

Designing Real Terrorism Futures

Robin Hanson*
Department of Economics
George Mason University†

August 2005

Abstract

In July 2003, the Policy Analysis Market (PAM) was described as terrorism futures, and immediately cancelled. While PAM was not in fact designed to be terrorism futures, I here consider five design issues with implementing and using real terrorism futures: combinatorics, manipulation, moral hazard, hiding prices, and decision selection bias. As neither these nor other problems seem insurmountable, terrorism futures appears to be a technically realistic possibility.

Introduction

In July 2003, a great furor arose over a Pentagon research project called the Policy Analysis Market (PAM). Billed as betting on terrorism, PAM was widely denounced by politicians as morally repugnant and immediately canceled, which quickly led to the resignation of Ex-Admiral John Poindexter (of Iran-Contra fame).

At the time, observers noted many potential problems with using betting markets to learn about the details of terrorist attacks. These included the difficulty of getting the few people with inside knowledge to participate, the possibility of inducing terrorists to do bad things, and the possibility of misinformation from manipulative trades.

In fact, however, PAM was not intended to forecast the details of terrorist attacks. It was instead intended to forecast aggregate measures of geopolitical stability in the Middle East. PAM would have used speculative markets to estimate economic growth, political stability, and military activity four times a year in each of eight nations, and how those measures would

*For their comments I thank Colleen Berndt, Bryan Caplan, Rebecca Jarnas, Peter McCluskey, and Charles Rowley. For their financial support, I thank the Center for Study of Public Choice, the Mercatus Center, and Microsoft.

†rhanson@gmu.edu <http://hanson.gmu.edu> 703-993-2326 FAX: 703-993-2323 MSN 1D3, Carow Hall, Fairfax VA 22030

depend on each other and on various U.S. policy choices. PAMs designers thought their plan ambitious enough without also tackling the added complexities of predicting terrorism.

But since the issue has been raised, it seems worth considering how one might design speculative markets in order to make more specific predictions about terrorist attacks. In this paper, I therefore consider in more detail these design issues:

- *Combinatorics* – Is it technically possible for market forecasts to distinguish dozens of attack methods, policy responses, terrorist demographics, thousands of different times and locations, and perhaps billions or more combinations of all these?
- *Manipulation* – Can bad guys use their trades in order to mislead those who rely on market prices? If so, what can we do about it?
- *Moral Hazard* – How can we avoid letting traders have too large a positive betting stake in bad outcomes, so that people will not want to encourage bad events in order to win bets?
- *Hiding Prices* – The public might be unduly alarmed, and bad guys might be unduly informed, if prices told what we know about what bad guys are likely to do. How can we hide the most problematic prices, while still allowing markets to function?
- *Decision Selection Bias* – Conditional market estimates can be biased when traders expect decision makers to know more than traders do when making key decisions. How can we avoid this bias?

After reviewing the history of PAM, and the fears that people expressed when they thought PAM was intended to be terrorism futures, I discuss these five design issues and how we might deal with them if we were to create real terrorism futures markets.

Prediction Markets

Scholars have long noticed that speculative markets do a great job of aggregating relevant information. In fact, in every known field comparison, speculative markets have been at least as accurate as other forecasting institutions. Orange Juice futures improve on National Weather Service forecasts (Roll, 1984), horse race markets beat horse race experts (Figlewski, 1979), Oscar markets beat columnist forecasts (Pennock, Giles, & Nielsen, 2001), gas demand markets beat gas demand experts (Spencer, 2004), stock markets beat the official NASA panel at fingering the guilty company in the Challenger accident (Maloney & Mulherin, 2003), election markets beat national opinion polls (Berg, Nelson, & Rietz, 2001), and corporate sales markets beat official corporate forecasts (Chen & Plott, 1998).

