

innovations

TECHNOLOGY | GOVERNANCE | GLOBALIZATION

Collaborative Innovation and Collective Intelligence

The Singular Insights of Many Minds

Lead Essays

Leadership and Innovation in a Networked World

Diego Rodriguez and Doug Solomon

Collective Intelligence to Address Global Climate Change

Thomas Malone and Mark Klein

Cases Authored by Innovators

Second Life: Collapsing Geography Cory Ondrejka

commentaries by Philip Evans, Paul Verkuil, and Thomas Malaby

Policy Analysis Market: A Thwarted Experiment Robin Hanson

commentary by Michael Abramowicz

Analytic and Policy Articles

The Principles of Distributed Innovation Karim Lakhani and Jill Panetta

Innovation Without Borders Bhaskar Chakravorti

Open Sourcing Social Solutions Charlie Brown

Science Common: Materials Transfer Agreement Project Think Nguyen

innovations

TECHNOLOGY | GOVERNANCE | GLOBALIZATION

Lead Essays

- 3 Leadership and Innovation in a Networked World
Diego Rodriguez and Doug Solomon
- 15 Harnessing Collective Intelligence to Address
Global Climate Change
Thomas W. Malone and Mark Klein

Cases Authored by Innovators

- 27 Collapsing Geography: *Second Life*, Innovation, and
the Future of National Power
Cory Ondrejka
- 55 *Case discussion: Second Life*
A Silicon Silicon Valley? Virtual Innovation and Virtual Geography
Philip Evans
- 62 *Case discussion: Second Life*
Contriving Constraints: The Gameness of *Second Life*
and the Persistence of Scarcity
Thomas M. Malaby
- 68 *Case discussion: Second Life*
The Values and Limits of Substitutional Sovereignty
Paul R. Verkuil
- 73 The Policy Analysis Market: A Thwarted Experiment
in the Use of Prediction Markets for Public Policy
Robin Hanson
- 89 *Case discussion: The Policy Analysis Market*
The Politics of Prediction
Michael Abramowicz

Analysis

- 97 The Principles of Distributed Innovation
Karim R. Lakhani and Jill A. Panetta
- 113 Innovation Without Borders
Bhaskar Chakravorti

125 Open Sourcing Social Solutions: Building
Communities of Change
Charlie Brown

Perspectives on Policy

137 Science Commons: Material Transfer Agreement Project
Thinh Nguyen

Letters

144

Organization of the Journal

Each issue of *Innovations* consists of five sections:

1. **Invited essay.** An authoritative figure addresses an issue relating to innovation, emphasizing interactions between technology and governance in a global context.
2. **Cases authored by innovators.** Case narratives of innovations are authored either by, or in collaboration with, the innovators themselves. Each includes discussion of motivations, challenges, strategies, outcomes, and unintended consequences. Following each case narrative, we present commentary by an academic discussant. The discussant highlights the aspects of the innovation that are analytically most interesting, have the most significant implications for policy, and/or best illustrate reciprocal relationships between technology and governance.
3. **Analysis.** Accessible, policy-relevant research articles emphasize links between practice and policy—alternately, micro and macro scales of analysis. The development of meaningful indicators of the impact of innovations is an area of editorial emphasis.
4. **Perspectives on policy.** Analyses of innovations by large scale public actors—national governments and transnational organizations—address both success and failure of policy, informed by both empirical evidence and the experience of policy innovators. The development of improved modes of governance to facilitate and support innovations is an area of editorial focus.
5. **Letters.** Readers comment on essays and papers published in previous issues of the journal.

Robin Hanson

The Policy Analysis Market

A Thwarted Experiment in the Use of Prediction Markets for Public Policy

Innovations Case Narrative: Policy Analysis Market

The past few years have seen an explosion of interest in prediction markets. We have long had speculative markets in gold, currency, pigs, and other commodities, which as a side effect do a remarkable job of aggregating information. Prediction markets turn this side effect into the main effect: if you want to know more on a topic, create and subsidize betting markets on that topic to elicit more accurate estimates. I have long been interested in how prediction markets can be used to improve decisions in the public arena. From 2001 to 2003 I had the opportunity to guide research on such markets that was sponsored by the U.S. government. The project, run by the Defense Advanced Research Projects Agency (DARPA), showed that general acceptance is still a long way off. Yet the academic support for the concept of prediction markets is old. In addition to the large literature on the information efficiency of financial markets (see Text Box 1), for several decades economists have been creating markets in the laboratory, showing since 1988 that markets with just a few traders trading for a few minutes can aggregate trader-held information.¹ Also since 1988, researchers at the University of Iowa have used a special legal exemption (which no one else has obtained) to run a series of real money betting markets on U.S. elections. Although these were far from the first election betting markets,² the added researcher-control they allow has led to new insights and academic attention.

Starting with the Foresight Exchange in 1994, a series of web-based markets were created that used play money and prizes to avoid the legal barriers that prevent the public from betting real money (see Text Box 2). The most successful of

Robin Hanson is associate professor of economics at George Mason University, and a research associate at the Future of Humanity Institute of Oxford University. After receiving his Ph.D. in social science from the California Institute of Technology in 1997, Robin was a Robert Wood Johnson Foundation health policy scholar at the University of California at Berkeley. In 1984, Robin received a masters in physics and a masters in the philosophy of science from the University of Chicago, and afterward spent nine years researching artificial intelligence, Bayesian statistics, and hypertext publishing at Lockheed, NASA, and independently.

