traders anticipate that decisionmakers on average would give these concerns less weight. Second, it avoids both giving agency officials an additional parameter that they can use to obtain the results that they seek and the appearance problem associated with such discretion. Kniesner and Viscusi could complain that predictive cost-benefit analysis would heighten the regulatory community's awareness of disagreement about cost-benefit analysis and thus perhaps undermine support even for a form of it that effects a compromise on such disagreements. Even if we find admissible arguments that we should hide the public from the sausage factory in which policy is made, this argument is of a more tenuous sort, suggesting that we should not improve the workings of the sausage factory lest the public realize that it is sausage they are eating.

By leaving the technicalities of cost-benefit analysis to predictors and retrospective decisionmakers, predictive cost-benefit analysis answers some broad critiques of cost-benefit analysis. Lisa Heinzerling, for example, has criticized cost-benefit analysis and numerical assessment of regulation more generally,³⁰⁰ arguing that numbers tend to obscure regulatory debates rather than elucidate them.³⁰¹ As Sunstein points out in response to a critique of a

unlikely that the public will tolerate some methodologies for calculating costs and benefits but not others, given the general lack of public awareness of methodological issues.

³⁰⁰ Lisa Heinzerling, *Regulatory Costs of Mythic Proportions*, 107 YALE L.J. 1981 (1998).

³⁰¹ Heinzerling argues:

[S]ome, probably many, people will be fooled into believing that numerical estimates of risks, costs, and benefits are

specific cost-benefit analysis that she offers in a separate article,³⁰² Heinzerling has no principled objection to considering costs and benefits in the abstract, but only to particular approaches that agencies have taken to counting them under a false claim to scientific truth.³⁰³ Predictive cost-benefit analysis creates an ongoing discussion, both among predictors and among retrospective assessors,³⁰⁴ about the appropriate measurement of costs and benefits. Predictive cost-benefit analysis results in a final number, but this number will reflect a compromise among various decisionmakers' anticipated views, rather than serve as a pseudo-objective justification for policy.

Predictive cost-benefit analysis also answers critiques that cost-benefit analysis focuses excessively on narrowly economic values, to the exclusion of aesthetic and other significant values.³⁰⁵ Proponents of cost-benefit analysis have developed techniques for considering such values, such as contingent valuation, which relies on surveys to assess how much citizens would pay to save a species or preserve a forest.³⁰⁶ Such surveys, however, are notoriously vulnerable to framing effects, with the amount surveyed individuals indicate that they would pay to save a forest bearing little relation to the size of the forest.³⁰⁷ Predictive cost-benefit analysis cannot determine what value should be assigned to a species or a forest, but it can allow for incorporation of such values into regulatory assessments without entrenching any particular flawed methodology. Aesthetic, economic, and other values may seem incommensurate, but agencies cannot avoid at least implicitly balancing different values in making decisions.³⁰⁸ Predictive cost-benefit analysis provides a method for achieving such balancing without putting a thumb or a value judgment on the scale.

impartial reflections of factual reality, in which case the likely result of increased reliance on quantification in setting regulatory policy will be that the side that best obscures the value choices implicit in its numbers will prevail. This will not produce more sensible regulation, but it will produce a more dishonest debate about regulation.

Id. at 2068.

³⁰² Lisa Heinzerling, *Markets for Arsenic*, 90 GEO. L.J. 2311, (2002) (responding to Sunstein, *supra* note 120).

³⁰³ See Cass R. Sunstein, *In Praise of Numbers: A Reply*, 90 GEO. L.J. 2379, 2383 (2002) ("Heinzerling's attack on cost-benefit analysis seems to be based not on a belief that costs and benefits are irrelevant, but on the willingness-to-pay criterion, which she identifies with cost-benefit analysis.").

³⁰⁴ Predictors will have incentives to critique methodologies and offer improvements to cost-benefit analysis. To the extent that they can convince other predictors that later decisionmakers will consider the improved methodologies, the predictors can profit from their innovations. See *supra* text accompanying note 250.