While speculative markets have traditionally been created for other purposes, recently some have created new markets specifically to take advantage of this information effect. Called prediction markets (Wolfers & Zitzewitz, 2004), information markets, virtual stock

markets (Spann & Skiera, 2003), artificial markets (Pennock et al., 2001), or idea futures (Hanson, 1990, 1995a, 1995b), such markets are now used to estimate things like product sales, project completion dates, and election outcomes.

During the Clinton Administration, the Defense Advanced Research Projects Agency (DARPA) became intrigued by this concept. In 2000, Michael Foster, who ran (and still runs) the National Science Foundation quantum computing research program, convinced DARPA (the Defense Advanced Research Projects Agency, the blue-sky research arm of the U.S. Defense Department) to fund research on prediction markets starting in 2001.

This research program was eventually named “FutureMAP,” but the first DARPA call for proposals went out 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.” The call went out in May 2001, for proposals due in August, and by December two firms had won SBIR (small business independent research) grants. The winners were Neoteric Technologies, subcontracting to Martek and professors at the University of Iowa, and Net Exchange, founded by Caltech professor John Ledyard, and subcontracting to myself and David Porter at George Mason University, and later also to the Economist Intelligence Unit (EIU).

The Policy Analysis Market

Net Exchange named its project for DARPA the “Policy Analysis Market” (PAM), and the media later called it “terrorism futures.” Here is the story of PAM (Polk, Hanson, Ledyard, & Ishikida, 2003; Hanson, 2006).¹, of which the author was a chief architect.²

DARPA’s initial plan was for two firms to get \$100,000 each for a six month Phase I, after which one of them would be awarded \$750,000 to continue a Phase II over two more years. There was also the possibility of getting \$100,000 for the six months between these phases. More money became available than initially planned, so in fall 2002 both firms were funded to continue to Phase II, and Net Exchange applied for and won interim funding. Also during 2002, the infamous John Poindexter, convicted in the Iran-Contra scandal (and who I have never met), 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, ICT, GMU, Sparta, and BBN, were awarded \$100,000 each.

Neotek sponsored an end of phase I conference in June 2002, and showed a few demonstration markets, using their pre-existing software, on SARS and the color security threat level. When FutureMAP was canceled, Neotek had still not identified their market topics, and had probably spent less than half of their Phase II funding. Net Exchange spent about

¹An archive of information on PAM is at <http://hanson.gmu.edu/policyanalysismarket.html>.

²Foster contacted me in December 2000. I asked David Porter about small businesses I might team with, and he pointed me to Net Exchange, who I then approached.

two thirds of their 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 over the next few years.

From the very start, the Net Exchange team began laboratory experiments on price manipulation, as this was a widely expressed concern (Hanson, Oprea, & Porter, 2007). Also from the start, we planned to forecast military and political instability around the world, how US policies would effect such instability, and how such instability would impact US and global aggregates of interest.³ The reasoning behind this choice was that since the cost to create markets does not depend much on the topic, while the value of estimates varies enormously with the topic, 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?

Our focused later narrowed to a smaller region, the Mideast, because the Economist Intelligence Unit charged a high price to judge after the fact what instability had actually occurred in each nation. The final plan was to cover eight nations. For each nation in each quarter of a year (over the two year final phase), traders would estimate five parameters: its military activity, political instability, economic growth, US military activity, and US financial involvement. In addition traders would predict US GDP, world trade, US military casualties, and western terrorist casualties, and a few miscellaneous items, to be determined by traders and the EIU. This would require ($8 \times 8 \times 5 =$) a hundred or so base markets.

In addition, we planned to let traders predict combinations of these events, such as how moving US troops out of Saudi Arabia would effect political stability there, how that would effect stability in neighboring nations, and how all that might change oil prices. Similar trades could have predicted the local and global consequences of invading Iraq, had such markets been ready then. (More on this in the later section on combinatorics section.)