Text Box 1. How Speculation Creates Insight

Prediction markets take advantage of an effect long known in finance: speculative markets do a good job of aggregating information. In any market, if one person makes an offer to buy or sell, someone else can accept that offer if they want. In a speculative market, a person can buy or sell something today in the hope of reversing her trade later for a profit. This doesn't work well with quick sales of houses because the process of buying and selling a house is so expensive, but it can be profitable in "financial" markets where many people frequently buy and sell items like currency, precious metals, corporate stocks, or grain futures. And when people can trade a durable standardized commodity at a low cost, speculation becomes cheap.

When speculation is cheap, all you need to succeed as a speculator is a bankroll and some method for predicting, with better than random accuracy, whether the prices will soon rise or fall. But if you find such a method and use it, it tends to become less useful. For example, if you find that prices generally rise on weekends, you might start to buy on Fridays and sell on Mondays. Yet doing this will make the price go up less on weekends, and so you will reduce the effect you are profiting from. Once many people look for and use such methods, it becomes hard to predict if the price will rise or fall, and so hard to profit by speculating in such markets.

Speculators thus compete to find information not yet embodied in market prices—so that it becomes hard to find information that such market prices do not embody. After all, anyone who finds neglected information can profit by trading on it, and thereby reducing this neglect. This idea that speculative markets do well at aggregating information has long been studied under the label "the efficient market hypothesis," but it is easy to be distracted by straw-man claims that such aggregation is perfect. Whatever else they disagree on, most everyone in finance agrees that such aggregation is very good.

While the information-collecting abilities of speculative markets are impressive, such markets have historically only been created for other reasons, such as to allow traders to hedge risks or have fun; the information-collecting effect has so far been a side-effect. If there is anything new in the current wave of interest in "prediction markets," it is the idea of creating and encouraging speculative markets on particular topics primarily to gain information on those topics. For example, a company that wants to know when its project will be completed can set up a market for employees and others to bet on such questions. If the current price is \$20 for the asset "Pays \$100 if project X completed by January 1," that can be interpreted as an estimate of a 20% chance of completing the project by that date.

these has been the Hollywood Stock Exchange, where over a million traders have since 1996 bet on movies and movie stars.³ In addition, real money betting markets based outside of the United States have grown enormously with the web, and these markets have sometimes offered interesting public policy-related topics. While most of these markets were created to help traders get rich, have fun, or collect bragging rights on topics of interest to them, the existence of and publicity about informative prices has at least helped to marginally shift social norms to favor believing betting market consensus.

Corporations are beginning to find internal prediction markets useful as a new kind of “forecasting institution.” When such markets are designed to directly inform particular decisions I call them “decision markets.” Many firms have created internal markets on topics important to the firm, such as input prices, sales, and project completion dates. Firms have created public markets that anyone can join, and also private markets limited to select employees. They have created both play money and real money markets in which employees are given their initial stake to circumvent anti-gambling laws. Such mechanisms seem to have great potential to remake corporate information flows.

This trend started when the first known internal corporate market was created at Xanadu in 1990, but sped up greatly in 2004 when *New Yorker* columnist James Surowiecki published *The Wisdom of Crowds*, a book that was soon widely read and cited. Another impetus was the DARPA’s “Policy Analysis Market,” which suddenly burst into public view in July 2003—and was immediately cancelled. In this paper I review the origins and development of this Policy Analysis Market.⁴

PROJECT TIMELINE FOR THE POLICY ANALYSIS MARKET

In 2000, Michael Foster, the program manager of the National Science Foundation’s quantum computing research program, learned about work by myself and University of Iowa researchers on prediction markets. He convinced DARPA to fund research on the topic starting in 2001.

DARPA’s first call for proposals went out in May 2001 under the name “Electronic Market-Based Decision Support.” The call basically said, “We’ve heard this works elsewhere; show us it works for problems we care about.” Proposals were due in August, and by December two firms had won SBIR (small business independent research) grants. The first was Neoteric Technologies, subcontracting to Martek and professors at the University of Iowa. I brought the call for proposals to the attention of the second firm, Net Exchange, founded by Caltech professor John Ledyard. Net Exchange subcontracted to two George Mason University professors (myself and David Porter), and later also to the Economist Intelligence Unit, a division of the *Economist* magazine. The Net Exchange project was later named the “Policy Analysis Market” (PAM), and my main role was as a system architect.

Text Box 2. From Xanadu to Graduate School, via “Idea Futures”

In 1983 I quit school to move to Silicon Valley, where by day I worked on artificial intelligence at Lockheed Research, and by night I worked on the Xanadu project attempting to create the web. The Xanadu folks were idealistic software engineers with contrarian opinions on many issues, such as on nanotechnology, cryonics, space, software security, libertarian politics, and of course the web. They thought their opinions were unfairly neglected, and hoped those opinions would gain wider acceptance in their envisioned future, where all writings would be online, readers would easily find rebuttals directly linked to criticized writings, and Google-like evaluations would combine reader-ratings of both writings and of other raters.