³⁰⁵ See, e.g., Thomas O. McGarity, *Regulatory Analysis and Regulatory Reform*, 65 TEX. L. REV. 1243, 1294-95 (1987).

³⁰⁶ See generally CONTINGENT VALUATION: A CRITICAL ASSESSMENT (Jerry A. Hausman, 1992) (collecting papers on contingent valuation). For a critique of the conceptual basis of these methods, see Boudreaux et al., *supra* note 31.

³⁰⁷ See, e.g., Frank B. Cross, *Restoring Restoration for Natural Resource Damages*, 24 U. TOL. L. REV. 319, 330-31 (1993).

³⁰⁸ See Matthew Adler, *Incommensurability and Cost-Benefit Analysis*, 146 U. PA. L. REV. 1371, 1383-89 (1998).

C. Comparative Benefit Analysis

Predictive cost-benefit analysis does not by itself, however, provide an escape from a related criticism of cost-benefit analysis, that the monetization of life diminishes human dignity.³⁰⁹ This claim does not insist that cost-benefit analysis will produce poor policy recommendations, but that the act of conducting a cost-benefit analysis itself may be debasing or have negative consequences. Partly in response to such claims, economists have offered alternatives to cost-benefit analysis. Risk-risk analysis, for example, compares the risks that a policy reduces with those that the policy will create.³¹⁰ This alternative, however, is of little use in a context in which risk is not the central issue of concern. Suppose, for example, that the Defense Department wishes to consider various alternative weapons systems.³¹¹ The relevant question in such a case is the extent to which different systems will advance national security goals. While it is possible to imagine translating this problem into cost-benefit terms or risk-risk terms, the conversion is awkward. Identifying the dollar benefits of an improvement to national security is a difficult exercise. Predictive cost-benefit analysis could perform this exercise and produce a compromise among what different people would believe, but the need to consider this factor will require traders to expend resources and impose risk on them. Because the aggregate defense budget is unlikely to depend on the recommendations of predictive cost-benefit analysis, the conversion is moreover unnecessary. All that matters in practice is the relative effectiveness of different weapons programs, not how improvements to national security should be valued in dollar terms.

With some modification, however, predictive cost-benefit analysis can easily measure comparative benefits without quantifying those benefits in dollar terms. Suppose that the Defense Department wishes to assess hundreds of possible weapons programs and is willing to spend some fixed sum to subsidize a predictive information market. To do so, the Department would designate one weapons program, presumably a well understood one such as the program

³⁰⁹ See, e.g., ELIZABETH ANDERSON, VALUE IN ETHICS AND ECONOMICS 190-95 (1993).

³¹⁰ For discussions of risk-risk analysis, see John D. Graham & Jonathan Baert Wiener, *Confronting Risk-Risk Tradeoffs*, in RISK VERSUS RISK: TRADEOFFS IN PROTECTING HEALTH AND THE ENVIRONMENT 1, 1-41 (John D. Graham & Jonathan Baert Wiener eds., 1995); Cass R. Sunstein, *Health-Health Tradeoffs*, 63 U. CHI. L. REV. 1533, 1538-52 (1996); and W. Kip Viscusi, *Risk-Risk Analysis*, 8 J. RISK & UNCERTAINTY 5 (1994).

³¹¹ The web site for the FutureMap program promises to use information markets to consider "analysis of the outcomes of advanced technology programs." <http://www.darpa.mil/iao/FutureMap.htm> (last visited June 19, 2003) (web site no longer available).

for building F-15 fighter jets, as the baseline program. It would then auction off the right to be the initial predictor for each of the other weapons programs and run an information market as before. After the market closed and the delay period of a decade or more has elapsed, the government would then retrospectively assess each weapons program relative to the baseline.³¹² A program, for example, that provided twice the national security bang for the buck of the baseline program would have a rating of 2.0, while one that provided only one tenth the bang for the buck would have a rating of 0.1.³¹³ The government would then divide the preset subsidy among the information markets corresponding to the various weapons programs in proportion to their ratings, and it would then further distribute the amount for each information market in accordance with a market scoring rule.³¹⁴

