

USING INFORMATION MARKETS TO IMPROVE PUBLIC DECISION MAKING

ROBERT W. HAHN* AND PAUL C. TETLOCK**

I. INTRODUCTION.....	214
II. AN INTRODUCTION TO INFORMATION MARKETS.....	219
III. GETTING BETTER INFORMATION FOR MAKING POLICY CHOICES.....	224
A. A New Approach	224
B. Extending the Framework.....	234
C. The Potential for Improving Fairness.....	238
D. Potential Problems with This Approach....	241
1. Versatility of the Approach.....	241
2. Project Governance.....	243
3. Measurement Issues	246
4. Market Design Issues	251
IV. A BENEFIT-COST ANALYSIS OF INFORMATION MARKETS.....	261
A. A More Informed Assessment of Policy Proposals	262
B. Greater Transparency and Accountability in Decision Making	263
C. Greater Availability of Assets for Financing Projects and Spreading Risks	265
D. Cost of Information Markets.....	266

* Co-Founder and Executive Director, American Enterprise Institute-Brookings Joint Center for Regulatory Studies; Resident Scholar, American Enterprise Institute.

** Assistant Professor of Finance, University of Texas at Austin, McCombs School of Business.

The authors would like to thank Michael Abramowicz, Matthew Adler, Jason Burnett, Peter Cramton, Christopher DeMuth, Robin Hanson, Randall Kroszner, Lester Lave, Robert Litan, Paul Milgrom, Terry Murray, Sam Peltzman, Eric Posner, Richard Schmalensee, Cass Sunstein, Philip Tetlock, and Scott Wallsten for helpful comments and Laura Goodman and Rohit Malik for research assistance. This article represents the views of the authors and does not necessarily represent the views of the institutions with which they are affiliated.

E. Comparison of Direct and Indirect Approaches for Information Markets.....	268
V. SIGNIFICANCE FOR POLICY DESIGN AND EVALUATION	270
A. Information Markets in the Policy Process.....	270
B. Solving Some Difficult Government Oversight Problems.....	273
C. The Potential Role for Government and Researchers.....	277
VI. CONCLUSION	283
VII. APPENDIX.....	287

I. INTRODUCTION

Many legal scholars have studied how to improve public decision making. Justice Breyer, for example, argues that technical problems could benefit from greater scientific expertise, and suggests using such analysis to help prioritize among competing social needs.¹ Cass Sunstein argues for the judicious use of cost-benefit analysis in a variety of areas, but also points out its limitations.² Sunstein's proposal, and the proposals of other scholars, would rely heavily, albeit not exclusively, on cost-benefit analysis to evaluate public policy decisions.³

Cost-benefit analysis is a tool used by decision makers to help inform the policy process. Cost-benefit analysis examines how different policies affect the overall level of net benefits to society, or benefits minus costs. A cost-benefit analysis may also be used to explore equity issues, examining how the distribution of net benefits varies across key groups, such as minorities or small businesses.⁴

1. *See, e.g.*, STEPHEN BREYER, *BREAKING THE VICIOUS CIRCLE: TOWARD EFFECTIVE RISK REGULATION* (1993); Stephen Breyer, Assoc. Justice of the U.S. S. Ct., *Economic Reasoning and Judicial Review*, AEI-Brookings Joint Center 2003 Distinguished Lecture (Dec. 4, 2003).

2. *See, e.g.*, CASS R. SUNSTEIN, *RISK AND REASON: SAFETY, LAW, AND THE ENVIRONMENT* (2002).

3. *See, e.g., id.*; KENNETH J. ARROW ET AL., *BENEFIT-COST ANALYSIS IN ENVIRONMENTAL, HEALTH, AND SAFETY REGULATION: A STATEMENT OF PRINCIPLES* (1986), available at http://aei-brookings.org/publications/books/benefit_cost_analysis.pdf.

4. A project can be said to yield net benefits if the economic benefits less the economic costs are positive. These calculations frequently ignore distributional implications; that is, a policy will typically have winners and losers. *See generally*

A fundamental problem with cost-benefit analysis of new policies is that the analysis is conducted before such policies are implemented. When conducting *ex ante* analyses, it is difficult to predict the future values of key variables that could be affected by a policy.⁵ For example, an analyst might predict that a worldwide carbon tax of \$100 per ton would reduce world GDP by 1% in 2010.⁶ How confident should we be in such a prediction?

In this paper, we present a new framework for addressing such uncertainty; this framework has the potential to substantially improve public decision making. We argue that decision makers can be more confident in analytical results if these results are based more directly on market data. Our framework introduces “information markets” that allow people to profit from superior knowledge about the future.⁷ For example, if an information market suggested that expected GDP would fall by 1% with a carbon tax,⁸ this estimate would theoretically incorporate all publicly available information about that policy’s effects. We also argue that if these information markets are designed well, information from the prices in these markets is likely to be much more accurate than other forecasts.

An information market allows individuals to purchase contracts, using real money, that yield returns to their owners contingent upon the uncertain outcome of a future event.⁹ With the

E.J. MISHAN, *COST-BENEFIT ANALYSIS* (2d ed. 1976); D.W. PEARCE, *COST-BENEFIT ANALYSIS* (2d ed. 1983); EDITH STOKEY & RICHARD ZECKHAUSER, *A PRIMER FOR POLICY ANALYSIS* (1978). For an analysis of the appropriate use of cost-benefit analysis in federal regulations, see ARROW ET AL., *supra* note 3.

5. We use the terms “project” and “policy” interchangeably. From the standpoint of a firm, it is undertaking a project. From the standpoint of the government, it is implementing a policy. Policies include regulations, research, standards, laws, programs, and public works projects.

6. For a similar prediction, see William R. Cline, *Meeting the Challenge of Global Warming*, in *GLOBAL CRISES, GLOBAL SOLUTIONS* (Bjorn Lomborg ed., forthcoming), available at <http://www.copenhagenconsensus.com/Default.asp?ID=165>.

7. See Robin Hanson, *Shall We Vote on Values, but Bet on Beliefs?* (September 2003) (unpublished manuscript, available at <http://hanson.gmu.edu/futarchy.pdf>); see also Michael Abramowicz, *Information Markets, Administrative Decisionmaking, and Predictive Cost-Benefit Analysis*, 71 U. CHI. L. REV. 933 (2004).

8. When we say the market may “know,” “believe,” or “suggest,” we are referring to the knowledge and beliefs of speculators in the market, which will be reflected in the market price. In this article, when we ascribe a view to the market, such as “the market expects,” we use such phrases as shorthand.

9. For a useful definition of information markets, see Justin Wolfers & Eric Zitzewitz, *Prediction Markets*, J. ECON. PERSP., Spring 2004, at 107, 108 (“Anal-

advent of the Internet, information markets are becoming more common. They are used in a number of contexts, ranging from assessing the likelihood that the Federal Reserve will raise interest rates to assessing the odds that a particular presidential candidate will be elected.

As an example, consider the online exchange at TradeSports.com. This exchange allowed its members to trade contracts that yielded \$10 to their owners if President Bush was reelected in November 2004. The contracts yielded \$0 if Bush was not reelected. The prevailing price of these contracts on July 6, 2004 (\$5.42)¹⁰ revealed the price at which the supply for Bush contracts equaled the demand for these contracts. Assuming that the market was efficient, this price suggested the probability of Bush's reelection was 54.2% on that date.¹¹

The idea of making greater use of information markets to promote social objectives is not new. Indeed, several scholars have suggested using information markets in a number of different contexts. Robin Hanson suggests that governments use information markets to identify whether particular policies will improve national welfare; he proposes relying on such markets exclusively, for example, when information markets predict a

ytically, these are markets where participants trade in contracts whose payoff depends on unknown future events."). The literature also refers to these markets as "speculative markets" and "betting markets." For a discussion of speculative markets, see Hanson, *supra* note 7, at 6 ("Most markets for stocks, bonds, currency, and commodities futures are called *speculative markets* because they allow people to bet on future prices by buying or selling today in the hope of later reversing such trades for a profit."). For a discussion of betting markets, see Paul W. Rhode & Koleman S. Strumpf, *Historical Presidential Betting Markets*, J. ECON. PERSP., Spring 2004, at 127; Hanson, *supra* note 7, at 6 ("*Betting markets* are speculative markets that trade assets that are specifically designed to allow people to bet on particular matters of fact, such as which horse will win a race.").

10. This price is as of 3:18 p.m. EDT, July 6, 2004.

11. Efficient prices equal "risk-neutral" probabilities, which are equal to actual probabilities if either the marginal trader is risk-neutral or the contract in question is not exposed to *systematic* (or aggregate) risk. It is possible, for example, that traders believe that future aggregate labor income will be lower if Bush is reelected. If they hold this belief, risk-averse traders will be willing to pay a premium for contracts that insure against the risk that aggregate income will fall in the future. The equilibrium price of the Bush contracts will therefore exceed the probability that he is reelected. However, through the clever use of other information markets, one can still infer the probability that Bush will be reelected, even in this more complicated example. The price of a contract yielding payments dependent upon future labor income conditional on Bush being reelected would identify the necessary adjustment between Bush's risk-neutral reelection probability and his true reelection probability.

policy will increase GDP.¹² Michael Abramowicz also notes the potential of such markets and advocates using them to predict the results of a cost-benefit analysis that would be done in the future by a designated expert.¹³ Ronnie Horesh offers a novel proposal for creating incentives for private investors to aid the government in achieving social policy objectives: government-issued “Social Policy Bonds” that pay a fixed amount after a certain performance objective is met, such as achieving a particular unemployment rate.¹⁴ His proposal has some similarities to ours, but does not address how to obtain information on the costs and benefits of a policy prior to implementation.¹⁵

To this point, no one has explicitly linked the improved information obtained through these markets to the subsequent implementation of policies that can maximize net benefits. Our central contribution is to propose an efficient way to implement well informed policy decisions.¹⁶ Our hope is that this general approach will induce the government to make more efficient decisions—and, in particular, curb the appetites of

12. See Hanson, *supra* note 7, at 10 (“Therefore, if one is willing to recommend policies that statistical studies suggest will increase (a time average of future) GDP, one should be willing recommend [sic] policies that speculative markets estimate will increase GDP, and so one should be willing to consider a form of government which relies more on such market estimates in choosing policies.”).

13. See Abramowicz, *supra* note 7, at 997 (“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.”); *id.* at 940 (“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.”).

14. See Ronnie Horesh, Social Policy Bonds: Explained in 4200 Words, <http://www.geocities.com/socialpbonds/aeu.html> (last visited Nov. 20, 2005).

15. Horesh’s proposal differs from ours in two other ways. First, it does not pay for incremental progress toward a specific goal. That is, the bond would only pay the fixed amount if a specific objective were achieved. Second, in the absence of costless bargaining and coordination among firms holding social policy bonds, individual firms will not have adequate incentives to help meet the performance objective under his proposal because of problems with free riding. By contrast, our approach explicitly deals with the free-rider problem.

16. By the word “efficient,” we mean relatively efficient but not necessarily perfectly efficient. Our proposal enhances the efficiency of existing policies, but may only be optimal in special circumstances.

politicians for introducing policies that have costs far exceeding their benefits.¹⁷

Our new approach builds upon the literature on information markets as well as the literature on the design of efficient policy mechanisms. Our claim is that the prices in information markets can inform the mechanism design process, thereby making previously infeasible mechanisms feasible for the policy maker. Specifically, information markets substantially broaden the scope for offering pay-for-performance contracts that yield useful information on the net benefits of a policy.

First, we show how it is generally possible to design markets whose prices will convey useful information on the costs and benefits of a number of policy choices, ranging from regulation to public works projects. Second, we describe one way of providing incentives for self-interested agents to implement policies that maximize net social benefits. Third, we show how information markets can be used to provide a stronger foundation for implementing a variety of government oversight mechanisms that, up to this point, have been stymied because of difficulties in estimating costs and benefits. We also show how legislators can use traditional budgetary controls in conjunction with information markets to exercise more effective oversight.

Finally, we identify and analyze the strengths and limitations of using information markets to help improve policy. While we focus on public-sector decision making, our analysis also holds for cost-benefit analyses used to inform private-sector and not-for-profit sector decision making.

This Article proceeds as follows: Part II provides a brief introduction to information markets. Part III provides a template for designing contracts that can supply better information on costs and benefits. Part IV discusses key costs and benefits of using information markets. Part V addresses the significance of

17. We are not so naïve as to think that a particular mechanism advanced by economists will be sufficient to stem the tide of inefficient government programs. On the other hand, the introduction of policy tools that increase transparency and provide better information on efficiency could help hold elected officials and bureaucrats more accountable. We do not, however, wish to imply that all regulation or taxation policies are inappropriate. For example, some environmental regulations can be expected to have benefits in excess of their costs if designed well. The key is to assess such regulations carefully with the best policy tools available. For a general discussion of the strengths and limits of economic analysis in designing more effective policy, see George J. Stigler, *Economists and Public Policy*, REGULATION, May–June 1982, at 13.

these contracts for policy design and evaluation, and suggests a new approach to legislative and regulatory oversight. Part VI presents our conclusions.

II. AN INTRODUCTION TO INFORMATION MARKETS

All markets can be thought of as providing some kind of information. We use the term “information market” to denote a market for contracts that yield payments based on the outcome of an uncertain future event.¹⁸ In this paper, we specifically consider contracts that yield financial payoffs to their owners contingent on the status of key policy variables. The relevant policy decisions could be private or public, but this paper focuses on public decisions. Information markets could provide information related to costs, benefits, net benefits, or the likelihood that a certain event will occur, such as the probability that a President is reelected.¹⁹ They can also be used to address potential market failures in the provision of information on public policy matters.²⁰

Information markets differ from traditional equity markets in that they are not typically tied to a claim of an ownership stake in a firm. Instead, the assets are claims that will pay off an amount that depends upon the state of the world, such as the monetary value of actual policy benefits.²¹ Although information markets for claims on benefits do not correspond directly to traditional equity markets for claims on corporate profits, there is a clear analogy. By monetizing policy benefits, information contracts allow organizations implementing policies to

18. Information market contracts are a subset of futures market contracts. See WILLIAM F. SHARPE, *INVESTMENTS* 521 (3d ed. 1985) (“Whenever something is ordered instead of purchased on the spot, a *forward* or *future* contract is involved. The price is decided at the time the order is placed, but cash is exchanged for merchandise later. . . . *Futures contracts* (*futures* for short), provide a standardized means of engaging in such transactions for agricultural and other commodities and for financial instruments and stock indices.”); see also MCGRAW HILL ENCYCLOPEDIA OF ECONOMICS 449 (Douglas Greenwald ed., 2d ed. 1994) (providing a similar definition of a futures market contract).

19. See generally Robert Forsythe et al., *Anatomy of an Experimental Political Stock Market*, 82 AM. ECON. REV. 1142 (1992).

20. See discussion *infra* Part V.

21. See Abramowicz, *supra* note 7, at 934 (“The securities in such a market do not serve as claims to corporate ownership, but rather offer payoffs contingent on the occurrence of some future event specified by the market’s sponsor.”).

transfer risks from these policies, just as traditional equity contracts allow corporations to transfer risks from their projects.

Information markets have already been used in a variety of contexts. The most well known information markets are for small-stakes political contracts; researchers at the University of Iowa, for example, conduct an electronic market for political futures contracts. In the corporate world, Hewlett Packard has experimented with information markets to forecast sales, while Eli Lilly has used these markets to help predict the success of pharmaceuticals. TradeSports.com offers information contracts in a number of areas including sports, politics, finance, law, entertainment, and even the weather. Goldman Sachs supports an exchange called *economicderivatives.com*, which hosts call auctions for securities based on economic indices.

The use of information markets is becoming more widespread.²² *Hedgestreet.com* recently received regulatory approval from the Commodity Futures Trading Commission to list several contracts based on economic indices; Macro Securities Research has developed another contract design that enables claims on economic indices to be bought and sold; and Case Shiller Weiss has constructed real estate indices that could provide the basis for information market contracts whose prices would convey the public's knowledge of housing prices.²³

The prices in information markets not only provide an additional source of information for use in cost-benefit analyses, they also dominate existing sources of information in many applications. This price discovery process seems to work well in theory and in practice. Las Vegas odds and point spreads predict the outcomes of sporting events better than sports ex-

22. In 2003, there was a controversial proposal to use information markets to predict terrorist events. A discussion of this proposal is beyond the scope of this paper. The issue is complicated because both terrorists and the government can affect the outcome. In addition, the government may not wish to disclose the kind of sensitive intelligence that an information market would reveal. For a discussion of the issues surrounding this proposal, see David M. Pennock, *The Good Side of the "Terror Futures" Idea (Yes, There is One)* (2004), <http://dpennock.com/pam.html>; Scott Wallsten, *Congress Shorts Future Terror-Fighting Innovation*, AEI-Brookings Joint Center (July 2003), <http://www.aei.brookings.org/policy/page.php?id=147>.

23. For example, the owner of such a contract might receive a payment proportional to the real estate index if a certain policy were implemented.

perts.²⁴ The prices in Iowa political markets are more accurate than concurrent opinion polls in forecasting elections 451 times out of 596.²⁵ Hewlett-Packard information markets beat official forecasts in predicting printer sales in two separate trials: fifteen times out of sixteen in one trial and six out of eight in the other.²⁶ Even play-money markets are a dominant source of information, outperforming four out of five columnists at forecasting Oscar winners in 2000.²⁷

Markets are generally the best available mechanism for gathering and aggregating dispersed information from private, self-interested economic agents.²⁸ People have something to lose if they are wrong, unlike in a standard cost-benefit analysis. That is, if the prices in information markets are poor predictors of the future, speculators have a direct economic incentive to get better information and trade on their superior information, moving prices toward the expected value of the contract payments.

The prices in properly designed information markets provide information on key cost-benefit analysis parameters in different states of the world. The economic theory and evidence cited above suggests the error of the market's implicit net benefits forecast will be no more than the error from a traditional cost-benefit projection.²⁹ Information markets not only provide in-

24. See PerformanZ Football Ratings, <http://tbeck.freeshell.org/> (last visited Nov. 20, 2005) (comparing experts, polls and the updated Las Vegas betting line in college and professional football and basketball); see also Raymond D. Sauer, *The Economics of Wagering Markets*, 36 J. ECON. LIT. 2021 (1998).

25. See Joyce Berg et al., Results from a Dozen Years of Election Futures Markets Research (Mar. 2003) (unpublished manuscript, <http://www.biz.uiowa.edu/faculty/trietz/papers/iemresults.pdf>).

