



J O I N T C E N T E R
AEI-BROOKINGS JOINT CENTER FOR REGULATORY STUDIES

Big Ideas: The Market's Last Frontier

Robert W. Hahn and Paul C. Tetlock*

Related Publication 04-17
February 2005

This paper was published in the *Milken Institute Review* in 2005.

An earlier version of this paper was published in December 2004 on the
AEI-Brookings Joint Center website.

* Mr. Hahn is co-founder and executive director of the American Enterprise Institute-Brookings Joint Center for Regulatory Studies and a resident scholar at AEI. Mr. Tetlock is an assistant professor of finance at the University of Texas at Austin, McCombs School of Business. The authors would like to thank Roger Bate, Adam Lerrick, Robert Litan, Allan Meltzer, Roger Noll, Mary Shirley and Scott Wallsten for helpful comments, and Katrina Kosec for excellent research assistance. This paper represents the views of the authors and does not necessarily represent the views of the institutions with which they are affiliated.

Executive Summary

Information markets are markets for contracts that yield payments based on the outcome of an uncertain future event, such as a presidential election. They can provide real-time information on the likely benefits and costs of different kinds of policies and projects.

We argue that information markets combined with pay-for-performance contracts have the potential to revolutionize the way the government, the non-profit world, and the private sector do business. Moving to a performance-based policy paradigm could have great benefits for consumers and the economy. In addition to providing economic benefits, this approach could also promote greater accountability and transparency in the development of policy.

BY ROBERT W. HAHN AND PAUL C. TETLOCK

Back in the 1980s, a handful of academics at the University of Iowa's business school came up with an idea for giving students some hands-on experience in trading markets like the stock and commodities markets. Rather than using play money to simulate trading, they created a real market in which anyone could bet modest sums on the outcome of future events – for example, on who would be the next occupant of the White House. And they convinced government regulators that, because the market would be primarily a teaching device, it would not require oversight by the Commodity Futures Trading Commission.

Mighty oaks (well, in this case, a sapling) from acorns do sometimes grow: today, the Iowa Electronic Market (www.biz.uiowa.edu/iem/) is a thriving nonprofit enterprise, offering markets to wager on events ranging from the outcomes of presidential elections to the periodic interest-rate decisions of the Federal Reserve's Open Market Committee.

To see how the IEM works, consider the winner-take-all 2004 presidential election market. On Oct. 16, 2004, the price of a \$1 Kerry contract was 39 cents, implying that the market "believed" Kerry had a 39 percent chance of beating Bush (in previous weeks, the price of a Kerry contract had fluctuated from a high of 54 cents to a low of 28 cents). Thus, on Oct. 16, if you believed that Kerry's chances of winning were better, you had the

opportunity to put your money where your intuition was.

It is tempting to assume that the Iowa Electronic Market is just entertainment to make the sometimes-dreary task of learning to be a good securities speculator (or corporate commodities hedger) more palatable. In fact, so-called "information markets" are beginning to get respect from the pros – and for good reason. Just as Las Vegas parimutuel betting markets have a record of predicting the outcomes of sporting events better than professional gamblers, prices on the IEM have proved more accurate than pollsters in forecasting elections roughly 75 percent of the time.

Why do information markets work as well as they do? No individual expert is likely to know everything there is to know about the probabilities. Moreover, experts have complex motives – for example, the personal consequences of making the wrong bet may vary considerably from choice to choice. If a mu-

BOB HAHN is co-founder and executive director of the American Enterprise Institute-Brookings Joint Center for Regulatory Studies and a resident scholar at AEI. **PAUL TETLOCK** is an assistant professor of finance at the McCombs School of Business at the University of Texas.

BIG IDEAS

tual fund manager goes along with the crowd by betting the stock market will go up, he or she is unlikely to be fired if proved wrong. On the other hand, if he or she bets against the crowd and is wrong, the error stands out.

By contrast, anonymous markets are more likely to process all the available information, and to reward and punish bettors in a straightforward fashion. So the market price reflects what *The New Yorker* columnist James Surowiecki calls “the wisdom of crowds” in his fascinating book of the same name.

The corporate world is catching on to this insight. Indeed, firms ranging from Hewlett-Packard to Microsoft to Goldman Sachs are experimenting with information markets in making business decisions. We believe the superior ability of markets to amass and process information could also be harnessed in the service of public policymaking. These markets could improve the quality of information on which government bases decisions, as well as make politicians more accountable to the electorate. And in the process, information markets could make it more practical to enlist private enterprise in pay-for-performance arrangements to meet societal needs. Such a policy shift could result in a genuine renaissance of government in an era of scarce financial resources and even scarcer public confidence in the efficiency and benevolence of the political system.