On May 20, 2003, DARPA reported to congress on the IAO, and described FutureMAP in terms of predicting a bioweapons attack against Israel. In June 2003 the PAM team began to tell people about our webpage, and to give talks to drum up interest. Net Exchange president Charles Polk created the PAM website. This site used some faint sample screens as a backdrop to bold text, and a small (<2%) miscellaneous section of two of these background screens included the colorful examples of an Arafat assassination, a North Korea missile attack, and the king of Jordan being overthrown.

In the summer of 2003, the Senate but not the House had canceled IAO funding, which included all FutureMAP funding, because of privacy concerns with another IAO project, "Total Information Awareness." Due to this funding uncertainty, when the media storm hit the PAM plans were to start on September 1 with one hundred testers, who would each trade with \$100 given to them by PAM. Registration to be one of those testers was to open August 1, and public trading was to begin January 1, 2004.

The media storm hit on July 28, 2003, when senators Ron Wyden (Democrat, Oregon) and Byron Dorgan (Democrat, North Dakota) complained that the Pentagon was planning to letting Poindexter create a "Terror Market . . . designed to predict terrorist events" (Wyden

³It was my job to survey possible application areas and choose one.

& Dorgan, 2003). The sample screens and DARPA report were their evidence. As the DARPA public relations person was (atypically) out of town and unreachable that day, the fifty mostly negative media articles that appeared the following day were based mainly on the Senators misleading complaint. Politicians fell over themselves to denounce PAM. Within a few hours the Deputy Secretary of Defense told the Senate Foreign Relations Committee that he had just learned of PAM from the morning newspaper, and that it was being terminated. During that crucial day, no one from the government asked the PAM team if the accusations were correct, or if the more offending aspects could be cut from the project.

Why was PAM canceled? Since PAM was accused of crossing a moral boundary, a politician asking for time to investigate the details would have been considered almost as immoral as someone who endorsed the immoral crossing; the moral calculus of such situations requires an immediate response (Tetlock, Kristel, Elson, Green, & Lerner, 2000; Hanson, 2006). But no one at high decision making levels knew much about a one million dollar research project within a half trillion dollar defense budget. If PAM had been a one billion dollar project, representatives from districts where that money was spent might have considered defending the project. But there was no such incentive for a one million dollar project (spend mostly in California and London). So the safe political response was obvious: repudiate PAM, and all associated with it, especially Poindexter.

A day after PAM was canceled, John Poindexter resigned, and two months later all IAO research was canceled. Bush had defended Poindexter from attacks associated with the “Total Information Awareness” project, widely criticized as attempted to collect and integrate databases on the public. But defending PAM seemed beyond the pale. The attack appears to have been made anticipating this outcome, in order to embarrass the Bush administration via their association with the freshly vilified John Poindexter, and by tainting them as being a bit too mad about markets.

Over six hundred media articles mentioned PAM over the next few days, weeks, and years. A careful analysis shows that the more informed media articles gave a more favorable impression of PAM, and that eventually the typical article was favorable (Hanson, 2005b). The initial public reaction was rather negative, and public opinion probably remains negative, as the more recent positive coverage hasn’t reached most readers.

Since FutureMAP began under Clinton, it would probably have progressed similarly at first, had Gore beat Bush in 2000. But Gore would not have put the political lightning rod John Poindexter in charge, and the Republicans would have been far less likely to try to taint Gore as a bit too mad about markets. Thus but for the accidental outcome of the closest presidential election in U.S. history, PAM would probably have been tried. But as it is, PAM remains political poison; there seems little chance the government will do something like PAM anytime soon.

Terror Futures Fears

Since PAM was not intended to predict terrorist attacks, stopping PAM because it might fail at that task seems a bit misguided. But the question of whether speculative markets

could help to forecast terrorist attacks or improve terrorism policy remains valid. So why did people think that a research agency like DARPA should not even consider the possibility that speculative markets might be used to fight terrorism?