I had come to the Valley excited by two visions: one of intelligent computers' giving objective expert analyses of controversial issues,⁵ and the other of an internet-based world wide web organizing human knowledge and debate.⁶ However, I soon learned that “artificial intelligence” was many decades away, at least; what we had so far taught computers was dwarfed by oceans of unarticulated human knowledge. Regarding the web vision, I also questioned how eager readers would be to review rebuttals, and how often averaged reader-ratings would vindicate unpopular but valid claims.

My attention turned to markets. Contrarian opinions would find a fairer hearing in a world where opinionated folks were expected to put their money where their mouths were, and where betting market odds were accepted as a social consensus. In such a world, established insiders could no longer suppress contrarian outsiders merely by silence or ridicule, and outsiders could no longer claim the mantle of neglected rebel merely by making wild claims. The risk of losing bets would give both sides pause, and would give third parties more reasons to adjudicate the dispute.

I called this concept “idea futures,” and began to explore it by reading widely, giving talks, and writing in a Xanadu-like passionate style.⁷ I framed the concept mainly as a way to reform intellectual debate, especially public policy-related debate, via changes in social norms and research funding. But I also explored other applications. I developed and tested a board game, where players used a mechanical market maker to bet on an unfolding murder mystery. I wrote software to explore techniques for managing a combinatorial explosion of betting topics. And with Marc Steigler I created a market in 1991 within the Xanadu firm on when they would finish their software. These initial explorations came to an end around 1992. Sensing I would be able to make progress on this or other ideas for institutional innovations only with formal credentials and an expanded network of contacts, I returned in 1993 to graduate school, at Caltech.

The Xanadu group failed and disbanded, but Tim Berners-Lee elsewhere successfully developed his World Wide Web using a different approach. In 1994 Mark James and Sean Morgan adapted my design for an email-based play-money game to the web, starting the first web-based betting market, now known as the Foresight Exchange.⁸

The plan was for two firms to receive \$100,000 for a six-month phase I; one of them would be awarded \$750,000 to continue phase II over two more years. There was also the possibility of receiving \$100,000 for the six months between these phases. More money became available than initially planned, so in fall 2002 both firms were funded for phase II, and Net Exchange applied for and won between-phase funding. Also during 2002, John Poindexter, convicted in the Iran-Contra scandal, became a DARPA executive. Foster's FutureMAP program was placed within Poindexter's organization, the Information Awareness Office (IAO). In December 2002, DARPA called for proposals for related research, at this point using the name FutureMAP. In summer 2003, a half dozen teams—at Penn State, Metron, the Institute for Counter-Terrorism, George Mason University, Sparta, and BBN Technologies—were awarded \$100,000 each.

Neotek sponsored an end-of-phase I conference in June 2002 and showed a few demonstration markets, using its pre-existing software, on the spread of SARS and the U.S. Department of Homeland Security color security threat level. When FutureMAP was cancelled, Neotek had not identified its market topics and had probably spent less than half of its phase II funding. Net Exchange had spent about two-thirds of its phase II funding, and the new small projects had spent little of their funding. Michael Foster had asked for, but not received, \$8,000,000 more in FutureMAP funding for the next few years.

From the start, we planned to forecast military and political instability around the world, how U.S. policies would affect such instability, and how such instability would influence U.S. and global aggregates of interest, such as growth rates or oil prices.⁹ We reasoned that the cost to create markets does not depend much on the topic, but the value of estimates varies enormously with the topic; thus, the greatest benefit relative to cost would come from the highest value estimates. And what could be more valuable than to inform the largest defense policy decisions?

The importance of these decisions highlighted concerns about manipulation. Many had expressed concerns that interested parties might seek to distort market prices, to influence related decisions. For this reason the Net Exchange team from the start began laboratory experiments on price manipulation. We discuss these results in more detail below.

Charles Polk, Net Exchange's president, came up with the project's name, the Policy Analysis Market (PAM). The focus later narrowed to a smaller region, the Mideast, because of the Economists Intelligence Unit's high cost to determine what instability had actually occurred in each nation.

The final plan was to cover eight nations. Regarding every three month period during the two-year test, traders would create prices that estimate five parameters for each nation: its military activity, political instability, economic growth, U.S. military activity, and U.S. financial involvement. In addition traders would predict U.S. GDP, world trade, total U.S. military casualties, and total western terrorist casualties, and a few to-be-determined miscellaneous items. (The miscellaneous items were to be decided at the last minute based on suggestions from traders, with approval by the Economist Intelligence Unit.) This would add up to a few hundred

base parameters to estimate.

In addition, traders would be allowed to predict millions of combinations of the parameters, such as how moving U.S. troops out of Saudi Arabia would affect political stability there, how that would affect stability in neighboring nations, and how all that might change oil prices. Similar trades could, for example, have predicted the local and global consequences of invading Iraq, had the markets been ready then.

For many years before PAM, Net Exchange had specialized in combinatorial markets, where buyers and sellers could exchange bundles of items. So from the start, the plan was to develop combinatorial prediction markets. In phase I, Net

From the start, we planned to forecast military and political instability around the world, how U.S. policies would affect such instability, and how such instability would influence U.S. and global aggregates of interest, such as growth rates or oil prices... [W]hat could be more valuable than to inform the largest defense policy decisions?