26. 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, Mar. 2002), available at http://www.hpl.hp.com/personal/Kay-Yut_Chen/paper/ms020408.pdf; see also Ajit Kambil, *You Can Bet on Idea Markets*, HARV. BUS. SCH. WORKING KNOWLEDGE, Dec. 1, 2003, at <http://hbswk.hbs.edu/pubitem.jhtml?id=3808&t=innovation>.

27. See David M. Pennock et al., *Extracting Collective Probabilistic Forecasts from Web Games*, PROC. 7th ACM SIGKDD INT'L CONF. KNOWLEDGE DISCOVERY & DATA MINING 2001.

28. See Chen & Plott, *supra* note 26; Hanson, *supra* note 7; see also Berg et al., *supra* note 25. See generally THE HANDBOOK OF EXPERIMENTAL ECONOMICS (John H. Kagel & Alvin E. Roth eds., 1995); F.A. Hayek, *The Use of Knowledge in Society*, 35 AM. ECON. REV. 519 (1945).

29. This particular error measure is the mean squared error. The mean squared error from a traditional cost-benefit projection is the average squared difference

formation about the true mean of important quantities, but also have the potential to shed light on their entire distribution. So, for example, one could assess the probability that a policy would have a small, medium, or large impact on GDP, where the size of a small impact is defined as less than a 1% increase, and so forth. Abstracting from hedging considerations, the price of a contract that yields \$100 in the event that GDP growth falls below 1% will indicate the probability that the policy will cause GDP growth to fall below 1%.³⁰

Evidence from the psychology literature buttresses the case for using information markets instead of experts. In many instances, decision makers cannot outperform even simple statistical rules.³¹ In fact, many experts fail to outperform statistical rules based only on the experts themselves; the experts merely add noise to their own forecasting "rules."³² In contrast, financial markets outperform simple statistical rules; failure to do so would undermine the reason for their existence.³³ Other markets, such as sports wagering markets, also appear quite efficient at aggregating information. As previously mentioned, Las Vegas odds forecast sports game outcomes better than any expert, outperform any ranking or poll system, and perform as well as any betting rule.³⁴

This Article considers only markets for contracts that use real currency as a medium of exchange. Information markets for contracts that use alternatives, such as points, may also prove

between actual net benefits and the government's net benefit estimate. For the market forecast, the mean squared error is the average squared difference between actual net benefits and the implied market expectation of net benefits estimated from the current price.

30. With an estimate of the underlying distribution of net benefits, regulators would be better able to avoid regulatory projects that could be major failures, if that were viewed as desirable. This opens up the question of whether the government should be risk neutral in its choice of policy instruments, which is beyond the scope of this paper. See Kenneth J. Arrow & Robert C. Lind, Uncertainty and the Evaluation of Public Investment Decisions 60 *AMER. ECON. REV.* 364 (1970); see also DISCOUNTING AND INTERGENERATIONAL EQUITY (Paul R. Portney & John P. Weyant eds., 1999).

31. See Robyn M. Dawes et al., Clinical Versus Actuarial Judgment, 243 *SCIENCE* 1668 (1989) (highlighting the superiority of statistical judgments over clinical judgments).

32. See id.

33. See BURTON G. MALKIEL, *A RANDOM WALK DOWN WALL STREET* (8th ed. 2003); Eugene F. Fama, Efficient Capital Markets: A Review of Theory and Empirical Work, 25 *J. FIN.* 383 (1970); see also Michael C. Jensen & George A. Bennington, Random Walks and Technical Theories: Some Additional Evidence 25 *J. FIN.* 467 (1970).

34. See PerformerZ Football Ratings, *supra* note 24; see also Sauer, *supra* note 24.

flexibility has the advantage that the analyst can respond to changes in our understanding of cost-benefit analysis.

Although the approach Abramowicz suggests has merit, there is a great deal to be gained from using information markets to estimate parameters that would be *directly* useful for doing a cost-benefit analysis.³⁸ We will show how this can be done for a wide range of problems. In addition, we suggest an efficient mechanism for implementing such projects.

III. GETTING BETTER INFORMATION FOR MAKING POLICY CHOICES

This section presents a new approach to assessing and implementing policies that builds on two ideas. The first is to use information markets to help inform policy makers, the public, and interested parties about the likely benefits of policies. The second is to implement policies using a design that pays for actual results, so that the party implementing the policy has the appropriate incentive to produce the desired results. By connecting these two ideas, the decision maker can obtain valuable information on the costs, benefits, and net benefits of a particular policy prior to actually implementing that policy.

A. *A New Approach*

We begin by showing that it is possible to define a set of market mechanisms that yields useful information on net benefits, and that this information can be used to implement desired policies in an efficient manner. The essence of our approach is illustrated in Figure 1. The basic idea is for the government to gather information to decide if a policy proposal is worthwhile, then if it is worthwhile, to implement that proposal in a way that is efficient—that is, that maximizes net social benefits.³⁹

The approach consists of four steps. First, the government must specify and monetize all verifiable benefits accruing from

38. See discussion *infra* Parts IV, V.

39. We do not claim that our framework is optimal. For example, the performance-based policy design advocated in this Article has advantages and disadvantages relative to “book building” methods for selling project equity contracts. Book building, the informal process used by most U.S. investment banks to assess interest, allocate shares, and determine prices in initial public offerings, may not be feasible for the government because of the potential for unequal treatment of investors. A comprehensive comparison of our framework with all possible mechanisms is beyond the scope of this paper.

a possible project.⁴⁰ The specification of verifiable measures is a prerequisite for designing appropriate information market contracts.⁴¹ It is also a key to designing policies that reward those implementing a project based on actual performance.

Second, a contract that provides information on expected benefits with no change in policy would need to be issued. The expected benefits with no change in policy would provide the benchmark against which performance is measured if the policy is actually implemented. This benchmark is critical for getting a reasonable assessment of net benefits.

As part of this second step, the government would have a trading exchange list the information market contract for trading. The exchange would serve as a meeting place for buyers and sellers making contract orders on the not-yet-issued contracts.⁴² Generally, potential trades would occur when and if a particular policy is implemented. Unlike normal securities, however, government information market benefit contracts would be traded even if there is no change in policy because a project is not implemented.⁴³

Third, the government would conduct an auction for the right to implement and receive the monetized incremental benefits from that proposal.⁴⁴ The government would have the option of setting a minimum acceptable price for the auction.⁴⁵ This price would reflect unquantifiable benefits and costs as assessed by the government and would be set in advance.

40. See discussion *infra* Part III for a more formal model. A project could be thought of as any action that yields quantifiable and verifiable benefits. Examples include enforcing a policy or regulation, researching a new technology, or manufacturing a product.

41. We relax the requirement that the government must know the monetary value of benefits later in Part III. We also consider the case in which not all benefits or costs can be quantified.

42. The exchange could be subsidized or run by the government.

43. Trades in securities generally take place only when and if the security is issued and are referred to as "when-issued" contracts. Examples include treasury notes in the U.S. and initial public offerings in other countries, such as Germany.

44. The baseline for the incremental benefits would be determined by the estimate of benefits without the policy, which would be obtained from the information market. Without some estimate of the benefits that would result with no change in policy, it is not possible to measure the net benefits that are associated with a policy. It is, of course, possible to sell off a pay-for-performance contract without such a benchmark, but the resulting sale would not provide useful information on the incremental net benefits of the project.

45. In the economics literature, this is often referred to as a "reservation" price.

Fourth, the performance contract would be sold to the highest bidder if the auction revenues exceed the minimum acceptable price. That bidder would win the right to implement the policy and receive the monetized incremental benefits.⁴⁶ That right would be freely transferable, in whole or in part.

In a competitive bidding market with risk-neutral bidders, the revenue raised by this auction would approximate the project's expected net social benefits.⁴⁷ This is because different agents are competing for the right to get the social benefits from the project, as defined by the government. The winning bidder's profits would equal the project's social benefits less its private costs.⁴⁸

If the project does not go ahead because the bids are less than the minimum acceptable price, then trades in the information market benefit contracts would occur at the prices already agreed upon by market participants. The final value of the contracts would depend on actual measured benefits at some pre-specified point in time. If the project does go ahead, then these

46. Many types of firms could be winners in the auction. We hypothesize that certain firms with good access to information, funding, and an ability to resell securities would be prime candidates; examples include investment banks, venture capital firms, hedge funds, and institutional investors. Involvement of such firms would increase the likelihood that this process will work because these firms could serve as underwriters for the monetized benefits. Underwriters help solve the information asymmetry problem and coordinate buyers' and sellers' interests in the security in a manner analogous to their function in private securities markets.

47. The highest bid will not necessarily equal the highest bidder's valuation, but could be a close approximation, depending on how others value the contract. See Peter Cramton, *Ascending Auctions*, 42 EUR. ECON. REV. 745 (1998). This auction is efficient (that is, it allocates the good to the agent with the highest value) and optimal (that is, raises the most revenues) when there is perfect resale. It allows for the presence of private and public information on the value of a good. Furthermore, the auction is flexible enough to allow the sale of multiple benefits payments with different reserve prices. See Lawrence M. Ausubel & Peter Cramton, *Auctioning Many Divisible Goods*, 2 J. EUR. ECON. ASS'N 480 (2004). Depending on the structure of the social problem, it might be preferable for the government to auction benefits payments to multiple firms.

48. Perfect competition is not required. For example, if two risk-neutral firms value the project identically and there is Bertrand price competition, this result will hold. Without risk-neutrality, the winning bid will be less than expected net benefits by an amount reflecting the firm's cost of capital (that is, the project's systematic risk). Generally, the difference between the auction revenue and net benefits will be equal to the economic rent obtained by the winning bidder (that is, the difference between the project's net benefits for the high bidder and the second highest bidder). See William Vickrey, *Counterspeculation, Auctions, and Competitive Sealed Tenders*, 16 J. FIN. 8 (1961).

trades would never materialize because they were conditioned on the assumption that there would be no change in policy.

A stylized example shows how the framework could work in practice. Suppose policy makers are interested in improving average standardized test scores of public high school students in a major metropolitan area. In step one, the government decides that it is willing to pay a firm \$1 million for each point that average test scores improve in a year as a result of the firm's educational reform project.

In step two, to estimate the benefits, the government (or an exchange) can list a contract for trading that is issued only if there is no change in policy. If the contract is issued, it will pay \$ x when the average test score is x . Suppose the price of this contract is \$75 just before the government auctions off the project. This implies that the average test score is expected to be 75 with no change in policy.⁴⁹

In step three, the government auctions off the right to receive compensation for the education reform project. It offers to pay \$1 million for each one-point increase in test scores above 75 points, the market's estimate of points scored in the absence of the reform project. Assume the winning bid reflects the expected net benefits from the policy, and that this bid is \$7 million.

Finally, in step four, the government decides whether to proceed with the project. The government will fund the project if the winning bid and therefore the expected net benefits exceed some minimum amount. Suppose the government's decision rule is that net benefits must be greater than zero. In this case, the government will proceed with the project because the net benefits are estimated to be \$7 million.⁵⁰ The firm winning this auction would work with the schools to increase performance measures, with the firm paid on the basis of the performance increase it delivers.

The framework in Figure 1 allows the decision maker to implement only those projects with high expected net benefits. For these projects, the auction mechanism selects the firm that

49. The possibility that the market price is a biased measure of expected test scores with no change in policy is discussed below.

50. In the case of multiple proposals, the government only needs to choose whether or not to implement the proposal yielding the highest auction bid (less the reserve price). All other proposals will yield lower expected net benefits.

is likely to deliver the highest net benefits. In some cases, the decision maker may also want to know information on the expected benefits prior to making a decision, as well as the likely direct costs to the decision maker of undertaking the project.

By introducing one additional information market in benefits, the decision maker can get information on both the expected benefits and the expected costs of the policy. This idea is illustrated in Figure 2. There are two key differences between Figure 2 and Figure 1. First, in step 2 of Figure 2, we add an information market on the benefits that would result if the project were implemented. This allows an estimate of the expected benefits of the project based on the difference between the prices in the two information markets. If the government does decide to go ahead with the auction, it obtains an estimate of expected net benefits using the auction price. To obtain an estimate of expected costs, it can subtract the expected net benefits from the expected benefits.⁵¹

The second change in the figure is that the government has a choice whether to auction the right to receive benefits payments in step 3. We add this choice here because the government may not want to go ahead with the auction if the expected benefits from the project are small.⁵²

To illustrate how Figure 2 can be applied by extending the stylized example on test scores, suppose the government or an exchange issues a second type of contract that pays $\$y$ if the government chooses the pay-for-performance model and the average test score is y . Suppose the price of that contract is $\$85$ just before the government decides whether to implement the policy. Recall that the price of the contract associated with no change in policy was $\$75$. These prices would suggest that the market predicts a ten-point increase in average test scores (85 minus 75). Policy makers value this expected increase at $\$1$ million per point or $\$10$ million in total. The prices give the government advance knowledge of the likely benefits of a particular policy.

51. When costs are verifiable the government can measure costs directly rather than via benefits minus net benefits by creating contracts with payoffs dependent on the future costs with and without the proposed policy. This is not feasible, however, when costs are not verifiable.

52. This choice was omitted from Figure 1 to keep the framework simple. But, as discussed in Part III.B, there are many ways of extending the basic framework.

As before, assume the auction occurs and the winning bid is \$7 million. The government then has information on benefits and net benefits. By subtracting benefits from net benefits, it gets an estimate of the costs of the project. In this case, costs are estimated to be \$3 million.⁵³

Before dealing with some limitations of the framework, it is important to note a number of points about its advantages and general applicability. First, note that the government is only paying for the cost of the project. Paying for costs is better than paying for benefits because this requires less revenue to be raised by the government. Most politically acceptable means of raising revenues distort after-tax prices in the economy,⁵⁴ leading to inefficiencies in production or consumption, or both.⁵⁵ In the preceding example, the government only has to raise \$3 million to cover the expected costs of the project, as opposed to \$10 million to cover the expected benefits.

Second, the government does not need to monitor the costs of the project because the firm has every incentive to minimize costs. Although the government pays the firm on the basis of its actual performance—in this example, the improvement of average test scores—the firm also bears all costs incurred in the process of improving scores. The firm thus has an incentive to efficiently manage its productive inputs for improving test scores, which means minimizing its costs.⁵⁶

53. \$10 million in benefits minus \$7 million in net benefits equals \$3 million in costs.

54. Only lump sum taxes (sometimes called poll or head taxes) do not distort decision making.

55. See Martin Feldstein, *Tax Avoidance and the Deadweight Loss of the Income Tax*, 81 REV. ECON. & STAT. 674 (1999). Feldstein estimates that as much as \$2.06 per dollar of taxes is lost through deadweight efficiency losses. See *id.* at 678.

56. By design, this approach characterizes the problem of a social planner trying to maximize expected social surplus. See Mark A. Cohen & Paul H. Rubin, *Private Enforcement of Public Policy*, 3 YALE J. ON REG. 167 (1985) (applying this idea to private enforcement); Marc J. Roberts & Michael Spence, *Effluent Charges and Licenses Under Uncertainty*, 5 J. PUB. ECON. 193 (1976) (using a similar idea in the context of maximizing the net benefits of pollution control). That is, the firm maximizes social benefits less private costs, which are assumed to equal social costs. If there were externalities associated with the firms' investments, these could be dealt with in standard ways to internalize those externalities. See, e.g., WILLIAM J. BAUMOL & WALLACE E. OATES, *THE THEORY OF ENVIRONMENTAL POLICY* (2d ed. 1988). Externalities could also be dealt with by including them in the information markets as part of a social benefit function or social cost function. For example, suppose introducing traffic lights in a town reduced fatalities but added delays. Both of these factors could be considered in an expanded benefit

Third, the transferability of the project (or a component thereof) permits additional efficiency gains. This transferability allows the project manager to sell shares in the project conferring the right to some fraction of the benefits payments from a policy. Any party could purchase these shares, whether it was interested in managing the project differently, managing the project more efficiently, or perhaps terminating the project. Transferability allows for the transfer of project control between managers in response to changing market conditions.

These project shares are very similar to the corresponding information market contracts. Both project shares and information market contracts entitle their owners to the same cash flows—that is, a fraction of the project benefits.⁵⁷ The only material difference is that a project share confers project control rights upon its owner as well.⁵⁸ This difference is akin to the difference between voting shares and non-voting shares in public equity markets. Under certain conditions, the price between the project shares and information market contracts may differ.⁵⁹ This price difference poses no obvious theoretical or practical problems for the proper functioning of information markets.

The market for the transfer of projects would function analogously to the market for corporate control. Through efficient bargaining, the most capable firm will end up doing the project. The price of the security in the shares market and the par-

function. This equivalence implies that the profit-maximizing firm will solve the problem optimally from a social standpoint. In the interest of simplicity, we abstract from explicitly including the deadweight losses to the government associated with taxation. We will discuss this issue in more detail in a related paper. By purchasing the monetized benefits from the government, the firm is able to internalize the social benefits and reap the surplus. Of course, in a competitive bidding market, the firm must pay this surplus to the government in order to win the right to implement the project.

57. The precise relationship between shares and contracts will depend on the number of shares and number of contracts issued.

58. There is another, immaterial difference between the two claims. In the event of project liquidation, the owner of a claim on project benefits receives a fraction of the actual proceeds from the liquidation, whereas the owner of the information market contract receives the equivalent amount in cash from the counterparty to his transaction. Counterparty margins secure the payment in one case, whereas the value of the asset secures the payment in the other case.

59. See Luigi Zingales, *What Determines the Value of Corporate Votes?* 110 Q.J. ECON. 1047, 1052–55 (1995) (requiring that (1) there must be a contest over corporate control; (2) the participants in this contest must receive private benefits of control; (3) there must be some information asymmetry).

allel market for contracts would reflect the market's belief that the most efficient firm will undertake the project.

The market for transfer of projects differs from the market for corporate control in one crucial respect: The government could repurchase the rights to the project benefits. Government repurchase allows abandonment of a project or regulation, or assumption of control over project implementation. Thus, the government can revise its policy over time in light of new information about the social benefit function. Foreseeable changes in the social benefit function need not lead to government repurchase because these changes can be incorporated in the original government benefits payments at the outset. To return to the test score example, if the government expects the value of a one-point rise in test scores to increase from \$1 million to \$2 million, it can promise \$1 million per point initially and \$2 million per point after the date when it anticipates a change in the value of a one-point rise in scores. The market price for project shares will reflect the possibility of project takeover or termination, ensuring that all voluntary buyers and sellers in such a transaction will be adequately compensated for the transfer of control rights.