If the idea of using markets to replace experts in predicting the seemingly unpredictable sounds a bit familiar – and just a bit bonkers – go to the head of the class. It would be foolish to introduce a novel way of government decision making by picking the hardest case first. Yet that is exactly what happened. In July 2003, it was disclosed that the Pentagon’s Defense Research Project Agency was planning to create a market in which trad-

ers could bet on where and when terrorists would strike.

This isn’t as dumb as it looks on first glance. We know relatively little about the motives or capacities of terrorists. And we are not inclined to trust experts, who can gain politically from predictions, to amass and process the information that is available. But a market for terrorism may have been unworkable. The prices would inform the government about terrorists’ attack plans, but they would also inform terrorists about the government’s security plans. What’s more, terrorists could profit from inside information, leading to public outrage.

Whatever the merits, Washington proved unwilling to venture into this brave new world. “There is something very sick about it,” concluded Senator Barbara Boxer of California. “I think you ought to end the careers of whoever ... thought that up.”

We would offer a less damning judgment: those promoting information markets should start simpler – and in an arena that is less controversial.

WHAT EVERY POLICYMAKER SHOULD KNOW

Suppose you were a policymaker assigned the task of improving standardized test scores in a failing urban public high school system, and decided it was worth \$1 million of public money for each point that test scores go up in a year. You could design a specific program with the goal of raising test scores by 10 points, and then ask private contractors to bid on providing the services. If someone were willing to undertake the task for less than \$10 million, you would presumably go forward with the deal.

Alternatively, you could create a pay-for-performance auction, in which contractors were invited to bid for the right to implement



a program to raise test scores and to receive \$1 million for every point that scores went up. Note there would be no bids unless at least one bidder expected to be able to improve the test scores for less than \$1 million per point. For example, if a bidder thought that a \$6 million outlay for after-school tutoring would raise average test scores by 10 points, it would have an incentive to bid up to \$4 million (\$10 million in pay-for-performance revenues, less \$6 million in costs).

This is a straightforward (if somewhat stylized) example of an approach that is oc-

casionaly employed in the public sector because it encourages contractors to use the most efficient means they can devise to get the job done. We think this generally superior approach could be made even better, however, if information markets were used in advance to inform both government officials and potential bidders.

For example, the government could create a futures contract that allowed people to bet on how much test scores will change in a year if a private contractor is engaged on a pay-for-performance basis. The terms would

BIG IDEAS

resemble those in futures contracts that currently allow speculators to bet on the future price of, say, jet fuel or wheat (and simultaneously permit airlines and food processors to lock in the future costs of vital inputs). If the market predicted that a pay-for-performance approach would raise test scores a year hence by X points, the government could use the information to decide whether to go forward with the program and potential bidders could use it to decide how much they could profitably offer.

Information markets could, in theory, do even more, providing insights into the net benefits of a policy initiative before it was even attempted. Here, futures contracts would be tied to the market's estimate of what bidders would be willing to offer for the right to a pay-for-performance contract.

In fact, social scientists have only begun to consider the potential value of using information markets to inform and complement innovation in the public sector. For example, with suitable tweaking, information markets could provide insights into who is likely to win and who is likely to lose from a given policy change. Thus, rather than just asking the market to predict how much average test scores would go up under a specified initiative, one could narrow the question to how much the test scores of children from low-income families would go up.

Information market contracts might even be used to hedge the risks associated with a public sector project, much the way commodities futures contracts are used to hedge against commodity price changes. For example, if a parent were worried that education quality would decline after some basic policy change, she could bet against the policy's success so she would later have enough money to send her child to private school. Information

market contracts would also help to finance worthwhile projects by allowing the winning firm to sell contracts at the market price.

One can certainly pick apart our examples – hypotheticals are always easy game. But we think the established idea of pay-for-performance contracting in the public sector augmented by the new idea of using information markets to make the pay-for-performance approach more practical offers some major opportunities.

TESTING THE LIMITS OF A NEW IDEA

For shorthand (and to satisfy the policy wonks' addiction to acronyms) we've named the combination "performance-based policy," or just PBP.