Exchange put together a combinatorial market similar to its previous combinatorial markets and ran a complex simulation in which a dozen students traded over a few days for real money. Unfortunately, only about a dozen trades occurred, a serious failure.

In the interim phase, the Net Exchange team prepared for and ran lab experiments that compared two new combinatorial trading mechanisms with each other and with a tradi-

tional mechanism. In these experiments three traders set seven price estimates in five minutes, and then six traders set 255 price estimates in five minutes. One of the combinatorial market mechanisms was the most accurate.¹⁰ Net Exchange spent most of Phase II was implementing a scalable production version of this market mechanism. Because this mechanism requires a net subsidy to traders, \$50,000 was budgeted for this subsidy, and individual bets were to be limited to a few tens of dollars.

The PAM team was concerned that they might not attract enough traders in the final phase to achieve a meaningful test. We had considered running markets within government agencies, but found strong legal barriers to conditional transfers of money between agencies. Since no single agency was strongly interested in collaborating with them, we choose to create public markets.

On May 20, 2003, DARPA reported to Congress on its Information Awareness Office (IAO), and described FutureMAP in terms of how it might predict a

bioweapons attack against Israel. In June 2003, we began to tell people about our new web site, created by Polk, and to give talks to drum up interest. On the web site, as a backdrop to bold text, were faint background sample screens. In a small (less than 2 percent) section of two such screens, Polk had included as colorful examples of possible miscellaneous items an assassination of Yasser Arafat, a missile attack by North Korea, and the overthrow of the king of Jordan.

In summer 2003, the Senate but not the House cancelled IAO funding, which included all FutureMAP support, because of privacy concerns with another IAO project, "Total Information Awareness." Because of this funding uncertainty, the plan was to start PAM on September 1 with 100 testers, who would each be given \$100 to trade. Online registration of was to open August 1, and public trading by up to 1,000 initial traders was to begin January 1, 2004. The PAM team would be protected against anti-gambling laws because it was an agent of the Department of Defense.

On July 28, 2003, Senators Ron Wyden (Democrat, Oregon) and Byron Dorgan (Democrat, North Dakota) held a press

conference at which they released an open letter to John Poindexter complaining that the U.S. Department of Defense was planning a "terror market" for people to bet on terrorist attacks.¹¹ The senators emphasized PAM's association with John Poindexter, who actually had little involvement with PAM, and described PAM as being "designed to predict terrorist events," when in fact it was focused on geopolitical trends. Their main evidence was the "miscellaneous items" in the background of the PAM website. Wyden and Dorgan's main complaint was that "terrorists themselves could drive up the market for an event they are planning and profit from an attack, or even make false bets to mislead intelligence authorities." Yet a few tens of dollars would hardly pay for an attack, and the PAM team had already told DARPA about its lab experiments showing that manipulators do not hurt price accuracy.¹² I discuss these complaints in more detail below.

A media storm immediately ensued. Although five mostly positive media articles mentioning PAM had appeared in the previous few months, fifty mostly negative articles appeared on July 29. DARPA's public relations person was out of town

On July 28, 2003, Senators Ron Wyden (Democrat, Oregon) and Byron Dorgan (Democrat, North Dakota) held a press conference at which they released an open letter to John Poindexter complaining that the U.S. Department of Defense was planning a "terror market" for people to bet on terrorist attacks.

and unreachable when the story broke, so DARPA was silent and initial media reports were based mostly on the senators' complaint.

On July 29, Deputy Secretary of Defense Paul Wolfowitz told the Senate Foreign Relations Committee that he first learned of PAM reading the newspaper that day and that the project was terminated.¹³ During that crucial previous day, no one from the government had asked the PAM team whether the accusations were correct or whether the more offensive aspects could be cut from the project. Many slow wheels of government were stopped in their tracks. For example, the president's Council of Economic Advisors was considering endorsing PAM, but the political decision to abandon PAM was made quickly, at a high level, and based on little information or lower-level input.

On July 30, seventy-eight media articles on PAM appeared, even more negative. Newspapers reported that Poindexter resigned that day, and two months later all IAO research was ended. Over the following days, weeks, months, and years, more than 600 more media articles have mentioned PAM, many at first, and then gradually fading in frequency. Interestingly, the coverage gradually became more positive, and the most recent fifty articles on average give readers a positive impression of PAM.

In a statistical analysis, eleven indicators of how informative an article is—including time from the events until the article was published, citing someone with firsthand knowledge, article length, a news or an editorial style, author anonymity, and the awards, circulation, frequency, and topic specialties of the periodical—individually predict that more informed articles give readers a more favorable impression of PAM. In a multiple regression model using six additional control variables, including media types, political leaning, and the author's gender, all six of the statistically significant variables predict that more informed articles favor PAM more.¹⁴ The more informed articles were more favorable, and eventually the average article was favorable, but the political decision to cancel PAM seems unlikely to be reversed anytime soon.