Government project takeover will not be necessary unless the social benefit function changes unexpectedly.⁶⁰ Even if new information about the likely success of the policy is revealed after the benefits payments have already been sold, the project manager will still optimally implement the project. For example, if it becomes clear that the project will not improve average test scores, the project manager will cease spending money without any prompting by the government.

In general, a project could give rise to other changes in government behavior that would not have occurred without the existence of the project. Thus, the market price of the education performance contract is a gross measure of benefits that includes the benefits from the policy changes that are expected to arise from the initial project. The measure of net benefits (auction proceeds) will also include the present value of all expected net benefits from anticipated policy changes.

60. The government may also repurchase the project when the government is the most efficient owner.

To see this point, reconsider the amount of money a project manager would be willing to bid in the step-three auction for the right to the benefits payments. Although the successful firm will not directly incur the costs of any other, unimplemented projects that could improve test scores, it will have to pay for these expected benefits in the auction for the right to the benefits payments. Upon winning the auction, the project manager becomes the claim holder of the benefits payments, which depend on the success or failure of this project and future projects. After the project manager commits to spending the funds on education reform, she can resell the rights to the benefits in the form of project shares to whoever can efficiently implement future education reform projects. In many instances, this repurchasing party is likely to be the government. Thus, the government will buy back the rights to the project shares when it appears that the manager has less use for these rights than the government. If no repurchasing party is likely to be interested in the rights, then the initial auction price and the subsequent project share and information contract prices will reflect only the benefits from the manager's project. In either case, the measure of project net benefits will be accurate.

The key point is that the benefits and net benefits measured in both the project share prices and the information market contract prices include all expected policy changes that result from the current decision. Information market prices signal the full marginal value of a social project by accounting for both its direct and indirect effects on welfare.⁶¹

Extending this argument, the anticipated benefits from one policy depend upon market participants' expectations about the government's future behavior *contingent on the policy being implemented*. If traders believe the government will choose more socially valuable projects following a particular initial policy, market participants will expect greater future benefits from that initial policy. Accordingly, the price of a claim on the policy's benefits, either a project share or an information market contract, will reflect these expectations.

Another advantage of transferability of shares is that the winning firm can hedge its investment in the project by selling

61. The prices can measure only the effects of policies that are not deterministically linked. That is, if one policy implies with certainty that another policy will be implemented, the government cannot construct an information market to measure the impact of each policy separately.

a fraction of the project payoffs. In the preceding education reform example, suppose that the project is implemented and awarded to the winning firm. The information contract yields one dollar for every one-point increase in test scores; so too, would rights to a 1-in-1,000,000 fractional share of the project benefits.

Suppose an educational testing company, say Princeton Review, is the winner of the right to collect \$1 million for each point that test scores increase above 75 points. Assume also that Princeton Review defines a share in the project as a 1 in 1,000,000 stake in the project's monetized benefits. For simplicity, suppose that the price of one share is equal to \$10, which is exactly the difference in price between the corresponding information contracts—that is, there is no premium for the control rights. Princeton Review could sell z shares to raise \$10 z (current price times z) and use these proceeds to finance the education reform project. If Princeton Review sells at least 300,000 shares, then it can fully finance the project (300,000 times \$10 = \$3 million).⁶²

Note that Princeton Review will earn only (\$1 million minus \$ z) per one-point increase in test scores, implying that increases in the number of shares it sells reduces its incentive to invest in increasing scores. There is a trade-off between allowing the firm to sell more project shares for hedging purposes⁶³ and wanting the firm to hold shares to maintain its incentive to invest the socially optimal quantity in education funds. For example, if Princeton Review sells 300,000 shares, it would then have expected profits of \$7 million (that is, the 10 point increase times \$1 million per point, less \$3 million). The net effect of agreeing to fund the project and selling these shares would be to pay nothing today in exchange for the right to receive

62. In general, the project manager can undertake myriad financial transactions to divide the rights to project control and cash flows. An important lesson from corporate finance theory is that these transactions will not affect efficiency unless they create or destroy value (for example, by affecting the implementation of the project). See Franco Modigliani & Merton H. Miller, *The Cost of Capital, Corporation Finance and the Theory of Investment*, 53 AM. ECON. REV. 261 (1958) (proving the sufficient conditions for this result to hold).

63. Compared to its profits under an unhedged position, Princeton Review's profits would be higher than they would have been if the actual point increase is less than ten, and lower than they would have been if the point increases exceed ten. By selling short, Princeton Review reduces the variation in expected profits associated with an unhedged position.

\$700,000 per one-point increase. In this case, Princeton Review would still have a strong incentive to make the project succeed. If Princeton Review performed research that led to further improvement of test scores, it would receive even more money from the government.⁶⁴

B. *Extending the Framework*

We consider three extensions of the framework illustrated in Figure 2: allowing for benefits to accrue over time; using the minimum acceptable price for different purposes; and adding information markets to supply more detailed information and improve hedging opportunities.

First, consider the problem of benefits accruing over time. If benefits accrue over time, the government can introduce a suitable discount factor into the performance measure. Suppose test scores improved over time. The contract could offer annual payoffs based on the average improvement in test scores each year. The price of this contract would now reflect the present discounted value of improvements in test scores, sending the appropriate signals to firms trying to value the project.⁶⁵ In principle, the timing of benefits is not critical, so long as the benefits are measurable.

Second, the government can use the minimum acceptable price to incorporate both information on unquantifiable benefits and costs and a budget constraint. If all costs and benefits are included in the framework, then an optimal decision rule would be to undertake the project if net benefits exceed zero. On the other hand, if there are some unquantifiable benefits and costs, then the government could add an amount to the minimum acceptable price that represents the government's best estimate of these unquantifiable costs and benefits.⁶⁶ The government could then impose the requirement that it will

64. See discussion *infra* Part III.C (examining the more general case of the incentive versus hedging trade-off).

65. If, for some reason, the social discount rate differs from the project discount rate, then the government should adjust future payments by the difference between the discount rates. This change induces the private sector to use the social discount rate.

66. If so-called unquantifiable benefits or costs were thought to be correlated with quantifiable benefits or costs, then the government could take account of this relationship directly in the benefit function rather than by adding an amount to the reserve price. In addition, if unquantifiable benefits or costs were expected to vary with the scale of the project and the expected benefits the project, then it would be more efficient to take this into account directly in the benefit function.

only sell the benefits payments if the winning bid in the auction lies above some pre-specified cutoff value. While the accuracy of these government estimates of unquantifiable costs and benefits is not guaranteed, there is some reason to think that a budget-constrained agency would behave reasonably in assessing these amounts.⁶⁷ Specifically, the agency would have some incentive, as under a regulatory budget, to maximize net benefits as it defined them.

The government can also use the minimum acceptable price to limit its expected payouts to the firm winning the contract.⁶⁸ Limiting payouts could be important if the government faces budget constraints. Suppose, as before, that the government estimates benefits of \$10 million from the increase in test scores, but it is not willing to pay out more than \$4 million. The government can then set the minimum acceptable price in the auction at \$6 million. Of course, even if the minimum acceptable price were \$6 million, this does not assure that the government would pay no more than \$4 million. To eliminate the risk entirely, the government could insure against the possibility that its payouts exceeded \$4 million.⁶⁹ If the government faces a budget constraint on several projects, the same logic applies.⁷⁰ It may, however, need to hold a series of auctions simultaneously to satisfy the overall budget constraint.⁷¹

67. A complete treatment of the unquantifiability issue is beyond the scope of this paper.

68. If the government wants to satisfy both a budget constraint and account for unquantifiable costs, then it should set its reserve price equal to the maximum of the reserve prices implied by the two cutoffs.

69. This transaction could entail the government swapping its obligation to pay the project benefits in exchange for the revenues from the net benefits auction and a bond with a present value of \$4 million. Corporations and financial institutions, such as investment banks and insurance companies, routinely conduct similar transactions in "over-the-counter" markets. There is little reason to expect the government would have difficulty executing such a transaction.

70. The same logic on budget constraints also applies to firms. A possible application for firms would be capital budgeting in a world in which they face capital constraints.

71. A good example is the case where the government considers two projects, both of which have expected benefits of \$10 million. Suppose the government cannot spend more than \$8 million total on the two projects. Then the government could impose a rule that both benefits payments are sold only if the combined revenues from the two auctions are greater than \$12 million (and each auction raises positive revenues). If the combined revenues fell short of \$12 million, then the government could implement whichever project raised greater revenues,

At first glance, it may seem inconsistent to allow the decision maker discretion in choosing a minimum cutoff value for auction proceeds while at the same time forcing her to use a quantitative rule to implement the project. These seemingly inconsistent directives merely require the decision maker to impose a binding minimum price prior to the auction of benefits. The good is not sold unless some bidder meets the policy maker's pre-specified price.⁷² This price must be determined in advance of the information market for net benefits and the auction to ensure there is no ex post manipulation of policy decisions. Although the minimum price could be private or public during the information market for net benefits and the auction, announcing that price has the advantage of transparency.⁷³

Third, we consider the role that additional information markets play in providing information for the decision maker and parties participating in the markets. Our framework can be modified to allow policy makers to assess the sensitivity of estimated net benefits to changes in assumptions about the value of policy benefits, which would be particularly important in cases where the benefit function is uncertain, or key parameters in the benefit function are uncertain.⁷⁴ The sensitivity analysis could also be useful when the government's budget is limited.

Consider, for example, the case where a government facing budget constraints is unsure whether to value a statistical life at \$1, \$5, or \$10 million.⁷⁵ In step one, the government could specify all three possible benefits values. In step two, the govern-

assuming these revenues exceeded \$2 million. If neither project raised \$2 million in revenues, then no project would be implemented.

72. There could be other restrictions on auction implementation, such as invalidation of the auction in the event of corruption or collusion in bidding.

73. On the other hand, a private reserve price will induce uncertainty whether bidders in the auction will meet the reserve price. This uncertainty may enhance the liquidity of the information market for benefits contracts by increasing the likelihood that trades will take place. Thus, uncertainty may be desirable in this instance.

74. More generally, information markets can help with sensitivity analysis on key parameters, such as costs, benefits, and net benefits to the extent that useful proxies for these measures are available. *See generally* M. GRANGER MORGAN & MAX HENRION, UNCERTAINTY: A GUIDE TO DEALING WITH UNCERTAINTY IN QUANTITATIVE RISK AND POLICY ANALYSIS (1990) (discussing different approaches to dealing with uncertainty).

75. *See generally* W. KIP VISCUSI, FATAL TRADEOFFS: PUBLIC AND PRIVATE RESPONSIBILITIES FOR RISK 17-98 (1992) (discussing the value of a statistical life); W. Kip Viscusi & Joseph E. Aldy, *The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World*, 27 J. RISK & UNCERTAINTY 5 (2003) (reviewing different estimates of the value of a statistical life).

ment would establish liquid information markets for contracts based on the benefits from each possible benefit value. In step three, the government would conduct multiple simultaneous auctions—one for each benefit value—in order to estimate the net benefits under the three alternative assumptions about the value of a statistical life.⁷⁶ The government could choose its benefit payment based on how the amount of the winning bid changes as a function of the value assigned to a life. Only the benefits payments corresponding to the value chosen by the government would be implemented in step four, and only the trades in the market corresponding to these benefits payments would take place.

More generally, this modified framework allows the government to assess how changes in its declared monetized benefit function would affect the estimated social surplus from a given policy. In contrast to sensitivity analyses performed in traditional cost-benefit analyses, this sensitivity analysis adjusts net benefits to account for general adjustments by firms to changes in the assigned monetary value of benefits.

The ability of information markets to accommodate a wide range of assumptions regarding the monetary value of benefits is particularly important for certain applications where this value is uncertain and likely to vary across the voting public. For example, the implied value of a statistical life depends strongly on the context in which the life is saved. Important factors include whether the program saving the life is voluntary or involuntary, whether the risk is “dreaded” or not, and whether the risk is familiar or not.⁷⁷ Also, different people value statistical lives differently, suggesting that the optimal policy depends upon the person selecting that policy. Sensitivity analysis using information markets allows policy makers to observe the dependence of optimal policies on different preferences.

The government may want to consider introducing more information markets on both benefits and net benefits if that information were useful. For example, if a proposed regulation

76. In fact, the government could explore an entire range of possible values of a statistical life by asking firms to submit demand functions for the project.

77. See generally Paul Slovic et al., *Regulation of Risk: A Psychological Perspective*, in REGULATORY POLICY AND THE SOCIAL SCIENCES 241 (Roger G. Noll ed., 1985) (discussing more thoroughly how and why the implied value of a life can vary).

were expected to reduce mortality and morbidity, the government could issue contracts that yield payments to their owners based on the level of these two quantities with and without the policy. Similarly, the government could use multiple information markets if it wished to compare more than two policy alternatives, such as a comparison of the status quo with a regulation that provides varying degrees of flexibility in implementation.

Finally, the government could also introduce an information market contract to estimate net benefits, as measured by auction revenues, prior to the auction. Using the education example, an information contract could pay \$1 for each \$1 million of revenues that would be received, contingent on proceeding with the auction. There are a number of benefits to having such a market, including providing early information to decision makers on the expected benefits and costs of a policy, allowing hedging opportunities for firms bidding in the auction, and reducing the likelihood of underbidding in the auction due to the winner's curse.⁷⁸

C. *The Potential for Improving Fairness*

It is often difficult to estimate the distributional impact of a proposed policy and, indeed, many cost-benefit analyses fail to do so.⁷⁹ Fortunately, information markets can be used to help assess the impact of proposed policies on equity. Furthermore, under certain circumstances, information market contracts can assist policy makers in implementing equitable policies without harming efficiency. In economic terms, it may be possible for the government to approximate *any* efficient policy using transfer payments determined by information market prices. Although there are practical difficulties with obtaining reliable information on the equity implications of individual projects, it may be possible to obtain better information on general sets of policies and use this information in deciding on appropriate redistributive policies. We use two tax policy examples to illus-

78. The winner's curse occurs in auctions where the bidders have private information about the value of the good. The winning bidder is likely to be the one with the most optimistic private information. Adjusting for this fact, all bidders shade their bids downward to prevent over-bidding. See generally RICHARD H. THALER, *THE WINNER'S CURSE: PARADOXES AND ANOMALIES OF ECONOMIC LIFE* (1992).

79. See, e.g., ARROW ET AL., *supra* note 3 (explaining the desirability of addressing such issues when data are available).

trate the potential of information markets to estimate the distributional impact of policy.⁸⁰

The first example is a case in which the government wants to know the impact of a certain tax proposal on a specific income group, say those at or below the tenth percentile in income. The government could design two contracts that yield \$1 for every \$1,000 in that group's average after-tax income: One would be conditional on no change in policy and the other would be conditional on a change in policy. From the difference in the prices of these contracts, one could infer the estimated impact of the proposal on those with the lowest incomes. Perhaps more important, if the poor have access to such markets, they could hedge the risk that their income would fall by assuming a short position on this contract.⁸¹ In addition, the contract price could inform policy makers seeking to offset the distributional impact of a project.⁸²

A second example reveals how government can meet redistributive objectives without sacrificing efficiency by using information markets to uncover valuable information about the effect of a policy in advance of its implementation.⁸³ Consider a proposed income tax cut in a simple hypothetical society with just two types of workers. The current tax rate is 50% for the more productive type of worker, who earns \$1 million per year before taxes (\$500,000 after taxes); the tax rate is 0% for the less productive type of worker, who earns \$10,000 per year (before and after taxes). Policy makers would like to reduce all mar-

80. We choose tax policy because distributional issues are salient in this area. Assuming availability of data, these ideas are applicable to other areas of policy, such as the impact of a policy on educational test scores for different income groups.

81. We recognize that the poor often have limited access to asset markets, limited means to invest in these markets, and limited awareness of opportunities in these markets. Notwithstanding these considerations, this particular strategy requires very little capital or specialized knowledge. Policy makers could alert the poor to this opportunity to purchase inexpensive income insurance, or could even purchase it on behalf of the poor. *Cf.* Cass R. Sunstein & Richard H. Thaler, *Libertarian Paternalism Is Not an Oxymoron*, 70 U. CHI. L. REV. 1159 (2003).

82. See ROBERT J. SHILLER, *THE NEW FINANCIAL ORDER: RISK IN THE 21ST CENTURY* 237–41 (2003) (emphasizing the value of hedging markets for addressing distributional concerns, but not suggesting hedging markets conditional on policy implementation); Hanson, *supra* note 7 (proposing conditional information markets, but not for hedging purposes).

83. Our example relies on lump sum transfers that may not be politically feasible.

ginal tax rates to zero in order to reduce distortionary impacts on work effort, but they are concerned that this will increase inequality.

Suppose policy makers would also like to retain the current absolute difference in after-tax income between the two types of workers (\$490,000). To redistribute income efficiently, policy makers want to utilize lump sum taxes. But policy makers do not know the amount to tax the productive type after the tax cut, because this requires advance knowledge of the productive type's behavioral response to the tax cut, that is, his increase in work effort. To solve this problem, policy makers could create an information market contract that yields a payment equal to the average pre-tax income of the productive type after the tax.⁸⁴

Suppose the price of the contract is \$2 million, implying the productive workers will double their pre-tax incomes in response to the tax cut. Then the government could levy a per-capita lump-sum tax of \$750,000 on the productive workers and give the proceeds to the unproductive workers (assumed to be equal in number to the productive workers and assumed not to alter their work effort choice). This transfer would leave the productive workers with \$1.25 million in after-tax income and the unproductive workers with \$760,000 in after-tax income. The two types can share equally in the efficiency gains only because the government is able to uncover valuable information about the effect of its policy in advance of the policy's implementation.⁸⁵

These examples reveal that information markets can provide new insights into the equity impact of proposals. There is a broader question, however, of how equity concerns should enter into decision making. We do not have an answer to that question, but offer some suggestions for accommodating equity concerns in different ways.

The government could, for example, add some kind of equity constraint to its decision-making algorithm. If a policy has rela-

84. Assuming there are many productive workers, the contract's payment is only barely influenced by each individual worker's effort choice. Thus, the contract poses negligible moral hazard problems.

85. Although the example appears to rely heavily on the absence of uncertainty, this procedure can be adjusted to accommodate random shocks to each income group. See SHILLER, *supra* note 82, at 237–41 (noting that the contracts described above are ideal hedging vehicles for the productive workers, enabling exact redistributive transfers even in the presence of uncertainty).

tively modest expected net benefits, then the government could add a requirement that it not be harmful to groups at the lower end of the income range. This constraint need not apply to policies with higher levels of net benefits.⁸⁶ Government could also incorporate the equity constraint by altering the benefits payment formula to account for the distribution of the benefits.⁸⁷ Still another approach is for the government to measure the equity impact of particular policies, but not to address them on a case-by-case basis. Instead, government could address questions of income redistribution more broadly through fiscal policy.⁸⁸

D. *Potential Problems with This Approach*

This section examines four groups of potential problems: versatility of the approach, project governance, measurement issues, and market design issues.