The approach easily could be adapted to any context in which the government must engage in priority setting. In the environmental context, for example, the government could consider the relative benefits of a number of different environmental programs, ranging from Superfund to the Endangered Species Act, perhaps even scrutinizing individual manifestations of such programs for individual sites or species. That would eliminate the need for contingent valuation procedures to translate environmental benefits into dollars. The relevant question instead would be how people do (or should) value different environmental goods, as well as what the social cost that these goods impose. As with any form of information market, agency officials and Congress might still ignore the predictive evaluations, assuming that no mechanism required priority setting to follow such recommendations. Comparative benefit analysis, however, provides a means for forward-looking analysis of priority setting, just as predictive cost-benefit analysis offers forward-looking assessments of whether individual policies should be adopted.

IV. CONCLUSION

In *Motor Vehicle Manufacturers Ass'n v. State Farm Mutual Automobile Ins. Co.*,³¹⁵ then-Justice Rehnquist's separate opinion suggested that the National Highway Traffic Safety

³¹² Once again, no retrospective assessment would be necessary in cases in which no transactions occurred in the second phase of a two-phase information market. *See supra* text accompanying note 251.

³¹³ An alternative would be to measure the relative total benefits of various programs without adjusting for cost. The government could then factor these benefit assessments into a predictive cost-benefit analysis conducted in the usual way.

³¹⁴ *See supra* text accompanying note 85.

³¹⁵ 463 U.S. 29 (1983); *see also* text accompanying notes 240-242 (discussing the facts and holding of the case).

Administration's about face on passive restraint policy should be allowed in part because the change in policy was the result of a change in Presidential administrations.³¹⁶ Rehnquist's realist approach may have been admirably candid, but critics insisted that although politics may affect agency policy decisions, courts should not simply abdicate their review responsibilities when confronting such decisionmaking.³¹⁷ Despite Rehnquist's view, courts do not yield to politics. Instead, they seek in administrative law cases to ensure public participation, consistency of regulations with statutory requirements, and careful consideration of policy decisions by agency officials. Doctrine, however, renders judicial review relatively deferential.³¹⁸ Part of the explanation for deference may be a Rehnquist-like theory of winner's spoils. Yet more often deference is justified by relative institutional competence,³¹⁹ that is by a conclusion that agencies make policy better than courts.

Agencies' general superiority at policymaking need not leave courts without a role. If doctrine grants courts too much deference, agencies may adopt policies that thwart congressional intent or represent ideological agendas that most of the population would reject. On the other hand, if it grants too little, then judges may do the same. The balance depends in part on the effectiveness and the objectivity of the tools that courts use to assess agency action. The only tool that courts currently use is plain old logic reasoning, as judges assess whether agencies have offered responses to significant counterarguments³²⁰ and whether agency interpretations are consistent with statutory authority.³²¹ Reasoning is a powerful tool but an inherently subjective one, and while judges presumably seek to avoid writing opinions that appear disingenuous,³²² reputation furnishes only a limited constraint. The extent of deference that doctrine affords

³¹⁶ *Id.* at 59 (Rehnquist, J., concurring in part and dissenting in part) ("The agency's changed view of the standard seems to be related to the election of a new President of a different political party. . . . A change in administration brought about by the people casting their votes is a perfectly reasonable basis for an executive agency's reappraisal of the costs and benefits of its programs and regulations.")

³¹⁷ *See, e.g.*, William F. Funk, *To Preserve Meaningful Judicial Review*, 49 ADMIN. L. REV. 171, 177-78 (1997).

³¹⁸ *See, e.g.*, *Chevron, U.S.A., Inc. v. Natural Resources Defense Council, Inc.*, 467 U.S. 837 (1984) (providing for deference to agency legal conclusions where those conclusions are reasonable).

³¹⁹ *See, e.g.*, Richard J. Pierce, Jr., *Chevron and Its Aftermath: Judicial Review of Agency Interpretations of Statutory Provisions*, 41 VAND. L. REV. 301 (1988).