FEARS OF "TERROR FUTURES"

Why did decision-makers feel that a research agency such as DARPA should not conduct research to see whether speculative markets might be useful in government intelligence? The dominant initial reaction to PAM seemed visceral and intuitive rather than analytic. People used descriptors like "absurd," "bizarre," "lunacy," "repugnant," "shocking," "sick," "turn the stomach," and "unbelievably stupid." Many pundits offered more specific reasons to justify this reaction. Here are three examples of the sets of reasons offered. The initial press release by Senators Byron and Dorgan said:

Senators . . . called for the immediate end of a . . . project ostensibly designed to predict terrorist events through the online selling of "futures" in terrorist attacks. . . . Surely, such a threat should be met with intelligence gathering of the highest quality—not by putting the ques-

tion to individuals betting on an Internet website . . . as wasteful as it is repugnant. . . . DARPA will not have access to their identities or funds. This promise creates the possibility that terrorists themselves could drive up the market for an event they are planning and profit from an attack, or even make false bets to mislead intelligence authorities. . . . the basics of communication and follow-through ought to be our primary weapons against the terrorist threat . . . Make-believe markets trading in possibilities that turn the stomach hardly seem like a sensible next step . . . We need to focus our resources on responsible intelligence gathering, on real terrorist threats. Spending millions of dollars on some kind of fantasy league terror game is absurd and, frankly, ought to make every American angry. What on Earth were they thinking?¹⁵

An early Washington Post news analysis said:

[The idea that] anyone with a credit card, a password and Web access . . . would be able to predict the future better than all those spies and experts out at the CIA . . . is also the latest and loopiest manifestation of a near-religious belief within the Bush administration in the power of markets to solve all problems . . . Would-be assassins and terrorists could easily use disinformation and clever trading strategies to profit from their planned misdeeds while distracting attention from their real target. Clever insiders like Jeffrey Skilling and Dennis Kozlowski made millions by fooling markets and manipulating prices, and I suspect Osama bin Laden could do the same with the Pentagon's proposed futures market. . . . The war against terrorism is not likely to be won by hiring more economists. It is going to have to be won the old-fashioned way, improving the government's intelligence network one spy at a time.¹⁶

Finally, a recent Nobel prize winner in economics said:

The Bush administration's naive belief in free-market economics reached a new level of absurdity . . . there are severe limitations in the ability of markets to provide accurate predictions; for instance, where markets have few participants and can be easily manipulated, or where there are large asymmetries of information, with some participants . . . having far more information than others. . . . what was [John Poindexter] thinking? Did he believe there is widespread information about terrorist activity not currently being either captured or appropriately analyzed by the "experts" in the FBI and the CIA? Did he believe that the 1,000 people "selected" for the new futures program would have this information? If so, shouldn't these people be investigated rather than rewarded? . . . If trading is anonymous, then it could be subject to manipulation, particularly if the market has few participants providing a false sense of security or an equally dangerous false sense of alarm. If trading is not anonymous, then anyone with information

about terrorism would be, understandably, reluctant to trade on it. In that case, the market would not serve its purpose. . . . the lack of intellectual foundation or a firm grasp of economic principles or the pursuit of other agendas has led to a proposal that almost seems a mockery of itself.¹⁷

These commentators seem to have three main complaints, beyond incredulity and revulsion: They disliked replacing skilled professionals with unskilled self-chosen amateurs; they feared that bad guys would mislead us via their trades; and they feared that bad guys would be rewarded for doing bad things.

ASSESSING THE FEARS

The first concern expressed—that of replacing professionals with amateurs—reflects a basic misunderstanding. PAM was a research project to test a new forecasting institution being considered as an addition to existing intelligence institutions, one that would help combine individual insights into a consensus forecast. Successful intelligence requires not only the collection and interpretation of pieces of information, but also that the information be combined into consensus forecasts and passed up the chain of command. A forecasting institution is a social context in which people spend resources and then produce forecasts. If we hold constant the participants, their resources, and the topics they forecast, and then we vary the institution, the institution that produces more accurate forecasts is the best. The question of who should participate in an institution is separate from the question of which institution works best for a given set of participants.

While amateurs might be able to contribute much more to defense intelligence than they do now, it is not unlikely that they can not. PAM was a preliminary test of a new institution, not as a test of who should participate in that institution. The concept was not to replace experts with amateurs—at most there was a hope that amateurs might be able to add more than they do now.

Of course there was no guarantee that the PAM test would have been successful. But at DARPA, an agency known for taking big chances, a guarantee of success was not the relevant standard. If PAM had a only one part in a thousand chance of improving by one part in a thousand the value we get from our half trillion dollars a year defense budget, PAM would have been well worth the investment.

Even if PAM had been wildly successful, it certainly would not replace existing intelligence agencies. At most it might have led to a new forum within which agencies could combine their efforts into a consensus forecast.

The second fear expressed was that bad guys would be willing to make losing trades in order to mislead us. Even if such manipulation increased the error in market prices, decision markets might still be useful. After all, every forecasting method has error; as long as you know roughly the level of error you can weigh how much to trust the forecast.

However, unless other traders have very shallow pockets or are prevented from

trading as much as they want, manipulation efforts should on average make prediction markets *more* accurate. This has been seen in lab and field experiments, and is predicted by solid theory.¹⁸

The key thing to understand is that all speculative markets have “noise traders,” and a manipulator is just another kind of noise trader. The main reward that attracts informed traders to markets is the opportunity to win against fools, mistakes, and people trading for other non-information reasons. If other trading behavior is held constant, adding more noise trades would increase price errors. However, when informed traders expect more noise trading, they change their behavior in two ways: they increase the size of their trades, and they spend more effort to become informed on the topic of the market.