1. *Versatility of the Approach*

We begin by examining the versatility of the approach from the government's perspective. The approach is *designed* to be performance based, and in that sense, would represent a paradigm shift in many areas, including regulation. This performance-based paradigm has been endorsed by a number of leading regulatory scholars, including Justice Breyer.⁸⁹ Our approach differs from that of Justice Breyer in that it relies less on the expertise of government regulators and more on actual markets.

Although the benefit function explicitly pays out on the basis of performance, it is still possible for the government to pursue traditional command-and-control regulation if it so desires.

86. The rationale for exempting policies with high net benefits is that they yield a "bigger pie," some of which could be used for redistribution if desired. The definition of "high" is somewhat arbitrary. From an economic standpoint, the most logical measure would be return on investment.

87. The downside to this approach is that the government could not separately identify the expected impacts of the policy on efficiency and equity, which is crucial for certain applications.

88. See, e.g., ARROW ET AL., *supra* note 3, at 8 (suggesting that regulation is a blunt instrument for redistribution); see also SHILLER, *supra* note 82, at 149–64 (suggesting ways of reducing inequality).

89. See BREYER, *supra* note 1, at 61–65 (describing and advocating a partially performance-based system).

This goal, if desired, could be met by modifying the framework slightly. Suppose, for example, the government wished to introduce an environmental standard with limited flexibility.⁹⁰ The general procedure would remain as described in Figure 1. The government could, however, auction the benefits to the project with a restriction on its implementation—for example, by specifically requiring a technology-based approach.⁹¹

Furthermore, the winning bidder would be allowed to contract with the government to actually implement the regulation if this contract were the most cost-effective way to do so. In this scenario, a private firm could negotiate with the Environmental Protection Agency to pay for the EPA's administrative, monitoring, and enforcement costs in exchange for a promise from the EPA to implement the regulation as outlined by the firm. The winning bidder can be thought of as a project manager, consultant, and overseer hired by the government. The performance-based compensation scheme gives the managing firm the appropriate incentives to implement the project efficiently. Thus, there are ways to use the framework to implement traditional and nontraditional regulatory approaches.⁹²

Another issue that may be a problem from the government's perspective, or an agency's perspective, is the potential liability incurred. Many government agencies face strict budgetary requirements and may be reluctant to take on projects where they are promising a very large or unbounded amount based on undetermined future benefits. In the discussion of the minimum acceptable price, we suggested one approach for dealing with this problem. There are others that are worth considering. For example, the government could adjust the amount it is willing to pay for results—say for a point increase in average test scores. Alternatively, the government could consider put-

90. A technology-based standard for power plants might be an example. For a discussion of different kinds of standards, see Cary Coglianese & Gary E. Marchant, *Shifting Sands: The Limits of Science in Setting Risk Standards* (Regulatory Policy Program, Working Paper RPP-15, 2003), available at <http://www.aei-brookings.com/publications/abstract.php?pid=392>.

91. Examples of technology-based approaches to regulation include requiring the installation of scrubbers on power plants or catalytic converters on vehicles.

92. For example, one could apply the same idea to market-based approaches to reduce pollution such as the allowance trading used to limit pollution from sulfur dioxide emissions. In this case, a firm such as a trading exchange with expertise in trading systems could be hired to manage and oversee the process. The framework has the advantage that it provides measures of benefits, costs, and net benefits.

ting a cap on payouts. For example, a contract could promise \$1 million per point for test score increases over 75 points with a maximum payout of \$20 million (achieved when average test scores reach 95 points). Of course, this yearly benefit cap will blunt the incentive for the project manager to increase test scores if scores approach the 95 point cutoff. Finally, some of the problems the government wishes to address have potentially long time horizons or long latency periods. Examples include climate change and reducing exposure to carcinogens. Long latencies present no problem in principle for the information market mechanism described here. There are liquid financial markets for 30 year U.S. Treasury bonds and for 30 year residential mortgage backed securities. Traditional equities have a potentially infinite maturity.

2. Project Governance

We consider two problems related to the incentives of the project manager: one if the manager does not act as a profit maximizer and a second related to excessive hedging and the incentives of the manager.

A problem could arise if the project manager decides not to behave as a profit maximizer. This could lead to inefficiencies, such as lower average test scores or wasteful expenditures. The potential for such behavior is not new to information markets or corporations. For example, firm management may be risk averse, have empire building goals, receive private benefits from control, or be interested in diverting cash flows into their own account.⁹³ There are already a number of mechanisms, including corporate charters, external financial contracts, internal labor contracts, and the market for corporate control, which address these issues.⁹⁴

93. See Philippe Aghion & Patrick Bolton, *An Incomplete Contracts Approach to Financial Contracting*, 59 *REV. ECON. STUD.* 257 (1992); Michael C. Jensen, *Agency Costs of Free Cash Flow, Corporate Finance, and Takeovers*, 76 *AM. ECON. REV.* 323 (1986), available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=99580.

94. For a discussion of these and related issues, see Douglas Gale & Martin Hellwig, *Incentive-Compatible Debt Contracts: The One-Period Problem*, 52 *REV. ECON. STUD.* 647 (1985); S.J. Grossman & O.D. Hart, *Disclosure Laws and Takeover Bids*, 35 *J. FIN.* 323 (1980); Sanford J. Grossman & Oliver D. Hart, *One Share-One Vote and the Market for Corporate Control*, 20 *J. FIN. ECON.* 175 (1988); Bengt Holmstrom, *Moral Hazard in Teams*, 13 *BELL J. ECON.* 324 (1982); Bengt Holmstrom, *Moral Hazard and Observability*, 10 *BELL J. ECON.* 74 (1979).

Conversely, if the winning bidder tries “too hard” to maximize profits, this could create other problems. For example, because the winning bidder is the only firm with the incentive to solve a particular social problem, it could threaten to fail to perform on the contract with the government after it has won the auction unless the government increases the benefits payments. Even if the government deems this threat to be credible, several mechanisms could curb the firm’s incentive to renegotiate the terms of the contract.⁹⁵ First, the government could put some requirements in the contract specifying minimum levels of performance, such as requiring a certain level of expenditures over time. Second, the government could refuse to do business with such a firm in the future. Third, consumers could punish the firm by refusing to buy its products once it became clear that the firm was harming the public by not acting in good faith; alternatively, consumer concerns could drive politicians to act. Finally, a regulatory body, such as the Federal Trade Commission, could be enlisted to identify cases in which firms significantly underperform and impose appropriate penalties.⁹⁶ All of these actions could negatively impact current and future revenues of the firm, and also affect its share price if it is publicly traded.

Another problem could arise during the implementation of the project if the project manager creates a hedge that reduces the firm’s incentive to obtain social benefits. Suppose the project manager in our earlier example would like to completely hedge the project’s risk by selling 1 million project shares at \$10 each. If the price of shares remained constant as the manager sold its shares, the manager would then receive an up-front cash flow of \$10 million in contract sales minus the \$7 million in auction payments. If this hedge is possible, it would create a huge problem, because the project manager would have *no* incentive to spend the money on the education reform project. For each point increase in test scores, the manager would receive \$1 million from the government, but also have to pay out \$1 million to fulfill its obligations to project shareholders (\$1 for each share sold). In this scenario, the manager would put forth

95. Because this is a bilateral bargaining situation, the government could try to hold up the firm as well. See generally Benjamin Klein et al., *Vertical Integration, Appropriate Rents, and the Competitive Contracting Process*, 21 J.L. & ECON. 297 (1978).

96. The agency would, of course, need to be acting within the law, and this should only be done as a last resort.

no effort, implying that the holders of project shares would expect to receive no payoff.

Fortunately, there is a natural market mechanism that would prevent this outcome. So long as the project manager must register its issuance of securities with the Securities and Exchange Commission (SEC), she would never be able to sell 1 million contracts at \$10 each. The reason is that investors would observe that the project manager is selling its entire stake, thus blunting her incentive to undertake the project. A rational investor would be willing to pay nothing for the benefit contracts if the manager were selling its entire stake, because the investor would correctly anticipate that the manager would not bother with the costly project. This implies that the manager would make more money by holding on to the contracts, doing the project, improving the test scores, and receiving the benefits payments from the government.

A far-sighted manager will keep a large stake in the project to inspire confidence in her investors.⁹⁷ The project manager does not care about the investors per se, but she does care about the price received for stakes in her project. She will sell the project or hold claims on project benefits, depending upon which option will create greater profits. Recall that profits have been designed to equal social net benefits. So we have really shown that the firm that is the sole claimant to the project will take the necessary actions in securities markets to maximize social net benefits. The manager's project finance decision is analogous to the decision faced by an entrepreneur choosing what fraction of her project to sell to venture capitalists. This analogy suggests the mechanism will work in practice just as well as it does in theory.

There is an additional mechanism that could prevent the project manager from shirking her implementation of the project. The shirking problem exists only to the extent that the manager holds a net short position in the net benefit contract market. This short position implies that other contract holders are effectively financing the project and have strong incentives to monitor its completion. These contract holders could elect a board of directors and use standard corporate governance techniques to

97. In fact, the price of the project stake will already reflect the market's expectation of benefits given the fraction of the project that the auction winner intends to sell.

keep project management from enriching themselves at the expense of the project shareholders. For example, bonuses for management could be based on project performance and management could be replaced if it performed poorly. The governance mechanisms that operate effectively at the corporate finance level could operate equally effectively at the project finance level.

3. *Measurement Issues*

There are several problems with measurement that the government may want to address. First, it needs to decide on appropriate indicators of performance for particular projects. Second, it may want to assess the amount of noise in key variables and consider the gains from reducing uncertainty in measurement error.

If policy makers have access to multiple unbiased measures of benefits, they will need some guidance on which to use. We think a reasonable approach is to specify performance measures as broadly as possible, subject to the constraint that projects can be implemented on an efficient scale without wasteful coordination costs.⁹⁸ Typically, we would advocate specifying benefits on a project-by-project basis, depending on the objectives of the project. Examples might include having performance measures for improving educational outcomes or for reducing fatalities from auto crashes.

Having one firm buy a right to a broad package of benefits, such as GDP, is unlikely to be efficient for several reasons. First, a single firm is not likely to be well informed about the entire set of possible social projects needed to increase GDP. Firms specialize in different areas. Returning to our education example, Princeton Review has expertise and knowledge in improving children's test scores. It presumably lacks expertise in other areas, such as improving the environment.

Second, if a firm wins a performance contract based on increases in GDP, it will have to incur potentially huge costs to coordinate thousands of firms. If there were several performance contracts based on GDP, the coordination problems would be even worse. Narrowing the benefit function reduces these coordination problems. So, for example, Princeton Re-

98. The problem is analogous to determining the optimal size for a firm. See R.H. Coase, *The Nature of the Firm*, 4 *ECONOMICA* 386 (1937).

view may not need to coordinate with any other firms to efficiently increase educational test scores.

Third, benefit functions linked to what a firm or small group of firms can produce allow firms to hedge their risks more easily. Princeton Review may not have much control over GDP, but it may have a lot of control over test scores. If firms are able to hedge their risks, then projects will become more attractive and more firms may participate in the auction. A firm may decide to avoid buying a residual claim on GDP, because it would fear bankruptcy if a general economic downturn occurred. For example, if Princeton Review owned a residual claim on GDP and GDP unexpectedly fell by just 1%, Princeton Review would lose roughly \$100 billion, resulting in certain bankruptcy.⁹⁹

To limit the bankruptcy risk, Princeton Review could hedge against risks in GDP beyond its control, but only if these risks are traded. This is really an argument in favor of introducing narrow project-based contracts, which would allow firms to hedge the risks from other projects that are beyond their control. Moreover, if the decision maker funding the performance contract divided it into narrow projects, then firms implementing the projects would not face any of the practical problems associated with hedging.¹⁰⁰

In addition, there is an important theoretical reason to worry about the efficiency of reselling parts of project benefits to others. The prospective buyers of these hedging contracts could infer that the seller is not motivated by hedging, but by unfavorable private information.¹⁰¹ This problem is particularly acute when only broad GDP contracts are available because the prices of these contracts do not convey information about the

99. Princeton Review could limit its risk by owning only a small fraction (say 0.1%) of the residual claim on GDP, but this would diminish drastically its incentive to reform education, rendering the "performance-based" contract ineffective.

100. These include transactions costs, liquidity risk, and basis risk. *See generally* JOHN C. HULL, *OPTIONS, FUTURES AND OTHER DERIVATIVES* (5th ed. 2003).

101. *See* George A. Akerlof, *The Market for "Lemons": Quality Uncertainty and the Market Mechanism*, 84 Q.J. ECON. 488 (1970) (classic treatment of how private information can lead to the sale of only "lemons" in the used car markets); *see also* Stewart C. Myers & Nicholas S. Majluf, *Corporate Financing and Investment Decisions when Firms Have Information that Investors Do Not Have*, 13 J. FIN. ECON. 187 (1984) (discussing models of adverse selection and its detrimental impact on efficiency).

individual components of GDP being resold. This adverse selection problem could be mitigated by creating information markets for the narrow, project-based contracts. A third party, such as the decision maker, could sell these contracts.

Another problem with issuing only broad contracts is that the information markets may not be very efficient. Distinct pieces of information about U.S. GDP are held by millions of consumers and thousands of firms in the economy. If there are even small fixed costs to participating in the GDP information market, then it is unlikely that everyone with information about GDP will participate. Even though together they may hold substantial information about GDP, most individuals hold an insignificant fraction of the total information. The information available to individuals is likely to address only their specific economic circumstances or expertise; in fact, it could be project-based. Individuals would search for more efficient ways to exploit specific project-based private information by trading in securities other than GDP.¹⁰²

Fourth, there is little cost to the decision maker in specifying a set of narrower benefit functions. Suppose the decision maker is interested in GDP and that she knows that increases in GDP are best accomplished through increasing educational achievement and making better environmental investments.¹⁰³ Then, performance contracts could be auctioned off for education and the environment, reducing the systemic risks to which bidders would be exposed if they were rewarded in terms of GDP.

In certain cases, there may be reasons to broaden the scope of the performance contract. First, an excessively narrow performance measure could mean that the project manager ignores important externalities. For example, the education reform project described above could also reduce law enforcement costs in areas where uneducated high school dropouts resort to crime rather than finding a job. In this case, policy makers would need to value the impact of improved test scores

102. Traders exploiting the connection between the GDP and project-based markets would ensure that the GDP contract incorporates all project-specific information. Of course, the project-specific securities must exist for this mechanism to work.

103. If the decision maker does not know which categories of GDP are most important, she can design derivative securities to estimate the impact of establishing information market contracts in different areas on GDP. For more details on how this would work in practice, see discussion *infra* on the value of climate change research.

on criminal activity and reward these improvements. This would give the firm implementing education reform an adequate incentive to focus its efforts on areas with the greatest benefits from improved test scores. Second, there may be somewhat greater transaction costs associated with implementing more performance contracts. But these costs should not be significant in comparison to the benefits of writing narrower performance contracts.

In short, there are many factors that should be considered in determining the optimal scope of a performance contract. Our preference would be for the performance measure to include all benefits that can be reasonably quantified and no more. In addition, it should include externalities to the extent reasonable. The same firm should reap the social benefits and bear the full costs of its actions.

Another potential problem with the information markets framework relates to measuring the key quantities. Fortunately, for many policy interventions, there appear to be reasonably accurate proxy measures for benefits to form the basis for information market contracts. Without accurate, verifiable measures of such quantities, it is difficult to construct information markets that provide useful information.

Consider the previous example on education reform. The estimated increase in test scores is 10 points if the reform project is implemented. Because that result may be subject to a great deal of measurement error, it would be useful to know the error bounds. For example, if the variance in realized benefits is quite high, then the policy could just as easily result in a 5-point or 15-point increase in test scores. In this case, the government may not wish to go ahead with the reform project.¹⁰⁴

The government can establish a market for another contract to estimate the uncertainty in benefits. Consider an option contract that entitles its owner to a payment of \$1 for each point increase above 10 points if the policy is implemented. The price of this option will reflect the likelihood that actual benefits will

104. Some have argued that the government should behave as though it is risk-neutral in its policy decisions. Rather than making this normative judgment, we note that real policy makers often behave as though they prefer less risky alternatives. Without endorsing this behavior, we outline a process that a risk-averse policy maker could adopt to accommodate his preferences. For a good introduction to these and related issues, see *DISCOUNTING AND INTERGENERATIONAL EQUITY* (Paul R. Portney & John P. Weyant eds., 1999).

deviate from expected benefits, allowing the policy maker to infer the degree of uncertainty in realized benefits.¹⁰⁵

Stock option prices in traditional financial markets already provide analogous measures of implied volatility for the prices of the underlying stocks. The Black-Scholes options pricing formula describes the mathematical relationship between the expected volatility of the underlying asset price and the price of an option on the asset, enabling policymakers to calculate the expected uncertainty in net benefits.¹⁰⁶ Such measures of uncertainty are particularly useful in situations where actual benefits have a large variance, such as climate change policy.

If information markets reveal that measurement error is a problem, they can also point to a solution by helping to assess the value of additional information.¹⁰⁷ If, for example, the relationship between greenhouse gas emissions and rising sea level is too uncertain to estimate accurately using information markets, the government might consider launching a research initiative to improve our understanding of this relationship. To measure the effectiveness of this proposed research, the government could establish markets for claims on the expected sea level with regulation of greenhouse gas emissions and the expected sea level without regulation. Consider the price difference between these two markets: This quantity measures the expected value of monetized sea level reductions from regulating emissions. The government normally would check whether this price difference exceeds cost to determine its optimal policy. In this case, however, the government would be concerned that this price difference and its optimal policy might change after-market participants learn more about the effects of greenhouse gas emissions on climate change.

105. Under the assumption that realized benefits are normally distributed, the price of this single option (along with the price of the benefits claims) will allow the government to infer the exact distribution of realized benefits. Even without this parametric assumption, the government can establish markets for multiple options with different strike prices (for example, point increases above and beyond 4 points) to estimate the distribution of realized benefits.

106. Fischer Black & Myron Scholes, *The Pricing of Options and Corporate Liabilities*, 81 J. POL. ECON. 637, 640–45 (1973).