³²⁰ Judges consider agency reasoning both under the hard look doctrine and under 5 U.S.C. § 553 (2000), which sets forth the requirements for notice-and-comment rulemaking. *See generally* Merrick B. Garland, *Deregulation and Judicial Review*, 98 HARV. L. REV. 505, 526-532 (1985) (discussing the relationship between hard look and procedural review).

³²¹ Courts also often invoke canons of construction to assess the reasonableness of agency interpretations, but judges sometimes disagree about their relevance and scope. *See, e.g.*, *Babbitt v. Sweet Home Chapter of Communities for a Great Oregon*, 515 U.S. 687 (1995) (providing an example of a case in which both sides used canons of construction to support their positions).

³²² *See supra* note 100 and accompanying text.

agencies thus reflects the difficulty of monitoring judges and the danger of government by judiciary.

To the extent that information markets provide for objective analysis, information markets both provide decisionmakers an opportunity to avoid poor decisions and offer others a metric by which to assess the decisions that are made.³²³ Information markets serve these two purposes both for agencies and courts. Just as an information market can help an agency avoid a decision based on a poor prediction, so too can it help courts avoid their own errors. An information market might lead judges to uphold policies that they would have struck down based on their own normative lights, or to strike down agency actions that they otherwise would have found to be within agencies' broad discretion. Equally importantly, just as an information market can allow courts and other governmental actors to assess agency officials, so too can such a market allow for scrutiny of judicial decisions. Although no court would be expected blindly to follow information market predictions and recommendations,³²⁴ information markets could make ideologically driven judicial reasoning more apparent, leading judges who care about their reputations for neutrality to hesitate before acting on their own political views.³²⁵ Used judiciously, information markets thus have the potential to improve the incentives and decisions both of agencies and of the courts reviewing them.

³²³ Information markets thus respond to Lisa Bressman's recent call for developing administrative law tools that reduce arbitrariness, rather than merely promoting accountability. See Lisa Schultz Bressman, *Beyond Accountability: Arbitrariness and Legitimacy in the Administrative State*, 78 N.Y.U. L. REV. 461 (2003).

³²⁴ This Article has advocated the use of information markets and predictive cost-benefit analysis as complements to rather than replacements for traditional decisionmaking forms. There are strong reasons for fitting new tools into existing legal frameworks. An interesting question, however, is whether it would be democratically illegitimate for the government to be required to follow the results of predictive cost-benefit analysis in all cases. Does democratic legitimacy depend on the representativeness of the *decisionmakers*, or just on the representativeness of the *decisions*? This question is not ordinarily asked since existing forms of decisionmaking require representative decisionmakers to achieve the goal of producing representative decisions. While the question is beyond the scope of this Article, it illustrates that predictive cost-benefit analysis can serve as a conceptual heuristic that sharpens analysis of the purpose of representative bodies.

³²⁵ If predictive cost-benefit analysis were commonplace, it would be easier to produce summary statistics on judges' decisions in judicial review. It would be possible to assess not only how frequently individual judges overturned agency action, but also whether they tended to do so in cases in which predictive cost-benefit analysis expressed skepticism of the decision. It might also be possible to design predictive cost-benefit analysis in such a way to identify judges' ideological tendencies. Existing studies use proxies for whether decisions are "conservative" or "liberal," depending in part on the identity of the party challenging the agency decision. See, e.g., Frank B. Cross & Emerson H. Tiller, *Judicial Partisanship and Obedience to Legal Doctrine: Whistleblowers on the Federal Courts of Appeals*, 107 YALE L.J. 2155 (1998). Predictive cost-benefit analysis could be modified so that in addition to the overall analysis, conditional securities predicted the result of the cost-benefit analysis conditional on the political party of the eventual retrospective decisionmaker. Such securities would provide an objective indication of expected differences in political parties' analysis of the decision. Across a large number of decisions, such information could improve analyses of the extent to which judges and courts are politically motivated.