As long as some other traders have deep enough pockets, or there are many traders with shallow pockets, there will be enough traders willing to accept the noise traders’ on average losing bets. (It is like poker players who prefer to play against inexperienced players; they are easier to win against.) If so, on average the overall increase in trading will completely cancel the increase in noise trading. And the increased information held by informed traders will then make the prices more accurate. Thus, markets with more noise trading have more accurate prices.¹⁹ Since manipulators trade for reasons other than their information about asset value, they are noise traders, and have the same effect as other noise traders.

The third main fear expressed was that bad guys might be rewarded for doing bad things. For example, some suspected that the September 11, 2001, terrorist attacks on the New York World Trade Center and Washington, D.C., were funded, in part, by trades of airline stock options. Similarly, many feared that the 1982 Tylenol poisonings were done to profit from short sales on the Tylenol stock. Airline stock prices did fall on September 11, as did the Tylenol stock at the 1982 poisonings. And a study has found that Israeli stock and currency prices respond to Israeli suicide bombings.²⁰ Nevertheless, we know of no examples of anyone using financial markets to profit from such sabotage. A thorough study of the September 11 attacks found nothing suspicious.²¹ Nor were any trades linked to the Tylenol poisonings.

But if this is not reassurance enough, the Policy Analysis Market had two additional protections. First it was not to be a terrorism market—the main topics were geopolitical trends, not the details of individual terrorist attacks. Second, trades there were going to be limited to a few tens of dollars—hardly enough to fund a major terrorist attack. Standard financial markets would have been much better place to try to profit from a terrorist attack.

EXPLAINING THE FEAR

Psychological research by Philip Tetlock and others on “taboo tradeoffs” can help us to make sense of the dramatic political reaction to the Policy Analysis Market.²² Tetlock and others study how people react upon learning that they or others have crossed a moral boundary. They find that people are not only outraged at anyone

who crosses such a moral boundary, but are almost as outraged at anyone who even thinks about crossing such a boundary.

For example, subjects might read a description of a hospital administrator who is considering cutting costs at the possible risk of patient lives. They are typically outraged at an administrator who chooses to cut costs, regardless of the financial situation of the hospital or the ratio of money saved to lives risked. Furthermore, they are almost as outraged at an administrator who takes several days to think about the decision, even if he ends up not cutting costs. Apparently someone who would even think of doing such a thing is considered nearly as morally bankrupt as someone who actually does it.

Bush had long defended Poindexter against attacks on his “Total Information Awareness” project, widely criticized as an attempt to collect and integrate databases on the public, but defending PAM seemed to have been beyond the pale. PAM appeared to some to cross a moral boundary, which can be paraphrased roughly as “none of us should intend to benefit when some of them hurt some of us.” (While many of us do in fact benefit from terrorist attacks, we can plausibly argue that we did not intend to do so.) So, by the taboo tradeoff effect, it was morally unacceptable for anyone in Congress or the administration to take a few days to think about the accusation. The moral calculus required an immediate response.

Of course, no one at high decision-making levels knew much about a \$1 million research project within a \$1 trillion government budget. If PAM had been a \$1 billion project, representatives from districts where that money was spent might have considered defending the project. But there was no such incentive for a \$1 million project (spent mostly in California and London); the safe political response was obvious: repudiate PAM, and everyone associated with it, especially John Poindexter. The Senators appear to have anticipated this outcome, attacking PAM in order to embarrass the Bush administration via its association with the freshly vilified Poindexter, and to taint that administration as mad about markets.

Since FutureMAP began under President Bill Clinton, it would probably have progressed similarly, had Al Gore beaten George W. Bush in the closest presidential election in U.S. history. Gore would not have appointed Poindexter, and Republicans were unlikely to have tried to paint Gore as mad about markets. Thus, but for an accident of history, PAM might have been tried.

The history of the aborted Policy Analysis Market demonstrates that the underlying “decision market” concept does not yet have much credibility with the public. The vulnerability of a market mechanism to such a political attack should not be too surprising: most familiar financial products, including stocks, insurance, futures, and options markets, were once prohibited by laws against gambling. It took a long time for the relevant industries to convince the public to see each of these products as not “just gambling.”

Of course, all of these products *are* ways to gamble, but, because they serve useful social functions, they have become politically and morally acceptable. Similarly, decision markets can also serve an important social function, especially when applied to government policy. Unfortunately, there is no large industry yet with an

interest in lobbying the public to see decision markets as more than gambling. In light of elected officials' fear of making the same political mistake twice, the PAM experience is unlikely to be repeated anytime soon. Nevertheless, decision markets remain a potentially valuable, if still maturing, tool for analysis of complex problems.

A revival of decision markets in the public sector will likely wait for a new generation of politicians, or perhaps some stunning successes with these mechanisms in the private sector. But eventually, perhaps after they have repeatedly demonstrated their superior information-producing capacities in relatively uncontroversial contexts, decision markets may yet be allowed to revolutionize how we make high-value policy decisions.

DISCUSSION

The Policy Analysis Market was one of the two main catalysts of the recent burst of interest in prediction markets. The other was James Surowiecki's popular 2004 book *The Wisdom of Crowds*.