107. Indeed, it is possible that information markets may reveal that more information is needed on many regulatory issues before the government can make an informed judgment on the net benefits of different policies. For a classic treatment of the value-of-information issue, see HOWARD RAIFFA, *DECISION ANALYSIS: INTRODUCTORY LECTURES ON CHOICES UNDER UNCERTAINTY* 157–61 (1968).

To assess the effect of the research initiative on this price difference and its optimal policy, the government would need to issue two more contracts. Both contracts would offer their owners the option at the projected completion time of the research initiative to receive a payment equal to the change in the price difference.¹⁰⁸ One contract would be issued if the research initiative is implemented, while the other would be issued if the research initiative is not implemented. The prices of these option contracts reflect market participants' beliefs about the new information that would be discovered with and without the research initiative. So, by comparing the prices of the options, the government can evaluate whether doing the research is likely to reduce uncertainty about the benefits of the policy, and by how much. Then it can proceed with steps three and four to choose whether to implement the research.

4. Market Design Issues

There are many important issues in market design.¹⁰⁹ Here we consider three: ensuring the market is liquid, ensuring market prices are not biased measures of policy benefits, and limiting market power in project implementation. This discussion is by no means exhaustive.¹¹⁰

Liquidity refers to the ability to execute large volume transactions in a market with minimal impact on the prevailing price. It is important for a number of reasons. First, liquidity is needed to get a reasonable price signal on the "true" underlying value of a quantity, because individual transaction prices in illiquid information markets may not be representative of market participants' beliefs. Second, liquidity allows well informed market participants to conduct profitable trades using their own information, providing incentives for traders to acquire

108. Formally, these securities are when-issued at-the-money European call options on the price difference, which expire at the projected completion time of the research initiative.

109. See, e.g., Hanson, *supra* note 7, at 15–25 (discussing thirty engineering design issues for information markets).

110. See, e.g., John O. Ledyard, *Designing Information Markets for Policy Analysis*, in *INFORMATION MARKETS: A NEW WAY OF MAKING DECISIONS IN THE PRIVATE AND PUBLIC SECTOR* (Robert W. Hahn & Paul C. Tetlock eds., forthcoming 2005) (focusing on several problems, including the potentially large number of markets that may be needed to inform policy decisions and the problems associated with getting an unbiased estimate of a policy's impact).

better information. Third, liquidity is important for information markets to serve their hedging function. Even though hedgers are long-term traders, they will be reluctant to take positions in contracts that they cannot reverse if they need cash at some point in the future.

Liquidity will likely be underprovided in fledgling information markets for several theoretical and practical reasons. First, liquidity has positive networking externalities. Each trader in a given market benefits from the addition of a new trader, because the additional trader offers all existing traders enhanced trading options.¹¹¹ Second, market liquidity in information markets allows for improved estimates of prices, implying that liquidity provides improved information to any economic agent that observes the prevailing market price. Because traders who supply liquidity cannot capture the returns from supplying valuable price information to other decision makers, traders will underprovide liquidity.

Liquidity may also be underprovided if traders do not expect the conditional information market contract trades to ever take place. So, for example, it might be hard to get bids and offers on a contract that was contingent on education reform if no trader expected the government to contract for an education reform project. Because traders can only profit when their conditional trades occur, they must believe that the contingency required for valid trades has a reasonable probability of occurring. Otherwise traders will not acquire information or trade.

One consequence of illiquidity is that individual transaction prices may not be representative of the market participants' beliefs. This can happen if the limit order book is so "shallow" that no trader has an incentive to trade on his superior information.¹¹² In this case, a great deal of information may have been revealed since the last transaction price, suggesting the last price is a poor measure of market participants' current beliefs. Also, when transactions do take place, they often occur at the bid price for sell orders and at the ask price for buy or-

111. See Nicholas Economides & Aloysius Siow, *The Division of Markets is Limited by the Extent of Liquidity (Spatial Competition with Externalities)*, 78 AM. ECON. REV. 108 (1988) (modeling market formation with positive liquidity externalities).

112. A shallow order book describes a situation in which there are only bids and offers in the order book for small quantities of a security and at prices that are not close to the quote midpoint.

ders.¹¹³ This leads to observed transaction prices fluctuating predictably back and forth between the two levels as buy and sell orders arrive at the market. In this case, neither the last bid nor the last ask price is necessarily a good measure of current beliefs. This phenomenon, known as bid-ask bounce, is particularly pronounced in situations where the bid and ask prices are far apart, such as in illiquid markets. In these markets, using the last transaction price as a measure of benefits is problematic.

Fortunately, bid-ask bounce can be mitigated through two simple means. First, most theoretical and empirical market microstructure models show that bid-ask spreads are increasing functions of asset prices.¹¹⁴ By designing contracts with returns that are very sensitive to variations in the measurement of net benefits, it is possible to minimize price noise, though at the expense of increasing measurement noise. Second, using average transaction prices or midpoints of bid and ask prices will mitigate bid-ask price noise. Because price noise is by definition temporary, this technique will reduce price measurement error. Averaging over transactions from multiple traders also blunts the incentive for an individual trader to manipulate prices.

A more general solution for liquidity-induced problems might ensure that information contract prices reflect the best available information on benefits, rather than manipulation or temporary liquidity shocks. Toward this end, the government could subsidize market-making activity either directly or indirectly.¹¹⁵ A market maker could post continuous bids and offers for the purchase and sale of each contract. The government could act as the market maker or it could solicit bids from pri-

113. See Richard Roll, *A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market*, 39 J. FIN. 1127, 1128 (1984).

114. See, e.g., Thomas E. Copeland & Dan Galai, *Information Effects on the Bid-Ask Spread*, 38 J. FIN. 1457, 1468; see also Lawrence R. Glosten, *Components of the Bid-Ask Spread and the Statistical Properties of Transaction Prices*, 42 J. FIN. 1293 (1987); Lawrence R. Glosten & Paul R. Milgrom, *Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders*, 14 J. FIN. ECON. 71 (1985).

115. Government subsidization of these markets could give incentives to lobby for more information markets that are subsidized, but it should not provide an incentive to manipulate the outcomes of particular markets. Only in the case that particular market outcomes were tied to the overall subsidies would this present a potential problem.

vate parties in an effort to minimize perceived conflicts of interest and the cost of providing this service.

The government also could try other methods of inducing price discovery. For example, it could offer payments to traders who post bids and offers at or inside the market quotes based on the duration of time that their orders remain inside the market quotes and the extent to which their orders improve upon the market quotes.¹¹⁶ In theory, this type of payment should only encourage liquidity provision and should not alter traders' incentives to take positions in the contracts. This mechanism could be used more generally to establish accurate prices for contracts that are not heavily traded.¹¹⁷

A second market design issue concerns the possibility that information market prices could be biased measures of policy benefits. Abramowicz specifically addresses biases in conditional information markets for policies, such as those proposed in this article.¹¹⁸ He considers the situation in which the policy maker has information unavailable to the market. This implies that the policy maker's decision to implement a policy signals to the market that the policy's net benefits are high, thereby inflating the market's estimate of the net benefits conditioned on the policy being implemented.

For the vast majority of policy issues, this potential selection bias can be solved by forcing the decision maker to disclose all private information or allowing the decision maker to trade with advance notices similar to those required by the SEC in U.S. equity markets.¹¹⁹ In practice, however, it is unlikely that

116. A bid at or inside the market refers to a bid at a price greater than or equal to the highest bid price in the order book.

117. There are other possible mechanisms to encourage liquidity. See Michael Abramowicz, *The Law-and-Markets Movement*, 49 AMER. U. L. REV. 327 (1999). Specifically there are the "market scoring rules" discussed in Robin Hanson, *Combinatorial Information Market Design*, 5 INFO. SYS. FRONTIERS 107 (2003), and the periodic call market idea outlined in Nicholas Economides & Robert A. Schwartz, *Electronic Call Market Trading*, 21 J. PORTFOLIO MGMT. 10 (1995).

118. Abramowicz even suggests that information markets will not necessarily perform better than traditional cost-benefit analysis. See Abramowicz, *supra* note 7, at 951 ("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. . . . 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.").

119. E.g., Form 3: Initial Statement of Beneficial Ownership of Securities, available at <http://www.sec.gov/about/forms/form3.pdf>; Form 4: Statement of Changes in Beneficial Ownership, available at <http://www.sec.gov/about/forms/form4.pdf>;

the decision maker has more information than the market, rendering this a moot point. Thus, properly designed information markets can overcome potential biases in prices caused by policy makers' private information.

Prices could also be biased measures of policy benefits if information markets fail to aggregate information properly. There is a growing literature that describes the situations in which the prices of asset markets do not accurately reflect publicly available information.¹²⁰ In general, the presence of irrational and unpredictable traders can cause inefficiencies in prices when there are limits on the extent of trading by rational traders.¹²¹

Of course, any process designed to aggregate information will perform worse when it incorporates the beliefs of irrational agents. The real question is whether information markets are more robust to the participation of irrational agents than other mechanisms. Judging by the experimental and field evidence cited above, the answer appears to be an overwhelming "yes." Empirically, market prices are better forecasters than not only the average belief in a population, but also the beliefs of only the experts.

There are also strong theoretical reasons to expect markets to perform better than other mechanisms for aggregating beliefs. First, if irrationality is a result of lazy thinking, then monetary rewards will encourage thoughtfulness. Conversely, the threat of monetary punishment could discourage the participation of irrational agents. Second, markets weigh most heavily the beliefs of the rational agents with the most precise information. These agents place the largest trades, so they have the biggest influence on market prices.

Another potential source of bias can arise because of the framework itself. The basic problem is that traders with information about benefits only profit when the trades in the infor-

Form 5: Annual Statement of Changes in Beneficial Ownership of Securities, available at <http://www.sec.gov/about/forms/form5.pdf>; Form 144: Notice of Proposed Sale of Securities Pursuant to Rule 144 Under the Securities Act of 1933, available at <http://www.sec.gov/about/forms/form144.pdf>.

120. See, e.g., ANDREI SHLEIFER, *INEFFICIENT MARKETS: AN INTRODUCTION TO BEHAVIORAL FINANCE* (2000); see also *ADVANCES IN BEHAVIORAL FINANCE* (Richard H. Thaler ed., 1993).

121. See J. Bradford DeLong et al., *Noise Trader Risk in Financial Markets*, 98 J. POL. ECON. 703 (1990).

mation market occur. So they have an incentive to trade not only to exploit their information but also to influence the government's decision to auction off the benefits of a policy so that trades take place. This incentive to influence the government's decision is strongest when the price of the benefits contract with no policy change corresponds to a level of net benefits that is close to the government's minimum acceptable price in the auction.

Suppose, for example, that the current price of the benefits contract with no policy change is too high and informed traders know this. The problem is that informed traders will limit their sales of the overpriced benefits contracts to increase the chances that the minimum acceptable price will not be met in the auction. They will pursue this strategy because they prefer that the contract with no policy change remains in force so they can profit from this information. This could result in a possible upward bias in the information market price for that contract.

The general problem is that if traders anticipate the government's use of the market prices, they will recognize that only programs with higher benefits will be implemented. This anticipation will alter their willingness to buy information market contracts based on policy benefits, potentially biasing the resulting market prices and the government's decision based on those prices.

In subsequent work, we propose a mechanism that deals with this concern by separating the information collection and decision tasks. Specifically, if the information market contracts do not depend on the government's decision rule, then their market prices will be unbiased measures of benefits that can be used by the government. The idea is to have the contract settlement depend on the benefits from a random decision, but to commit to using this random-decision rule infrequently. Most of the time, the government can simply implement its preferred policy, using the unbiased market prices as a guide. Some small fraction of the time, however, the government would need to commit to implementing a particular program, even if it turned out not to be in its best interest to do so. In our model, this is essentially the price of getting better information to improve decisions.¹²² In general, decision makers will need to

122. To see the problem caused by biased information and how it can be solved, consider the education reform example above. The problem with the mechanism is that the reform program will only be implemented if an organization is willing

weigh the cost of getting better information against the benefits of that information.¹²³

There are other ways to address the problem of biased information. The problem is mitigated, for example, when each individual trader has a small amount of wealth invested in the market relative to the size of the overall market. When no individual trader can influence the price, then all traders simply choose their demands without regard for strategic considerations. The problem is also mitigated when information about the policy is widely dispersed across traders. If no individual trader has enough information to be confident of the true pol-

to bid on the project. This will not happen if, for example, no organization anticipates being able to implement the project at a cost less than the expected benefits, which were estimated at \$10 million by the market. If traders know that no firm will implement the project, they will not bother trading the information market contract because its payoff depends on project benefits that will never be realized.

Traders will only trade the contract linked to project benefits if they expect the project to be implemented. Unfortunately, even if the contract is traded, the equilibrium price of the contract may be biased. For example, if a trader has information that the project is overvalued at the current price, he may hesitate to bet against the project to the extent that he otherwise might, for fear that the price of the contract would decline to the point that the project is not chosen. In that case, his contract would be worthless. Similar concerns about price bias apply to the information market for test scores that would prevail without education reform.

To avoid this bias, a decision maker can commit to implementing the education reform program with at least some positive probability, say 5%, *regardless* of what information is revealed by the market price. Because a trader's payoff no longer depends on the decision maker's choice, the bias described above will be eliminated. Most of the time (95%) the decision maker can still choose whether or not to implement the education reform program. The disadvantage is that the decision maker must sometimes choose the reform program when this choice will not yield the highest expected net benefits. This situation would arise no more than 5% of the time, and it may be a price that the decision maker is willing to pay for the ability to implement an informed and unbiased decision at least 95% of the time.

123. For an alternative approach to this problem, see Justin Wolfers & Eric Zitzewitz, *Five Open Questions About Prediction Markets*, in *INFORMATION MARKETS: A NEW WAY OF MAKING DECISIONS IN THE PUBLIC AND PRIVATE SECTORS* (Robert W. Hahn & Paul C. Tetlock eds., forthcoming 2005) (manuscript at 14–21, available at <http://faculty-gsb.stanford.edu/zitzewitz/Research/Five%20Questions.pdf>). These authors suggest introducing an “instrumental” variable to identify the causal effect of a policy choice on the price of an information market contract. This approach is often used in econometric analysis to separate correlation and causation. A key problem, as the authors note, is that it is difficult to find good instruments that would actually get rid of the bias. Our approach does not require a naturally occurring instrument.

icy benefits, then no trader can know whether the policy's implementation will occur because of favorable information about true benefits or because of noise trading. In a similar vein, to limit further the potential for bias, the government could induce uncertainty in its implementation decision by not revealing its minimum acceptable price in advance.

A more general concern is that information markets are susceptible to price manipulation by those with a vested interest in the policy decision.¹²⁴ This need not be the case, however. Hanson and Oprea show that one of the most common forms of price manipulation can actually enhance the accuracy of real-money information markets.¹²⁵ Price manipulation is most likely to occur in information markets when a certain interest group prefers that a particular policy is or is not implemented.¹²⁶ This opens the possibility that members of this interest group will try to manipulate the prices in information markets as a form of political rent seeking. These members could conduct trades to push prices in a particular direction, regardless of the true underlying value of the asset. When other traders expect prices to deviate from fundamental value because of manipulation, however, they have a strong incentive to acquire costly information.¹²⁷ These informed traders will assume large positions in order to profit from the value-independent trades of manipulators.¹²⁸

124. See, e.g., Abramowicz, *supra* note 7, at 937 ("Information markets provide objective predictions, though only if concerns such as the possibility of market manipulation can be overcome.").

125. See Hanson & Oprea, *supra* note 35.

126. Note that manipulation by interest groups is distinct from the above discussion of manipulation by profit-seeking traders who do not inherently prefer a particular policy. Groups with vested interests are willing to sacrifice profits to influence the implementation decision, whereas profit-seeking traders will only attempt to influence the implementation decision to maximize the expected trading profits they obtain from their private information.

127. See Hanson & Oprea, *supra* note 35.

128. Although it does capture some important cases, this argument does not apply universally to efforts to manipulate prices. The proposed market for terrorism, see *supra* note 22, is one such exception: Terrorists have the ability to manipulate information market prices in other ways because they have both insider information and an ability to affect the underlying outcome of this type of information market contract. In addition, the government may not wish to disclose the kind of sensitive intelligence that an information market would reveal. Thus, more research on information markets is necessary before establishing such markets in areas of national security and intelligence.

The interaction between informed traders and price manipulators has many consequences.¹²⁹ One effect is a large transfer of wealth from those trying to manipulate prices to those gathering information.¹³⁰ As a byproduct of this activity, more information will be reflected in the information market price and this result might discourage attempts to use information markets for rent-seeking purposes.

Consider the U.S. government's goal of deregulating local telecommunications services. Policy makers want to provide appropriate incentives to local incumbents to price their local networks efficiently. To realize this goal, it is necessary to measure accurately the riskiness of the telecommunications business. Companies will generally require a higher price to engage in riskier business activities. Not surprisingly, local incumbents tend to represent their business as risky to obtain a higher price for their resources, whereas prospective competitors represent the incumbents' business as riskless to buy the resources cheaply from the incumbents. Both sides of this debate spend substantial sums of money on legal and consulting fees in an effort to convince policy makers of their position.

Suppose, hypothetically, that the local incumbents are able to spend more money convincing policy makers that their business is risky, and therefore obtain a higher price for their networks. How would the existence of a liquid information market for claims on verifiable business risk alter this balance of power?¹³¹ It might seem that local incumbents would also

129. Informed traders may also have an interest in a policy decision and may therefore wish to manipulate the information market price when that price affects the decision. For a formal treatment of this issue, see Robert W. Hahn, Donald Lien & Paul C. Tetlock, *Designing Information Markets for Decision Making* (November 2005) (unpublished manuscript, on file with authors). The paper suggests how to design efficient information markets when the decision maker uses the price in the information market for a specific decision. The paper also shows that the mere act of linking the decision to the market price will typically enhance liquidity in the information market.

130. Price manipulation only refers to cases where manipulators have no control over the underlying contractual outcome. So, for example, this definition does not apply to terrorists participating in terrorist futures markets, but does apply to private citizens predicting the success of education reform.

131. The most important component of the cost of capital is the cost of equity, which is related to the riskiness of equity. The observed "beta" of the incumbent's stock is a reasonable proxy for this risk. Beta is a measure of how much the stock return varies in tandem with the market return; according to one prominent theory of asset pricing, beta is proportional to expected returns. See William F.

spend more in the information market, pushing the price of the security in a direction to suggest their business risk is high. However, if the security becomes priced too highly, speculators will sell the security to profit from the mispricing. By pushing the price back to its fundamental value and correcting the measure of business risk, these speculators will effectively act on the behalf of the prospective telecommunications competitors. In this way, information markets can limit the potential for rent seeking.