The basic PBP paradigm is straightforward. First, the decision maker defines a performance measure. In the case of the test score example, the government decided it was worth an additional \$1 million for each point that average test scores improved. Second, the decision maker uses information markets to assess what the policy will deliver in terms of benefits (average test score changes). Third, the decision maker opts for go or no-go. If it's a go, a pay-for-performance auction determines who gets the contract, assuming the winning bid exceeds some reservation price. If, for example, the government wanted to limit net expected payouts, it could set a high reservation price. Fourth, the decision maker pays for performance. If test scores go up, the government pays according to how much is achieved; if they go down, the contractor eats the costs.

As illustration of the potential advantages of performance-based policy, consider the recent "Copenhagen consensus" – a high-profile attempt to set priorities for solving societal problems that was orchestrated by the Danish environmental policy entrepreneur

Information market contracts might even be used to hedge the risks associated with a public sector project, much the way commodities futures contracts are used to hedge against commodity price changes.

neur Bjorn Lomborg. In May 2004, a group of eight distinguished economists (including three Nobel Prize winners) met in Copenhagen to consider the best ways to meet the biggest challenges facing the planet. To make the problem interesting, they assumed that governments had an additional \$50 billion to spend on solutions.

The experts ranked 17 social investments in four categories ranging from bad to very good. The very-good category included investments in controlling HIV/AIDS and malaria, reducing malnutrition and promoting free trade. The bad category included some investments in slowing climate change and instituting guest-worker labor transfer programs. To reach their decision, the experts relied on their collective wisdom, as well as research papers done by other experts and criticisms of those papers done by yet other experts.

There's nothing inherently wrong with that approach. But as we noted earlier, experts are only one source of information, and sometimes not a very good one. Now consider an alternative: using information markets. Suppose the Bill and Melinda Gates Foundation was considering investments in containing the spread of HIV in sub-Saharan Africa. It could use one market to predict the number of infections that would occur without intervention, while a second could predict the number of infections with a specific policy intervention – say, creating rural sex education centers. Suppose these predictions sug-

gest that one million HIV infections could be avoided if the foundation paid, say, a \$1,000 fee for every HIV infection below the baseline number. The foundation could then decide if it wants to go ahead with the project.

If it decides the project has promise, the foundation could use a pay-for-performance auction to select a contractor. In a competitive auction, bids would tend to reflect the difference between profits (defined here as monetary payments based on reducing the number of HIV infections) and the private costs of achieving reductions in the number of infections. Thus, the revenues from the auction would provide a measure of the net societal benefits from the project, which is a key criterion for allocating resources wisely.

Thinking big, health and safety regulation offer a variety of opportunities to apply the PBP paradigm. Suppose Congress decided to switch from command-and-control regulation – as in, automakers must install airbags for front-seat passengers – to an approach in which each agency is given a budget to meet broad performance objectives. The objectives could include everything from reducing highway fatalities to raising childhood vaccination rates to improving air quality.

The regulators would then translate these objectives into the pay-for-performance context by attaching a monetary value to performance improvements. And here, Congress could be explicit about priorities – for example, by putting a high value on air quality (hence, visibility) in national parks or a

BIG IDEAS

low value on coastal beach erosion. Next, the regulators would implement the PBP policy, using information markets to sharpen knowledge about costs and benefits and then choosing contractors through pay-for-performance auctions.

All this would amount to a revolution in the provision of government services. Policymakers would have straightforward ways of measuring success and failure. Moreover, the process would be totally open; it would be as easy to see market predictions on the success of air quality policies as it was to find the price of a Kerry contract on the Iowa exchange. And since the system would use public markets to award contracts, it would be relatively difficult for politicians to reward their friends or punish their enemies.

Finally, policymakers and the public could use information markets to assess the likely impact of prospective legislation, such as price controls on prescription drugs. Price controls, it is widely agreed, would reduce the incentives to develop new drugs; on the other hand, controls might well mean greater use of existing lifesaving drugs because more people (and government agencies) could afford to buy them.

Now, this is an arena in which it is hard to find an expert without a conflict of interest or an ideological bias. But information markets finesse the problem of generating an unbiased prediction – self-interest in the form of betting one's own money would presumably drive the results in the market. And if the public came to view markets as a more trustworthy source of information, lawmakers and bureaucrats would be less able to shield their own interests or biases behind the opinions of their favorite experts.

PBP hardly offers a better way in every circumstance. For example, if the government

can't figure out a practical way to pay for performance, PBP won't work. On the other hand, if policymakers can't define success and failure in concrete terms, perhaps we don't want them to spend public money in pursuit of ill-defined goals.

Arguably a more serious criticism is that PBP is, in effect, an effort to take the politics out of politics. If changing the way policy is made reduces the power of influential interest groups, they are bound to resist the change in the first place and attempt to subvert it once it is established. And since the groups most affected would be lawmakers and government bureaucrats, the natural enemies of PBP would be in charge of implementing it.