Surowiecki's book praised averaged crowd opinions as wise in general when contrasted with expert opinions, and described prediction markets as just one of many ways to see crowd opinion. Others have described prediction markets as another form of "crowd-sourcing" social software, like Wikipedia, MySpace, and open source software. The crowd-sourcing vision also compares experts to the crowd, but focuses less on crowd accuracy and more on low crowd wages. Many recent business plans have been based on the idea of replacing expensive experts with hobbyist amateurs motivated by community bonds, and thus willing to work for little pay.

Both the wise crowd and the cheap crowd visions have been oversold. Yes, on some topics experts may be overrated and the crowd may offer valuable and neglected contributions. And yes, people often put a lot of effort into hobbies that tie them to a community. But on many topics crowds are fools while the experts know best, and hobbyists will be picky about which communities they will be willing to support with free labor.

Instead, prediction markets are a robust way to gain an accurate consensus regardless of whether crowds or experts are wiser. When experts know better, I expect experts to dominate market prices, and when crowds know better, I expect the crowd to dominate. Information is valuable enough to decision makers that they should be willing to offer substantial market subsidies in order to induce knowledgeable traders to participate.

A better way to frame prediction markets is as a new approach to the ancient human quest for social epistemology, which asks: who shall we believe? Or more precisely: what mechanisms shall we use to combine our varied information, analysis, and opinions into useful consensus opinions, and to encourage appropriate participation?²³ This question goes to the heart of conflicts within families, clubs, businesses, cities, nations, and broader intellectual communities. It encom-

passes issues like whether managers should hire “yes men,” whether to allow hearsay evidence in trials, whether to teach intelligent design in public school, how drug approval should account for unpublished drug tests, whether congressional earmarks should influence research funding, and so on.

Many of the most central and contested institutions in our societies, from religious hierarchies to courts to legislatures to peer review panels and ad hoc committees, have all been justified and contested in terms of the accuracy of the “consensus” conclusions they produce. Prediction markets are making a bid to be an important part of that mix.

Endnotes

1. Plott & Sunder (1988).
2. Strumpf & Rhode (2004).
3. See <<http://www.hsx.com/about/history.htm>>.
4. Polk et al (2003).
5. Lenat (1983).
6. Nelson (1981).
7. Hanson (1990, 1992a, 1992b). The following excerpt provides a flavor: “Academia is still largely a medieval guild, with a few powerful elites, many slave-like apprentices, and members who hold a monopoly on the research patronage of princes and the teaching of their sons. ... Perhaps the core problem is that academics are rewarded mainly for telling a good story, rather than for being right. ... Imagine a betting pool or market on most disputed science questions, with the going odds available to the popular media, and treated socially as the current academic consensus.” Others developed related, but not identical, ideas. In 1975 John Brunner published a science fiction novel describing a world that took betting market odds seriously on important policy topics, but with government manipulation turning this into a propaganda tool. (Nicholls 2006) In 1984 Willem Hofstee published “A betting reconstruction of empirical research,” but later clarified “Do not take the betting model literally”. (Hofstee 1984, 1991) In 1990, Zeckhauser and Viscusi mentioned in *Science* that betting markets might reduce biases in our estimates of important risks, but did not explore this idea further. (Zeckhauser & Viscusi 1990).
8. Hanson, Morgan, and James (1995). The Foresight Exchange grew within a year to a few thousand users, which helped greatly to publicize the concept.
9. It was my job to survey possible application areas and to recommend one. I recommended this one.
10. I designed this mechanism, which is described in “Combinatorial Information Market Design.” The experiments are described in Robin Hanson, John Ledyard, and Takashi Ishikida, “An Experimental Test of Combinatorial Prediction markets,” forthcoming in *Journal of Economic Behavior and Organization*.
11. Wyden and Dorgan (2003).
12. Hanson, Oprea, Porter (2006).
13. Michele Norris and David Welna (2003).
14. Hanson (2005).
15. Wyden & Dorgan (2003).
16. Pearlstein (2003).
17. Stiglitz (2003).
18. Strumpf and Rhode (2004), Camerer (1998), Hanson, Oprea, Porter (2006), Oprea, Tila, Porter, Hanson (2007), Wolfers & Zitzewitz (2004), Kumar and Seppi (1992), Hillion and Suominen

The Policy Analysis Market

- (2004), and Hanson and Oprea (2004).
19. Kyle (1989), Spiegel and Subrahmanyam (1992), Berg, Forsythe, and Rietz (1996).
 20. Eldor & Melnick (2004).
 21. Kean and Hamilton et al. (2004)
 22. Tetlock et al. (2000).
 23. Fuller (1988).