A third issue relates to specifying the optimal number of firms for implementing a project. It could be undesirable to have only one firm receive project benefits payments, but beyond the hold-up problem addressed in the project governance section, it is not clear what kind of market power this “monopolist” would wield, as the firm is not given an exclusive right to do anything other than receive contingent benefits payments at a price already determined, in advance, at the social value of the good provided.

One possibility is that the owner of the benefit payments could exercise monopsony power in some input markets. For example, if the project manager is the sole buyer of certain pollution control devices, this could lead to a sub-optimal resource allocation. If a firm’s monopsony power depressed input market prices below the socially optimal level, then it would induce reductions in the production of inputs and market exit from suppliers of inputs. Monopsony is probably not a viable long-term position, however, because supply resources rarely remain specialized in the long term. Moreover, to maintain a steady supply of inputs, a forward-looking monopsony buyer will not wield its monopsony power even in the short run.¹³²

So far we have shown that it is possible to combine information markets with pay-for-performance contracts to improve decisions, at least in theory. At the same time, we have shown that information markets may not be ideal in some settings. These include when disclosure of information is costly, as it

Sharpe, *Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk*, 19 J. FIN. 425 (1964). All that is needed to calculate beta is time-series data on a firm’s stock price and the price of the market index. Both can be readily verified using monthly data from the Center for Research in Securities Prices.

132. For a discussion of monopsony, see DENNIS W. CARLTON & JEFFREY M. PERLOFF, *MODERN INDUSTRIAL ORGANIZATION* 107–10 (4th ed. 2005).

may be in the case of national defense;¹³³ when information markets may underperform experts, which appears to be unlikely in most settings;¹³⁴ when it is very costly to get an information market to open; and when there is no verifiable measure of performance.

IV. A BENEFIT-COST ANALYSIS OF INFORMATION MARKETS

In this section we provide a general description of the benefits and costs of information markets. We also compare the direct approach and indirect approaches to performing cost-benefit analysis using information markets.

At the outset, we consider the first-order question of whether such markets could, in principle, address a market failure.¹³⁵ From a social standpoint, there could be a failure to provide adequate information, because information is a classic example of a public good. The government is frequently trying to develop policies on the basis of limited information that is strategically supplied by interested parties.¹³⁶ There is no reason to think that the provision of such information is optimal. Indeed, there is little reason for the private sector to provide adequate information in government decisions, even when required by law to supply different kinds of information.¹³⁷

As argued in Part III, information markets can help address some of these information gaps. We now turn to a qualitative

133. Some information aggregation can be achieved without disclosing sensitive information if an organization adopts an internal information market. Corporations, such as Hewlett-Packard and Eli Lilly, have already begun to adopt internal information markets for obtaining information that could confer advantages over their competitors.

134. For example, if only one individual has a key piece of information, then an information market would not necessarily help if that person were known beforehand.

135. See, e.g., Francis M. Bator, *The Anatomy of Market Failure*, 72 Q.J. ECON. 351 (1958).

136. For example, automobile companies may lobby for pollution controls on oil companies and vice versa. See, e.g., Bruce Yandle, *Bootleggers and Baptists: The Education of a Regulatory Economist*, REGULATION, May-June 1983, at 12, 13.

137. A key problem is that firms neither want to spend the resources nor risk giving valuable information to competitors. See, e.g., BRUCE M. OWEN & RONALD BRAEUTIGAM, *THE REGULATION GAME: STRATEGIC USE OF THE ADMINISTRATIVE PROCESS* 4 (1978).

description of some of the benefits and costs of using these markets to inform and implement public policies.

There are three broad categories of benefits: a more informed assessment of policy proposals; greater transparency and accountability in the decision-making process; and greater availability of assets for efficient financing and risk sharing, which encourages the implementation of projects with positive net benefits.

A. *A More Informed Assessment of Policy Proposals*

One of the primary advantages of information markets is that they have the potential, if designed properly, to provide a much more informed assessment of policy proposals.¹³⁸ As noted above, they can be very helpful in determining net benefits, distributional implications, and the value of additional research in assessing net benefits. In this section, we explore some other benefits, including the potential for examining new policies, the role that markets can play as assimilators of information, and the benefits of real-time information.

Greater use of these markets could give rise to such proposals as well as a more informed assessment of new proposals. As better information is revealed about the potential impacts of policies, it could encourage policy makers to experiment and fine tune the menu of policies being considered. In addition, because the marginal cost of considering some new policies using information markets is likely to be modest, academics and private citizens may decide to advance proposals of their own and use information markets to estimate their effects.¹³⁹

Similarly, interest groups could advance and analyze their preferred proposals using this approach. Disputes about the likely impact of different policies, such as a gasoline tax, a minimum wage, or a free trade area, could be more easily resolved with data from information markets.¹⁴⁰

138. For a similar argument in support of information markets, see Hanson, *supra* note 7, at 1–9. See also Abramowicz, *supra* note 7, at 962–97.

139. See *infra* Part IV.E for a discussion of costs.

140. Disputes surrounding the effectiveness of policy are quite common when the economic stakes are large, with groups taking different sides on the impact of all sorts of proposed legislation and regulation. Some of these conflicts could be addressed by using information markets to measure the impact of various proposals on key variables, such as output and employment. The information from these markets could also be used more frequently by the media as a way to understand and resolve differences of opinion. It would provide a low-cost way

Still another benefit of information markets relates to assessing the effectiveness of different policies in something that approximates real time. Such information could be a potentially large benefit for policy makers who must make time-sensitive decisions. In some cases, agencies have a relatively short period of time to evaluate policy. Performing detailed analysis can often take months, while market prices can be generated from liquid information markets relatively quickly. For example, stock prices and derivative prices can reach equilibrium within minutes or even seconds.¹⁴¹

An additional advantage of nearly real-time information is that a comparison of policies can be made across time as key uncertainties are resolved. A major issue in cost-benefit analysis is the extent to which *ex ante* studies of a particular policy agree with *ex post* studies. With information markets, the policy maker could choose to examine the expected value of key variables at different points in time as the policy evolves and is implemented.

B. Greater Transparency and Accountability in Decision Making

One of the advantages of cost-benefit analysis is that it has the potential to make the analytical basis for public policy decisions more transparent.¹⁴² The same argument can be made for information markets. By supplying information on key parameters that are involved in decision making, information markets make the public policy process more transparent.¹⁴³

for journalists, who often operate under severe time constraints, to see how different groups are trying to shape issues.

141. See Joel Hasbrouck, *Intraday Price Formation in U.S. Equity Index Markets*, 58 J. FIN. 2375, 2387–88 (2003) (estimating that the prices of S&P 500 equity index futures respond to new information within roughly sixty seconds).

142. See, e.g., ARROW ET AL., *supra* note 3, at 1–4; see also ROBERT W. HAHN & ROBERT E. LITAN, *IMPROVING REGULATORY ACCOUNTABILITY* (1997), available at <http://aei-brookings.org/admin/pdffiles/phpf4.pdf>; SUNSTEIN, *supra* note 2, at 106–08.

143. Indeed, if the government wished to go a step further, it could either directly or indirectly help make information market prices, volume, and order books available at very low cost or no cost. However, this information is probably most efficiently conveyed by the exchanges listing information market securities, which need not be subsidized by the government. These exchanges have obvious incentives to convey some, if not all, relevant information about buying and selling activity in their securities. See discussion *infra* Part IV.E.

Greater transparency is likely to increase accountability for lawmakers and regulators. Making the data more readily available means that the cost of monitoring regulators and lawmakers decreases. Interested parties, then, will find it easier to hold them accountable for decisions. Abramowicz makes a similar point, arguing that information markets have the potential to help discipline the agency decision-making process.¹⁴⁴ Indeed, even if the government does decide to pursue a more narrowly self-interested agenda, as public choice theory might predict,¹⁴⁵ the introduction of such markets could place constraints on what the government can do.¹⁴⁶

Greater transparency could also reduce the scope for political manipulation. In traditional cost-benefit analyses, it is easy to imagine cases where the analyst may be biased or captured by a particular constituency.¹⁴⁷ Information markets, by improving the quality of information along with access to that information, reduce the potential for certain kinds of manipulation. It is possible, though by no means certain, that the level of rent-seeking would decrease if information markets for policy were used widely. This is because it would become more costly to change policy outcomes if something like our new approach were implemented.¹⁴⁸

Consider an example: Suppose the government is considering a law that would regulate the prices of pharmaceuticals. Economists frequently argue that price controls would reduce incentives to develop new drugs and, therefore, reduce the number of new drugs that are introduced into the market. At the same time, others argue that such controls would make little difference. Well designed information markets could shed

144. See Abramowicz, *supra* note 7, at 934 (“Information markets could help constrain administrative decisionmaking and limit ideological decisionmaking.”).

145. See, e.g., JAMES M. BUCHANAN & GORDON TULLOCK, *THE CALCULUS OF CONSENT: LOGICAL FOUNDATIONS OF CONSTITUTIONAL DEMOCRACY* (1962).

146. In what follows, we generally presume that the government is interested in maximizing some measure of social welfare as it is typically operationalized. Even if it is trying to achieve some other goal, information markets could be helpful. We abstract from the general difficulties of defining social welfare. The seminal work here is KENNETH J. ARROW, *SOCIAL CHOICE AND INDIVIDUAL VALUES* (1951).

147. See Abramowicz, *supra* note 7, at 970 (discussing how information markets could reduce the potential for interest group manipulation).

148. We recognize that rent seeking is difficult to measure. The general idea is simply to try to reduce the returns on rent seeking over a wide range of activities. See, e.g., Anne O. Krueger, *The Political Economy of the Rent-Seeking Society*, 64 *AM. ECON. REV.* 291 (1974).

light on this issue by estimating the rate of new drug approval under the two policies. If the public and the media came to see such markets as a trustworthy source of information, these markets could have a greater influence on policy, and the influence of interest groups on both sides of the issue could decrease.¹⁴⁹

Information markets could also improve decision making by increasing the level of expected net benefits from policy choices. Regulators, lawmakers, and the courts might find it in their interest to make greater use of data on benefits and costs if those data were more accurate. These policy makers would then become more accountable in the sense that they would be more inclined to use data on costs and benefits to justify their decisions. In addition, lawmakers and regulators may implement procedural changes that increase the likelihood that data on costs and benefits will be considered seriously.¹⁵⁰

Finally, the introduction of information markets could lead to a change in bureaucratic focus. Fewer resources may be allocated to traditional cost-benefit analyses and more to the construction of information markets to inform the process. Such a reallocation would not necessarily lead to a reduction in overall expenditures. That reduction depends on the value of the information that is produced from the information markets and the extent to which it compliments or substitutes for existing information. One could even imagine a case in which net time and effort spent on policy analysis actually increases, if time and effort spent analyzing data from information markets increases at a greater rate than time and effort spent on conventional cost-benefit analysis declines.

C. *Greater Availability of Assets for Financing Projects and Spreading Risks*

The creation of information markets for claims on the benefits from policy decisions has the potential to facilitate the im-

149. A current problem is that it is difficult for the public to know what may have occurred in situations that are not actually observed. Sam Peltzman argues this is one of the reasons that inefficient regulation persists. See SAM PELTZMAN, *REGULATION AND THE NATURAL PROGRESS OF OPULENCE* (2004). Information markets provide a solution to this problem.

150. Examples include a regulatory budget and regulatory net benefit accounts. See discussion *infra* Part V.B.

plementation of socially valuable projects. By aggregating information from diverse sources, the prices in these markets reveal the social value of policy proposals. Using these prices, the government can then monetize this value and package it in the form of a security. This security would be identical in its pay-offs to the claim on project net benefits, but it would also confer to its owner the right to vote on project implementation decisions. The government could collect the surplus on projects by auctioning off the right to execute the project: in other words, by selling the security. The government's position in such a transaction is analogous to that of an entrepreneur auctioning his project to venture capitalists.¹⁵¹

In addition, as discussed above, it is also possible to design contracts that enable private sector participants to reduce or eliminate much of the risk associated with public policies. This applies to the firms actually implementing the policies, but it also extends to consumers and workers. Those parties affected by economic policies have an opportunity to mitigate the uncertain impact of the policy by purchasing information market contracts. For example, consumers concerned about price increases caused by a policy could purchase contracts that yield payments when the policy results in price increases. Similarly, workers concerned about job losses resulting from a policy could purchase contracts that yield payments when the policy creates more unemployment.

Finally, in addition to serving as a risk management tool for firms, information markets can help firms raise cash when they are constrained. While some firms have ready access to capital markets, others may not. The information market can serve as a source of project funding for both types of firms.

D. Cost of Information Markets

Several costs are related to information markets. These costs can be minimized by setting up an electronic exchange for the purchase and sale of contracts, similar to the electronic communication networks that facilitate trades in U.S. stock mar-

151. Robert Shiller identifies several markets that could be used to help increase efficiency and equity. These include markets for claims on individual and national income and human capital. While Shiller's proposals are not strictly information markets as we have defined them, information markets have the potential to improve risk management substantially. See ROBERT J. SHILLER, *MACRO MARKETS: CREATING INSTITUTIONS FOR MANAGING SOCIETY'S LARGEST ECONOMIC RISKS* (1993).

kets. The fixed costs of establishing this electronic exchange will include payments to programmers who develop the web exchange application and hardware costs for servers and storage devices necessary for the daily operation of the exchange.

The operating costs would, at a minimum, include monitoring, maintenance, and upkeep costs for addressing ongoing exchange issues. Legal procedures and controls are needed to maintain a reputation for the market so traders perceive that they are treated fairly. Operating costs could also include direct or indirect capital outlays that ensure sufficient liquidity in exchange contracts to retain traders' interest. Outlays could take the form of commission rebates or negative commission on trades, order flow, or some other variable of interest.

These operating costs may be offset somewhat by exchange revenues, which include commissions charged by the exchange. If the exchange also operates as a market maker, its revenue would be dependent upon the difference between the prevailing bid and ask prices on each contract. In this case, the exchange (or any market maker) also incurs the risk that the value of its contract inventory (net contract position) declines. Through careful risk management and research into exchange trading behavior, this risk can be kept to a minimum. Alternatively, the exchange could hire a market maker to share or completely eliminate this risk.

It would be desirable for an exchange to set up a mechanism, such as a public Internet forum, for proposing new contracts. The direct cost of listing new securities on the exchange for trading would be trivial, but indirect costs could be large. Adding too many new securities too quickly could detract from the liquidity in existing markets, thereby undermining the purpose of the exchange: facilitating the meeting of buyers and sellers by coordinating their demands and supplies in time and space. On the other hand, if adding a new security encourages new trading volume, meeting this volume could require the exchange to upgrade its hardware and increase the scale of its operations. Exchanges will need to balance these considerations in choosing which securities to approve for listing.

E. *Comparison of Direct and Indirect Approaches
for Information Markets*

This paper has focused on a direct approach to estimating net benefits using information markets. An alternative to this direct approach is an indirect approach, where the market predicts what an expert will say about net project benefits at some time in the future. In other words, the market predicts the conclusion of an expert's cost-benefit analysis.¹⁵²

Here, we will briefly compare the two approaches. A summary of these approaches is provided in Table 1. The table reveals that both approaches can estimate net social benefits and also address equity concerns. In general, however, the direct approach is preferable because it can perform just as well as the indirect approach. Furthermore, it has the advantage that it avoids introducing additional uncertainty related to the selection of an expert.

The reason the direct approach possesses most, if not all, of the capabilities of the indirect approach is because we allow the government to change its policy. The government is allowed to buy back contracts and issue new ones in light of new information or improved measures of benefits. Thus, it can redefine the benefit function, and hence, the policy, over time as circumstances warrant.

The flexibility of the decision rule when using an indirect approach has numerous disadvantages. Because there are few, if any, restrictions on what the analyst may include in a cost-benefit analysis, the price of an information market contract is extremely difficult to interpret. It is impossible to know what sort of cost-benefit analysis market participants anticipate will be used in the future. In other words, the information market's data will compound both the expected benefits and the expected way in which the analyst will conduct the benefits determination, producing data that is potentially less valuable.

Abramowicz argues that if the selection of the analyst occurs sufficiently far in the future, the market beliefs about the analyst's future cost-benefit analysis will yield an unbiased measure of the future expert's estimates of net benefits.¹⁵³ This may

152. See generally, Abramowicz, *supra* note 7 (discussing the advantages of using information markets to forecast the retrospective cost-benefit analysis conclusions of such a future expert).

153. See *id.*, at 939, 991-92.

be true, but it overlooks some potential problems with the expert's estimate.

Although this procedure potentially overcomes expected political biases and policy preferences, it does not preclude biases introduced by the market's anticipation of the analyst's information access, information gathering, and information processing abilities.¹⁵⁴ For example, the market may believe that the future analyst will overlook certain types of economic benefits, perhaps because typical analysts performing cost-benefit analyses have tended to overlook these types of benefits in the past. More generally, there is no way to know whether the market expects the future analyst to omit certain types of benefits and costs that result from the policy. Essentially, the market's attempt to forecast the subjective decisions of the analyst, rather than objective measurements, needlessly introduces "noise" into the information conveyed by the market's prices. This noise reduces the value of the market as a general predictive mechanism.

The analyst is just one of many experts who could have different pieces of information that would inform a judgment of net benefits. The current price will reflect only the public's knowledge of the expected analyst's knowledge. The requirement that prices of information market contracts be based only on one imperfectly informed analyst's future judgment defeats much of the point of having information markets.

Also, the unspecified future analyst will not have a clear financial incentive to gather and process information accurately in her determination of net benefits.¹⁵⁵ In contrast, participants in information markets have strong monetary reasons to be thoughtful, resourceful, and meticulous in estimating future net benefits.

Finally, decision making based on the anticipation of an unspecified analyst doing an unspecified cost-benefit analysis is extremely opaque and, therefore, unlikely to inspire public confidence and trust in the decision-making process. The direct method, by comparison, provides a transparent way of valuing

154. This is true even for the case in which the analyst herself decides what should count as a benefit or cost. She is still not omniscient when it comes to measuring those benefits and costs.

155. In correspondence with the authors, Abramowicz notes that professional pride could serve as a partial substitute for financial incentives.

benefits and costs. Transparency is especially important to the success of an information market because traders are reluctant to participate in financial markets they do not trust or understand.

The indirect approach to estimating net benefits is inconsistent with the guiding concept of securities laws and rules. The SEC mission statement explains that “all investors . . . should have access to certain basic facts about an investment prior to buying it.”¹⁵⁶ Designed to restore public confidence in financial markets, the Securities Exchange Acts of 1933 and 1934 stress disclosure of all material risks involved in investing.¹⁵⁷ A transparent explanation of risks is extremely difficult for a security whose value depends on the whims of an unspecified future analyst. By contrast, the risks in information market contracts based on transparent quantitative measures could be estimated using mathematical models.