GETTING FROM HERE TO THERE

Start with what we've got. A variety of information markets are already flying below the enemy radar. As noted earlier, Hewlett-Packard has experimented with information markets, sponsoring markets among employees to predict sales. Eli Lilly has pursued similar experiments to predict successful drugs. The Iowa Electronic Market (www.biz.uiowa.edu/iem/markets) offers ways for anyone to bet on the presidential elections as well as the policies of the Federal Reserve. TradeSports (www.tradesports.com) is a far more ambitious for-profit effort to create information markets in sports, politics, economic indicators – even movie box office revenues. And Goldman Sachs (www.gs.com/econderivs) hosts auctions for derivatives based on the value of economic indices.

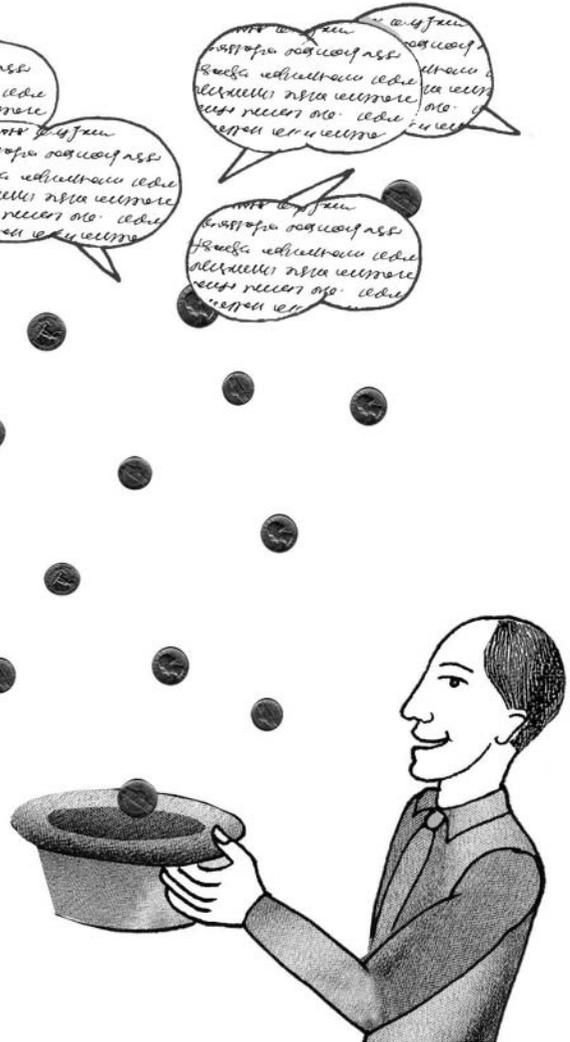
A major hurdle in making the public aware of the opportunities presented by information markets is (not surprisingly) government. The federal government does not have a clear policy on whether information markets should be regulated as futures markets – or worse, barred as interstate gam-

bling. And as information markets become better known, the states will be similarly tempted to regulate them through gambling commissions. Hence a critical step in advancing the use of information markets is to create a national policy toward them that both keeps federal regulators at bay and pre-empts all state regulation. We believe that, in the absence of federal legislation, the Commodities Futures Trading Commission should protect information markets from (other) government interference. Some federal legislation may be needed, however, to allow these markets to work.

Betting on presidential elections or box office results is still a fair distance, however, from using information markets to predict the impact of public policy. Making that leap would require some help – say, from government research agencies and nonprofits. The National Science Foundation has traditionally supported research in economics. It could provide seed money for research on ways to exploit information markets for purposes of policymaking. Currently, we don't know enough about the theoretical or practical properties of varying designs. Organizations like the World Bank could help with pilot projects. And the major philanthropic foundations might also be enlisted, as they have a big stake in improving the way public priorities are set and public services are delivered.

FINAL THOUGHTS

This is the place where you expect policy wonks like us to summarize the virtues of their proposals, belittle the nay-sayers and



leave the field of battle in alleged triumph. But we know the idea of using information markets to revolutionize public policymaking is a stretch. All we really ask of readers is to contemplate the failure of government-as-usual and to suspend disbelief that radical fixes could make a difference.

To date, social scientists have focused on all the reasons that government has expanded its reach, even as it has proved ever-less-capable of satisfying public demands. Surely it makes sense to experiment with ideas that, however modestly, offer hope of making government more responsive to public needs and more efficient at delivering services. **M**