References

- Berg, Joyce and Robert Forsythe and Thomas Rietz (1996). "What Makes Markets Predict Well? Evidence from the Iowa Electronic Markets," pp. 444-463, in Wulf Albers et al. eds. *Understanding Strategic Interaction: Essays in Honor of Reinhard Selten*. New York: Springer.
- Brunner, John (1975), *The Shockwave Rider*, New York: Harper & Row.
- Camerer, Colin (1998). "Can Asset Markets Be Manipulated? A Field Experiment with Racetrack Betting." *Journal of Political Economy*, 106: 457-482.
- Eldor, Rafi, and Rafi Melnick (2004). "Financial markets and terrorism." *European Journal of Political Economy*, 20: 2 (June): 367-386.
- Fuller, Steve (1988). *Social Epistemology*, Bloomington, Indiana: Indiana University Press.
- Hanson, Robin (1990) "Market-Based Foresight-A Proposal." *Foresight Update*, 10:1,3,4.
- Hanson, Robin (1992) "Could Gambling Save Science? Encouraging an Honest Consensus", in *Gambling and Commercial Gaming: Essays in Business, Economics, Philosophy, and Science*, Ed. William Eadington and Judy Cornelius, pp. 399-440, 1992, Institute for Study of Gambling and Commercial Gaming, U. Nevada, Reno. (This is the proceedings of the Eighth International Conference on Risk and Gambling, London, Aug. 90.)
- Hanson, Robin (1992). "Idea Futures", *Extropy*, 8.
- Hanson, Robin (2005). "The Informed Press Favored the Policy Analysis Market," mimeo.
- Hanson, Robin, Ryan Oprea, David Porter (2006). "Information Aggregation and Manipulation in an Experimental Market." *Journal of Economic Behavior and Organization* 60:4 (August): 449-459.
- Hanson, Robin, Sean Morgan, and Mark James, (1995). "The Story of Idea Futures," pp. 54-59 in *Prix Ars Electronica 95, International Compendium of the Computer Arts*, ed. Hannes Leopoldseder and Christine Schopf, published by Osterreichischer Rundfunk, Linz, Austria.
- Hillion, Pierre, and Matti Suominen (2004). "The Manipulation of Closing Prices." *Journal of Financial Markets*, 7:4: 351-375 (October).
- Hofstee, Willem K.B. (1984). "Methodological Decision Rules As Research Policies: A Betting Reconstruction of Empirical Research." *Acta Psychologica*, 56:93-109.
- Hofstee, Willem K.B. (1981). "Do not take the betting model literally." *Kwantitatieve Methoden*, 3: 70-72.
- Kean, Thomas H. Lee H. Hamilton et al. (2004). *The 9-11 Commission Report*, (July 22). <<http://www.9-11commission.gov/report/911Report.pdf>>.
- Kumar, Praveen and Duane J. Seppi (1992). "Futures Manipulation with 'Cash Settlement.'" *Journal of Finance*, 47:4 (September): 1485-1502.
- Kyle, Albert S. (1989). "Informed Speculation with Imperfect Competition." *The Review of Economic Studies*, 56:3, 1989 (July): 317-355.
- Lenat, Douglas B. (1983). "The Role of Heuristics in Learning by Discovery: Three Case Studies," pp.243-306 in Michalski, Ryszard S., Carbonell, Jaime G., and Mitchell, Tom M. (eds.), *Machine Learning: An Artificial Intelligence Approach*, Palo Alto, CA: Tioga Publishing Company.
- Nelson, Theodor (1981). *Literary Machines*, Mindful Press.
- Nicholls, Henry (2006). "Tall stories and the twisted history of science." *New Scientist*, July 15, 2006.
- Norris, Michele and David Welna (2003). "Reaction to Pentagon's Proposed Terror Futures Market," All Things Considered, National Public Radio, (9 p.m., July 29, 2003).

Robin Hanson

- Oprea, Ryan, David Porter, Chris Hibbert, Robin Hanson, and Dorina Tila (2007). "Can Manipulators Mislead Prediction Market Observers?" April, mimeo.
- Pearlstein, Steven (2003). "Misplacing Trust in the Markets," Washington Post, July 30; p. E01. <<http://www.washingtonpost.com/wp-dyn/articles/A64741-2003Jul29.html>>
- Plott, Charles, and Shyam Sunder (1988). "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Market." *Econometrica*, 56: 5; p: 1085-1118.
- Spiegel, Matthew and Avanidhar Subrahmanyam (1992). "Informed Speculation and Hedging in a Noncompetitive Securities Market," *The Review of Financial Studies*, 5:2: 307-329.
- Stiglitz, Joseph E. (2003). "Terrorism: There's No Futures in It," *Los Angeles Times*, p. B13, July 31.
- Strumpf, Koleman S., and Paul W. Rhode (2004). "Historical Presidential Betting Markets," *Journal of Economic Perspectives*, 2004 (Spring), 18:2, pp. 127-141. <http://www.unc.edu/~cigar/papers/BettingPaper_10Nov2003_long2.pdf>
- Tetlock., Philip E., et al. (2000). "The Psychology of the Unthinkable: Taboo Trade-offs, Forbidden Base Rates, and Heretical Counterfactuals," *Journal of Personality and Social Psychology*, 78: 853-70.
- Wolfers, Justin and Eric Zitzewitz (2004). "Prediction Markets." *Journal of Economic Perspectives*, 18:2 (Spring): 107-126.
- Wyden, Ron and Byron Dorgan (2003). "Wyden, Dorgan Call for Immediate Halt to Tax-Funded 'Terror Market' Scheme," July 28. <http://wyden.senate.gov/media/2003/07282003_terrormarket.html> last accessed July 2006
- Wyden, Ron, and Byron Dorgan (2003). "Wyden, Dorgan Call For Immediate Halt to Tax-Funded 'Terror Market' Scheme," Press Release, July 28. <http://wyden.senate.gov/media/2003/print/print_07282003_terrormarket.html>
- Zeckhauser, Richard Z. and W. Kip Viscusi (1990). "Risk within reason." *Science* 248: pp. 559-564.