V. SIGNIFICANCE FOR POLICY DESIGN AND EVALUATION

We have presented a framework for designing contracts that provides information on the costs, benefits, and net benefits of different policy proposals. In this section, we consider some broader issues: first, how this information could be used in the policy process; second, how it could substantially change the nature of legislative and regulatory oversight; and third, possible roles for the government in facilitating and evaluating information markets, and roles for researchers in evaluating such markets.

A. *Information Markets in the Policy Process*

First, the potential scope for information markets in decision making is enormous. In theory, information markets can address many gaps in our knowledge of both public and private

156. See U.S. Securities and Exchange Commission, *The Investor's Advocate: How the SEC Protects Investors and Maintains Market Integrity*, available at <http://www.sec.gov/about/whatwedo.shtml>.

157. The Exchange Act established the SEC in response to the investor disenchantment with securities markets that followed the stock market crash of 1929 and the Great Depression. See *The Securities Act of 1933*, 48 Stat. 74, (codified at 15 U.S.C. §§ 77a–bbb); and *The Securities Exchange Act of 1934*, 48 Stat. 881, (codified as amended at 15 U.S.C. §§ 78a–lll).

decision making.¹⁵⁸ More generally, liquid information markets facilitate the allocation of capital to value-maximizing projects in the public sector, inform resource allocation decisions in the private sector,¹⁵⁹ and enable unique risk-sharing possibilities throughout the economy.¹⁶⁰

The benefits of information markets are not limited to governments; there are a large number of potential applications to international institutions, non-governmental organizations, and foundations. It is quite common for these institutions to fund or design projects. The World Bank, for example, funds projects in developing countries. Information markets can help in the design of specific projects as well as with the ranking of those projects.¹⁶¹

Ventures in the private sector could also benefit from improved access to information and hedging vehicles. A natural extension of Germany's initial public offering (IPO) market would be an IPO market that conditions contracts on a firm's implementation strategy.¹⁶² Hanson has described a similar secondary conditional market in which traders could directly bet on the effect of firing a company's CEO on the stock price of the company.¹⁶³

There are also a number of potential applications of information markets to corporate finance. One of the current problems in corporate finance is delivering accurate information to the

158. See Abramowicz, *supra* note 7, at 962 ("Information markets in theory might be used in any decision making environment, and corporations have experimented to determine how effectively information markets aggregate information. . . . Information markets are best justified [for government] as a means for disciplining government decisionmaking.").

159. For more on applying information markets to both the private and public sectors, see Abramowicz, *supra* note 7; Hanson, *supra* note 7.

160. See SHILLER, *supra* note 151.

161. See, e.g., Robert W. Hahn and Paul C. Tetlock, *Making Development Work*, POL'Y REV., Aug.-Sept. 2005, at 27, 32-34. One of the applications considered in that article in some detail is the Copenhagen Consensus, which aimed to rank approaches for addressing some of the world's most serious problems. Specifically, in 2004, a group of eight distinguished economists assembled in Copenhagen to see if they could achieve some consensus on the best ways of advancing global welfare if an additional \$50 billion were made available to government. Information markets can serve as a valuable source of information in framing this important debate. See *id.*

162. For example, the when-issued contracts in benefits claims could be conditioned on alternative organizational structures or compensation systems.

163. See Hanson, *supra* note 7, at 25.

CEO so that she can decide which projects to fund. Self-interested project managers, desiring implementation of their own projects, do not necessarily have an incentive to supply accurate information. Information markets would be ideally suited to aggregating such information if it resided with several different employees in the company. This approach could reduce problems of rent seeking on the part of some key members in the firm.¹⁶⁴

Athanasoulis et al. describe the importance of markets based on key macroeconomic indices.¹⁶⁵ In addition to conferring the benefits stressed by the authors, these markets would facilitate the spread of accurate economic statistics to agents whose decisions could benefit from this information. In this respect, macro markets are similar to information markets, but they serve another important purpose. As emphasized by Athanasoulis et al., individuals, firms, and even national governments could derive enormous benefit from the formation of liquid markets that allow them to minimize unnecessary fluctuations in their incomes.¹⁶⁶

While there is a huge potential scope for information markets, the test will be in their application. One of the critical issues is how they are designed, researched, and applied in particular contexts.

A critical question is how information from these markets should be used in decision making. Hanson suggests that these markets should be decisive in certain contexts.¹⁶⁷ While we would not rule out making these markets decisive in certain settings, we think it is premature and that his proposal is overly ambitious.

Our general approach differs from Hanson's in a number of ways. First, we would allow decision makers to consider other factors in deciding whether to move forward with a project.

164. Managers and CEOs are sometimes modeled as having an interest in empire building, which is a form of rent seeking. See, e.g., Jeremy C. Stein, *Internal Capital Markets and the Competition for Corporate Resources*, 52 J. FIN. 111 (1997). Information markets have the potential to reduce rent seeking in both private and public decisions: By giving principals more information, these markets expand the set of feasible incentive contracts, enabling improvements in firm efficiency.

165. See Stefano Athanasoulis et al., *Macro Markets and Financial Security*, ECON. POL'Y REV., Apr. 1999, at 21.

166. See *id.*

167. See Hanson, *supra* note 7, at 2 ("The basic rule of futarchy would then be this: when speculative markets clearly estimate that a proposed policy would increase expected national welfare, that policy becomes law.").

Second, we would not focus on GDP as a measure of welfare in most instances. Third, we believe it is important to look for proposals that are likely to maximize net benefits as opposed to implementing those that simply increase net benefits. Fourth, Hanson focuses solely on the role of information markets in improving policy, whereas we highlight the link between information markets and using pay-for-performance contracts to provide incentives for efficient implementation of projects.

More knowledge is needed about how these markets perform in practice in different settings. Furthermore, one needs to weigh the costs and benefits of using a decision rule as opposed to allowing decision makers to exercise judgment. This is a well studied problem in a number of contexts, including the choice of whether to have a monetary rule.¹⁶⁸

In general, these markets will be most useful when they can closely approximate the variables of concern to interested parties and provide good measures of those variables. Thus, they should play a more prominent role when evidence indicates they can provide more useful information.

B. *Solving Some Difficult Government Oversight Problems*

One of the great problems facing lawmakers is how to design a legal and regulatory process that is transparent and also holds lawmakers and regulators accountable for the laws and regulations that they make. In addition to calls for greater transparency, several scholars have suggested mechanisms for gaining a better appreciation of the impacts of regulatory spending and overall spending. For example, Robert Crandall and Christopher DeMuth both propose a regulatory budget that would constrain the costs that an agency could impose on society with its regulations, or more generally with its policies.¹⁶⁹ Such a budget could encourage a focus on projects that have a high social payoff. Eric Posner suggests a similar idea,

168. See, e.g., Finn E. Kydland & Edward C. Prescott, *Rules Rather than Discretion: The Inconsistency of Optimal Plans*, 85 J. POL. ECON. 473 (1977) (discussing the advantages and disadvantages of a monetary rule).

169. See Robert Crandall, *Is Government Regulation Crippling Business?*, SATURDAY REV., Jan. 20, 1979, at 31, 34–40; Christopher C. DeMuth, *Constraining Regulatory Costs—Part II: The Regulatory Budget*, REGULATION, Mar.–Apr. 1980, at 29; see also ROBERT E. LITAN & WILLIAM D. NORDHAUS, REFORMING FEDERAL REGULATION (1983); HAHN & LITAN, *supra* note 142, at 13–19.

but extends the concept to net benefits.¹⁷⁰ In his formulation, each agency could have a net benefit account based on the estimated net benefits of the regulations it promulgated.¹⁷¹

While scholars have advanced these proposals, both Congress and the executive branch have introduced a number of procedures and requirements aimed at improving regulation and getting a better understanding of its impacts.¹⁷² For example, Congress now requires the Office of Management and Budget (OMB) to do an annual report on the costs and benefits of federal regulation.¹⁷³ In its report, OMB must estimate the total annual costs and benefits of federal rules, analyze the impacts of federal regulation on state, local and tribal government, small business, wages, and economic growth, and provide recommendations for reform.¹⁷⁴

In addition, all presidents since Jimmy Carter have implemented some form of regulatory oversight.¹⁷⁵ Beginning with President Reagan, as a result of the regulatory impact analysis requirement in Executive Order 12,291,¹⁷⁶ agencies are providing more information than ever before about the benefits and costs of federal regulation.

170. See Eric A. Posner, *When Reforming Accounting, Don't Forget Regulation*, POL'Y MATTERS, Aug. 2002, available at <http://www.aei-brookings.org/policy/page.php?id=104> (arguing that the institution of Net Benefit Accounts (NBA) will promote more efficient regulations and stating, "The well-run agency could build up a surplus in its NBA, allowing it to take risks with regulatory projects that might initially fail a cost-benefit analysis but, in the end, produce a net benefit for society").

171. See *id.*; see also Robert W. Hahn, *The Economic Analysis of Regulation: A Response to Critics*, 71 U. CHI. L. REV. 1021 (2004).

172. While the examples used here focus on regulation, this is done in the interest of avoiding abstraction. The applications to oversight could be much broader. The only requirement is that the government be able to specify some measure of the output or results of its resource allocation decisions.

173. See Robert W. Hahn & Mary Beth Muething, *The Grand Experiment in Regulatory Reporting*, 55 ADMIN. L. REV. 607 (2004), available at <http://www.aei-brookings.org/publications/abstract.php?pid=314>.

174. See, e.g., OFFICE OF MGMT. & BUDGET, REPORT TO CONGRESS ON THE COSTS AND BENEFITS OF FEDERAL REGULATIONS (1998), available at <http://www.whitehouse.gov/omb/inforeg/costbenefitreport1998.pdf>.

175. President Carter created the Office with the Paperwork Reduction Act of 1980, Pub. L. No. 96-511, 94 Stat. 2812 (codified at 44 U.S.C. §§ 3501-3520) but the Act was not implemented until the end of 1980. Before Reagan introduced the Regulatory Impact Analysis requirement in 1981, data on the benefits and costs of regulation were sparse.

176. Exec. Order No. 12,291, 3 C.F.R. 127 (1981), reprinted in 5 U.S.C. § 601 (2000), available at <http://www.archives.gov/federal-register/codification/executive-order/12291.html>.

The effectiveness of many of these requirements and proposals rests on their ability to gather useful information on the costs and benefits of particular regulations. Take the regulatory budget as an example. The Achilles heel of the regulatory budget has always been the difficulty of getting reasonable estimates of the cost of regulation. Should the government use its own estimates, those of the affected businesses, or some other source? A similar problem plagues net benefit accounts, except information on both costs and benefits is needed. Even if using direct measures of costs and benefits, such as compliance costs and benefits such as lives saved, the problem can be formidable.

Information markets can aid substantially in solving the problems that have plagued the regulatory budget, net benefit accounts, and similar proposals. They do this by presenting reasonably unbiased estimates of the costs and benefits of particular regulations. As noted above, these estimates are not perfect.¹⁷⁷ They are likely, however, to be substantially better than those generated by conventional methods in many settings.¹⁷⁸

Moreover, with the advent of widespread information markets, the nature of the regulatory oversight problem could change. It may no longer be necessary to have a regulatory budget if the expected net benefits of regulations can be readily estimated in a transparent manner that is viewed as legitimate.¹⁷⁹

Net benefit accounts may not be needed for a similar reason. While net benefit accounts give agencies some incentive to promote economic efficiency, a simpler rule could be enacted. For example, if the net benefits of a particular rule did not pass a benefit-cost test according to data from information markets, then the agency could be asked to provide more justification for implementing the regulation.¹⁸⁰

177. See *supra* Part III.D.4.

178. See *supra* Part II.

179. This presumes that such markets would be viewed as legitimate, which likely depends on how they evolve over time. In addition, if agencies still have a natural tendency to regulate in their area, perhaps due to "tunnel vision," a regulatory budget could be a helpful antidote. See BREYER, *supra* note 1, at 11–19 (discussing the problem of tunnel vision in some government agencies).

180. See ROBERT W. CRANDALL ET AL., AN AGENDA FOR FEDERAL REGULATORY REFORM 13 (2003).

In fact, greater use of information markets for assessing the benefits and costs of policies could fundamentally change the nature of the oversight process. Indeed, it could transform the process into something that is much more familiar to legislators: a budgetary process. At the core of contracts to elicit information on benefits and costs is a reward system. In the case of public policy, the government makes payments based on the actual benefits of a proposed policy. That is, the agency or the government is expected to pay out money in exchange for benefits.¹⁸¹

Thus, Congress, or other legislatures, may want to address the problem in a different way. Congress would need to allocate budgets so that agencies could essentially offer market-based contracts that increase expected social welfare or net benefits. Congress could, for example, pass budgets for particular policy areas such as telecommunications regulation, environmental, health and safety regulation, and health care. By imposing a budget constraint, Congress would be limiting the deadweight losses associated with raising revenues, and imposing discipline on agencies similar to a regulatory budget. Agencies could petition Congress in special cases if they could show that net benefits were particularly high for certain projects but they lacked adequate funds to support the projects. Alternatively, agencies with a very good track record of generating net benefits could be rewarded with larger budgets in subsequent budget cycles.¹⁸²

Consider the case of environment, health and safety regulation, addressed in Justice Breyer's *Breaking the Vicious Circle*.¹⁸³ Justice Breyer, among others, has suggested trying to save more lives with a given level of resources by considering reallocating expenditures to activities that would save the most lives.¹⁸⁴ Information markets could easily accommodate this problem. If another objective were viewed as desirable, such as

181. It is not necessary that the money be paid by the government or come from any particular source, but some credible promise to pay needs to be made.

182. This is similar to Eric Posner's suggestions for net benefit accounts. See, e.g., Eric A. Posner, *Controlling Agencies with Cost-Benefit Analysis: A Positive Political Theory Perspective*, 68 U. CHI. L. REV. 1137 (2001), available at http://www.law.uchicago.edu/Lawecon/WkngPprs_101-25/119.EP.Controlling%20Agencies.pdf.

183. See BREYER, *supra* note 1.

184. See *id.*; see also Tammy O. Tengs & John D. Graham, *The Opportunity Costs of Haphazard Social Investments in Life-Saving*, in RISKS, COSTS, AND LIVES SAVED: GETTING BETTER RESULTS FROM REGULATION 167 (Robert W. Hahn ed., 1996).

age-adjusted lives saved, quality-adjusted life years, or willingness to pay, information markets could be reconfigured to predict measurements of that objective under a variety of policy alternatives. Thus, if the value of a statistical life should differ across agencies,¹⁸⁵ contracts could be written to accommodate that characteristic by adjusting the payout depending on the lives expected to be saved.

Congress could apply the same logic to other areas of policy, such as public works projects. There is no reason, save perhaps politics, that Congress could not apply the same rigor to dam building or environmental restoration projects.¹⁸⁶

In summary, we are suggesting that more direct measures of the benefits and costs of different policies can help administrators solve some vexing problems of legislative and regulatory oversight. In addition, these measures may change the nature of the oversight process to something that is more familiar to lawmakers: allocating government expenditures by agency.

C. *The Potential Role for Government and Researchers*

Because the potential benefits from the use of information markets are so large, government needs to examine its potential role in facilitating information markets for promoting better public policy. We suggest that government use an iterative approach in which it initially experiments with information markets on a small scale.¹⁸⁷ It should strive to obtain more information about the value of information markets and the impact of different policy interventions. After learning more about which strategies work through experimentation, the government could revise its strategy accordingly.

To assess the appropriate role for government, we need to know something about the potential net benefits that could

185. See Cass R. Sunstein, *Are Poor People Worth Less Than Rich People?: Disaggregating the Value of Statistical Lives* (AEI-Brookings Joint Center, Working Paper No. 04-05, 2004), available at <http://www.aei-brookings.com/publications/abstract.php?pid=430>.

186. For an early treatment of the politics of such issues, see JOHN A. FERREJOHN, *PORK BARREL POLITICS: RIVERS AND HARBORS LEGISLATION, 1947-1968* (1974).

187. In many cases, it may be difficult to formulate an appropriate benefit function. Doing experiments on a small scale could help uncover problems in the specification of the benefit function, ensuring that policy makers get the benefits they intend. For example, municipalities eager to undertake policy experiments could serve as trial grounds for future statewide and nationwide policies.

flow from information markets. At this point, we know very little in a quantitative sense.¹⁸⁸ We also need information on the potential value of different methods of government intervention.

We offer one possible template for government involvement that consists of five parts. First, government could pursue a “no-regrets” policy in which it attempts to reduce inefficient regulatory barriers to information markets.¹⁸⁹ The current regulatory environment is highly uncertain for information market contracts, which is a barrier in itself. Neither the federal Commodity Futures Trading Commission (CFTC) nor the individual state gambling commissions have an obvious legal claim to regulating information markets. We recommend that the CFTC take decisive action to eliminate this uncertainty by assuming control of the review and approval process for newly introduced information market contracts.¹⁹⁰

Second, the government could facilitate the development of an online exchange that specifically addresses these kinds of issues. The costs for setting up the exchange itself should be relatively modest. Subcontracting the job to an existing exchange (or exchanges) would be desirable to obtain economies of scale and expertise. In principle, this would be similar to the way the U.S. government hires private underwriters, such as Goldman Sachs or Merrill Lynch, to assist in the issuance of treasury securities. The exchange could also handle contract clearing and settlement, adjudicate trade disputes, and verify the events that form the basis of its listed contracts.

188. We would advocate the use of a “meta” information market to determine the expected benefits of the use of information markets. For example, establish several when-issued markets for U.S. GDP, conditional on different amounts of funds allocated to government-sponsored information market research. However, a meta-market might not satisfy those skeptical of the utility of information markets due to the circular nature of the assessment.

189. Currently, the Commodity Futures Trading Commission is considering whether event markets fall under its jurisdiction. If it decides to assert jurisdiction over event markets, even non-profit markets established for research purposes would be forced to adhere to the Commodity Exchange Act. The law would require all event markets to endure a long and costly designation process in order to demonstrate that market transactions are safe and secure. *See* 7 U.S.C. § 6 (2005).

190. For a more detailed discussion arguing that state regulation aimed at Internet gambling is generally not appropriate for regulation of information markets, see Robert W. Hahn & Paul C. Tetlock, *A New Approach for Regulating Information Markets*, J. REG. ECON. (forthcoming), available at <http://www.aei-brookings.com/publications/abstract.php?pid=881>.

The government should consider whether it wants to facilitate the coordination of contract supply and demand by providing liquidity—that is, the posting of bids and offers that enable buyers and sellers wishing to trade the option to do so.¹⁹¹ By narrowing the bid-ask spread, the government would provide more profit opportunities for traders. This would encourage traders to reveal their existing private information through their trades and it would provide incentives for traders to acquire new information.¹⁹²

Third, the government should also consider conducting or funding specific policy experiments, or both. These experiments should be selected so that they have a reasonable chance of success and are not controversial. For example, a controversial market related to predicting the probability of a terrorist event in some time frame would not make sense, at least at the outset. A market for claims on GDP conditional on the passage of a proposed tax cut might make sense. This market would potentially attract a great deal of expert, popular, private, and public interest. Such a market would also be unlikely to incite a public outcry.

Fourth, the government could experiment with stimulating new policy suggestions that have a potentially high payoff. It might, for example, provide a reward for specific policy suggestions that turn out to be “winners” and are adopted. The value of the reward could be equal to a fraction of the difference between the net benefits on the best proposal and the second-best proposal. Alternatively, the government could simply work to lower the cost of introducing new policy ideas into the debate by setting up and monitoring online idea forums for proposals.

191. Offering liquidity is analogous to selling an option to trade. See Thomas E. Copeland & Dan Galai, *Information Effects on the Bid-Ask Spread*, 38 J. FIN. 1457 (1983). The price of the option is equal to one half of the bid-ask spread.

192. The government should try to avoid holding contract inventory because it could send misleading signals to market participants and lead to avoidable ethical objections to market intervention. To meet this goal and still satisfy its market making function, the government could adjust the bid-ask prices downward or upward so that its net inventory balanced to zero. If the government wished to mimic the behavior of a competitive market maker, the bid-ask spread would widen or narrow depending upon equilibrium loss or profit for the government. The spread required to maintain zero profits will depend upon the number of informed and uninformed traders in the market. More informed traders implies a larger spread.

Fifth, the government could help make price and order book information from information markets available at a very low cost. Requiring exchanges that list information market contracts to release real-time order book information publicly (via the Internet, for example) would be a non-intrusive and virtually costless method of providing a public good. On the other hand, it makes sense to allow private parties an opportunity to supply the information before imposing this requirement.¹⁹³

Even if there were widespread adoption of information markets to improve policy, the government would still need to consider how the new information should affect its analyses of costs, benefits, and other key factors. It is important to recognize that the process of assessing net benefits will not necessarily be automatic, even with extensive information markets. In many cases, the prices in information markets can only approximate the objective function of the decision maker due to the limitations of specifying contracts. The decision maker may want to include other factors or information in an analysis. For example, if only the direct cost to the government is estimated, some factor would need to be added for the deadweight cost of raising revenues.¹⁹⁴

Because the process of assessing net benefits from a policy will not necessarily be automatic, there is likely to be an important role for government analysts. That role should depend on the quality and scope of information obtained from information markets. If the quality is high and the information essentially answers the desired question, the decision maker may decide there is less of a role for the kind of government policy analysis that is done today. At the same time, independent analysts may be more inclined to provide such analysis because the basic information inputs to this analysis are easier to obtain.¹⁹⁵

193. It is unclear that private exchanges will display order book information to all parties whose decisions would benefit from this information. Not all such parties are prospective investors: liquidity-constrained consumers, for example.

194. See, e.g., Feldstein, *supra* note 55; see also ANTHONY B. ATKINSON & JOSEPH E. STIGLITZ, LECTURES ON PUBLIC ECONOMICS (1980). In principle, this factor could be added to the benefit function.

195. We might expect experts, hired by those with a stake in the outcome, to develop and publicize "independent" reports about policy consequences in an attempt to influence the information market's evaluation of those consequences. Based on the minimal impact of stock analyst reports on prices, the impact of such experts is likely to be small. At the same time, such experts might affect the policy options considered by the government.

As with any problem involving new technology or information, the government will need to learn how best to use that information in conjunction with other sources. It may, for example, want to issue guidelines for cost-benefit analysis that highlight the importance of using information markets when they are useful. These could be included in economic guidance provided by the OMB.¹⁹⁶

Both the government and researchers have a very important role to play in the design and evaluation of information markets. The government may choose to introduce procedures for evaluating the performance of these markets. An advisory board consisting of distinguished researchers could determine how to optimize contract structure, exchange rules, and government involvement to meet the objectives of information markets. Possible considerations include the ability of prices to forecast net benefits, the ex post net benefits of policies that are implemented, the costs of operating the markets, the scope for price manipulation, and the hedging opportunities afforded by contracts. The advisory board could examine, for example, how increases in market stakes and liquidity impact the forecasting accuracy of prices. Investigations such as this will be necessary to ensure that information markets are productively employed in public policy decision making.

Independent researchers can contribute by working to resolve other design and performance issues. First, much more thinking needs to be done about where such markets could be applied successfully. As there are a large number of potentially important applications, one of the problems in the short term may be to choose those with the highest potential payoff. For example, in the field of regulation, it probably makes sense to focus first on major regulations that have impacts on the economy of at least \$100 million annually. Over time, analyses can help shed light on the optimum number of such markets, particularly for purposes of policy design. In addition, more thought needs to be given to the kind of markets that are most useful. For example, the potential value of using information markets to provide information on the equity implications of

196. *See, e.g.*, Office of Mgmt. & Budget, Appendix C: OMB Draft Guidelines for the Conduct of Regulatory Analysis and the Format of Accounting Statements, 68 Fed. Reg. 5513–14 (Feb. 3, 2003), available at http://www.whitehouse.gov/omb/fedreg/2003draft_cost-benefit_rpt.pdf.

different policies depends in part on how policy makers wish to use them in decision making. At one extreme, information markets can be used simply to generate information on the incidence of different policies without affecting the actual decision. At the other extreme, information markets could be used to block policies that have adverse distributional consequences, or at least require that those policies only be implemented when certain redistributive objectives are met.¹⁹⁷

Second, research is needed to determine how such markets can inform important policy issues. For example, information markets could shed light on the nature of biases that may be introduced through ex ante analysis that does not take account of market data.¹⁹⁸ Similarly, one could attempt to identify patterns in prices and hence, net benefits, as information markets evolve over time. This analysis could reveal, for example, whether net benefits are likely to be overstated or understated when a policy is first proposed.

Researchers could also use information markets to compare different measures of net benefits.¹⁹⁹ Lawrence Goulder and Roberton Williams argue that in certain settings it is possible to estimate the difference in benefits resulting from partial and general equilibrium analysis.²⁰⁰ In principle it would be possi-

197. For a novel approach to promoting equity, see SHILLER, *supra* note 82, Part 3.

198. See, e.g., Winston Harrington et. al., *On the Accuracy of Regulatory Cost Estimates*, 19 J. POL'Y ANALYSIS & MGMT. 297 (2000) (examining trends in the biases of regulatory cost estimates). One of the problems with existing studies is that they often compare ex ante and ex post point estimates or make arbitrary assumptions about whether there is a significant difference between the estimates. Information markets can provide information on point estimates and variation that allow formal tests of significance. Another problem with many comparisons of ex ante and ex post point estimates is that they cannot easily address differences in the quality of the studies. Information markets could help address this deficiency by providing a common framework for making the comparison that is specifically aimed at aggregating diverse sources of private information. Finally, information markets could also eliminate problems that arise when different studies use different benchmarks for comparison.

199. In practice, all of the measures are likely to be imperfect for a variety of reasons. To take just one, the cost number used above only includes direct costs. It is well known that indirect costs and benefits could play an important role in certain policy arenas. See, e.g., Michael Hazilla & Raymond J. Kopp, *Social Cost of Environmental Quality Regulations: A General Equilibrium Analysis*, 98 J. POL. ECON. 853 (1990).

200. See Lawrence H. Goulder & Roberton C. Williams III, *The Substantial Bias from Ignoring General Equilibrium Effects in Estimating Excess Burden, and a Practical Solution*, 111 J. POL. ECON. 898 (2003) (showing that under typical conditions the simple "excess-burden triangle" formula substantially underestimates the excess

ble to compare partial and general measures of welfare based on information markets. For example, one could compare the impact of the education reform project identified earlier in terms of its impact on GDP versus a narrower definition of net benefits. The difference in the prices of the contracts with pay-offs dependent upon each measure of benefits would reveal the differences in the measures.

Finally, academics could study the potential impact and value of information from different sources, including information markets and conventional cost-benefit analyses done by experts.²⁰¹ The choice is not necessarily between using information markets or a trained analyst to provide most or all of the information on a cost-benefit analysis. Both sources of information could turn out to be useful. There is also an issue of how many information markets are necessary. For example, we argue that an information market that assesses benefits could be helpful in providing assessments of net benefits, and would possibly even raise revenues for the government, but that proposition needs to be examined more carefully.

The optimum utilization of different resources will likely vary with the situation. Where one is reasonably confident in the information provided by information markets, one could substitute markets for experts doing detailed cost-benefit analyses. Where one is less confident, one may want to use such experts or use them in tandem with information markets.

VI. CONCLUSION

This Article shows how a new decision framework for policy making can improve upon existing analytical tools, such as cost-benefit analysis. We explain how application of this framework can have important consequences for how government officials make public policy decisions. Using information markets in conjunction with pay-for-performance contracts can enhance both legislative and regulatory oversight activities. A

burden of commodity taxes, in some cases by a factor of 10 or more. This formula performs poorly because it ignores general equilibrium interactions, including interactions between the taxed commodity and the labor market.)

201. This could involve some combination of controlled experiments as well as natural experiments, as governments gain more experience with information markets.

key feature of our framework is that it would hold both lawmakers and regulators more accountable for their policies.

One might ask why a mechanism for implementing projects like the one proposed herein has not been designed or implemented already. The answer is likely that information markets are still relatively new and scholars and practitioners are just beginning to appreciate their potential. Further, there are strong forces that have an interest in maintaining the status quo, and the performance-based-policy paradigm could dramatically change the political landscape.

Our main contribution is to propose an efficient way to implement well informed policy decisions. We do this by linking and building upon the literatures on information markets and mechanism design. The prices in information markets can inform the mechanism design process, thereby making previously infeasible mechanisms feasible for the policy maker.

Information markets allow policy makers to estimate key quantities in the counterfactual world in which their policy is *not* implemented. This is an important advance. Among other things, it enables policy makers to assess the benefits and net benefits of a policy before it is actually implemented. This, in turn, allows the design of true pay-for-performance contracts in a wide array of policy domains.

Although we focus on public-sector decision making, the analysis is sufficiently general to apply to a wide range of problems in private-sector and not-for-profit decision making. The framework can be applied to any situation in which a decision maker has the resources, but not the necessary information and ability, to achieve her specified objective. There are a wide array of possible applications, including regulation, trade, taxation, and social policy. The private sector could use information markets to evaluate new investments, aid forecasting, improve decision making, and enhance risk-sharing opportunities. The not-for-profit sector could use these markets to evaluate new investments and inform decisions about how to design performance-based contracts.

To sum up the key points of our proposal: First, we show how it is generally possible to design contracts based on different contingencies whose prices will convey useful information on the costs and benefits of a number of policy choices, ranging from regulation to public works projects.

Second, we identify the strengths and limitations of using information markets to help improve cost-benefit analyses. With these markets, decision makers might do a much better job assessing critical public policy issues.

Third, we show how information markets can be used to provide a stronger foundation for implementing a variety of government oversight mechanisms that, up to this point, have been stymied because of difficulties in estimating costs and benefits. We consider two such mechanisms, a regulatory budget and the use of net benefit accounts, and show how our proposed mechanism can help make these concepts workable. We also show how legislators can use traditional budgetary controls in conjunction with information markets to exercise effective oversight and implement performance-based policies.

Fourth, we argue that information markets could promote greater transparency in governmental decision making, provide more accurate estimates of the efficiency and distributional impacts of different policies, provide a better understanding of uncertainties, help with sensitivity analysis, offer a low-cost way of assessing new policy proposals, finance government projects and regulations with positive net benefits, allow those affected by specific policies the opportunity to hedge risk, and aid in the design of policies. Information markets can also help assess the value of additional research on the decision to undertake a project. If these markets become more widely accepted as a way of informing policy decisions, they have the potential to change the political landscape and reduce rent seeking.

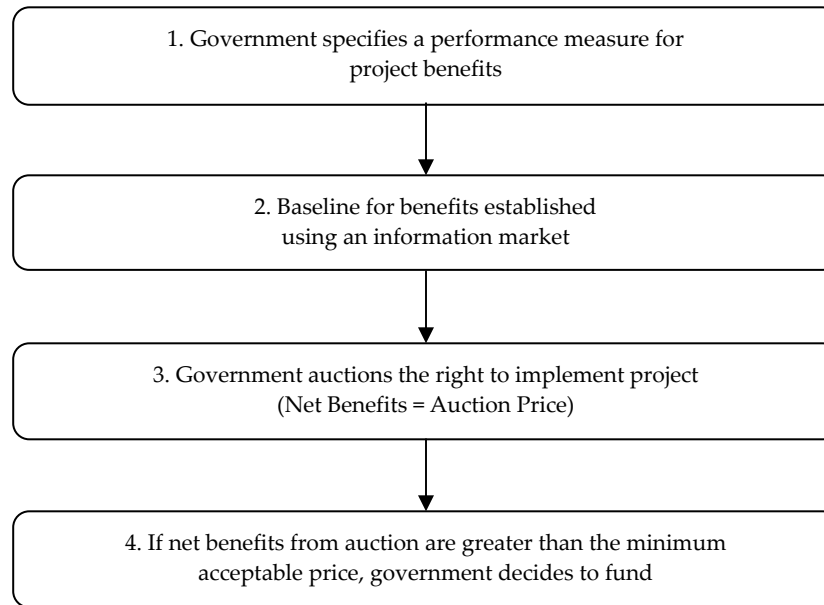
Fifth, we explain how government could play an important role in the expansion of information markets and how independent researchers could help in the development and assessment of these markets. Specifically, we recommend that government take a careful, iterative approach to facilitating the development of these markets.

Sixth, more research is needed to understand the theoretical potential of information markets as well as their practical relevance. The challenge for researchers will be to develop practical applications of these markets that can yield significant improvements in public policy. We think that goal is now within reach.

Finally, while we are enthusiastic about using information markets, they are only one of many tools available to policy makers for reaching decisions. The extent to which they can and should substitute for other tools should be examined carefully as their use in policy settings becomes more widespread. We would not advocate that information markets solely determine the outcome of key public policies at this point, but we would not rule out this possibility in the future.

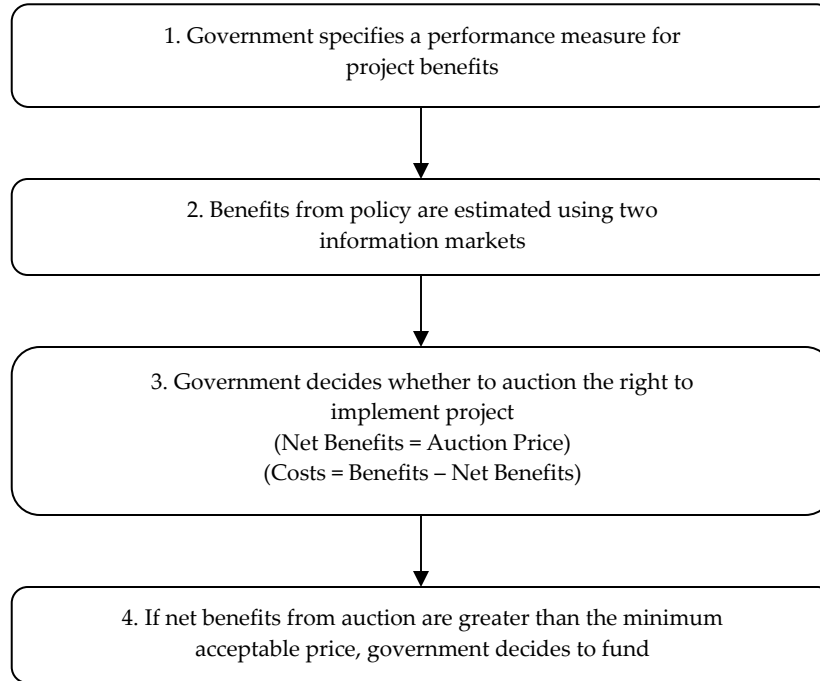
VII. APPENDIX

FIGURE 1.
Schematic for Project Selection



Note: The government must attach a monetary value to the performance measure. The baseline is the level of benefits expected in the absence of the project. In an auction with competitive bidding, the prevailing price will approximate expected net benefits from the project. The government's minimum acceptable price may or may not be publicly known to bidders. See text for additional details.

FIGURE 2.
*Schematic for Project Selection with
Early Information on Benefits*



Note: The government must attach a monetary value to the performance measure. The estimate of project benefits equals benefits with policy minus benefits without policy. In an auction with competitive bidding, the prevailing price will approximate expected net benefits from the project. The government's minimum acceptable price may or may not be publicly known to bidders. See text for additional details.

TABLE 1.
*Comparison of Direct and Indirect Cost-Benefit Methods
 Using Information Markets*

Issue	Direct Method	Indirect Method
Net social benefits estimated	Yes, by market participants	Yes, market estimates hypothetical analyst's net benefit estimate
Measure policy equity	Yes (same as above)	Yes (same as above)
Price includes all information available to market participants	Yes, assuming markets are efficient	No, unless market participants anticipate the hypothetical analyst possesses all information
Information contained in prices clear	Yes	No, cannot distinguish between information about the analyst vs. net benefits
Adapt to new information	Yes, through issuing new contracts	Yes, based on hypothetical analyst
Potential for bias in prices	Low (possibility of inefficient prices)	Depends on expectations of the hypothetical analyst's identity and motives and possible inefficiencies in prices
Estimate the distribution of net benefits	Yes, using options or other derivatives	No, cannot distinguish between variance in prices caused by changes in beliefs about analyst and new information
Assess changes over time in information about net benefits	Yes, by following the path of price over time	No, price changes could be caused by changes in beliefs about the analyst
Assess the <i>social</i> value of new information	Yes, using options	No, can only assess the value to a hypothetical analyst
Can incorporate unverifiable benefits	Yes, the policy maker can explicitly employ different cutoffs in his decision rule	Yes, the hypothetical future analyst can employ any arbitrary decision rule
Contracts can serve as a hedging vehicle	Yes, perfect hedge for project implementer and others affected by the project	Yes, but imperfect hedge because of price fluctuations from changes in beliefs about the analyst