The dominant initial reaction to PAM seemed to have been visceral and intuitive, rather than analytic. People used descriptors like “absurd,” “bizarre,” “lunacy,” “repugnant,” “shocking,” “sick,” “turn the stomach,” and “unbelievably stupid.” PAM seemed to violate a taboo, i.e., to cross a basic moral boundary, which might be phrased as “none of us should intend to benefit when some of them hurt some of us.” Of course many of us do benefit from terrorist attacks; but it does not appear that we intended to so benefit (Hanson, 2006).

Many pundits were uncomfortable simply declaring their revulsion, and so they tried to also offer more specific reasons to justify this reaction. Here are three examples of sets of reasons offered. The initial Senators’ press release 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. . . . in a letter to . . . John Poindexter . . . Surely, such a threat should be met with intelligence gathering of the highest quality – not by putting the question 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? (Wyden & Dorgan, 2003)

A 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 loopyest 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. (Pearlstein, 2003)

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. . . . a terrorism futures market . . . John M. Poindexter . . . what was he 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. (Stiglitz, 2003)

These commentators seem to have three main points, beyond incredulity and revulsion:

1. They disliked replacing skilled professionals with unskilled self-chosen amateurs.
2. They feared that bad guys would mislead us via their trades.
3. They feared that bad guys would be rewarded for doing bad things.

The last two points seem like sensible concerns, and are discussed in more detail below. The first point, however, seems to be based on a basic misunderstanding. PAM was a *research project* to test a new *forecasting institution* being considered as an *addition* to existing intelligence institutions, to help *combine* individual insights into a consensus forecast.

Successful intelligence requires not only that various smart people dig up relevant clues and interpret them given their specialized knowledge. It also requires that various such efforts be combined into consensus forecasts and passed up the chain of command. We have often failed badly at this last task, sometimes because of political pressures to give higher-ups the conclusions they want to hear, but perhaps more often from the usual tendency of organizations to fail to coordinate and to keep information to themselves.

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 (all else equal).

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 it is not crazy to imagine

that amateurs might be able to contribute much more to defense intelligence than they do now, it is also not crazy to think that amateurs can contribute little more than they do now. PAM was intended primarily as a test of a new institution, and much less 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.

It is possible to limit who can participate in a market. In fact, the initial PAM plan was for government intelligence agencies to pay for their employees to participate. Strong legal barriers, however, were found prohibiting the kind of contingent payments between government agencies that would be required to settle bets. Doing the test entirely within a single large agency remained a possibility, but since no large agency showed much enthusiasm, and a hard time deadline loomed, the PAM plan switched to a public market. Marketing of PAM still focused on intelligence experts, in the hope of getting some of them to participate as private individuals. As long as some experts participated, PAM might have reasonably tested the concept of using speculative markets to create intelligence forecasts.

There was of course no guarantee that the PAM test would have been successful. But as a research project of 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, there was no realistic prospect of it *replacing* existing intelligence agencies. At most we might have added a new institution, a new forum within which agencies could combine their efforts into a consensus forecast.

Basic Terrorism Futures

Prediction markets can forecast more than just possible events, like elections or terrorist attacks. *Decision markets* can also advise policy choices via decision-conditional outcome estimates (Hanson, 1999). By forecasting important outcomes conditional on different particular choices, such markets can directly inform crucial decisions. For example, a decision market might say that unemployment is expected to be lower if you raise interest rates, compared to not raising interest rates.

Since the cost of creating a market is largely independent of the topic, the obvious place to look for high benefit-relative-to-cost applications is in advising our most important policy questions. Directly informing the highest-value decisions would seem to be the highest-value applications for prediction markets.

PAM was designed to use this decision market approach to forecast the aggregate effects of our largest Mideast policy choices, such as who we give money to and where we deploy our troops. Since one of the aggregate outcomes PAM would have looked at was total terrorist-caused deaths in the west, PAM might have given us some insight into whether actions like invading Iraq deter terrorism, by showing our resolve, or instead encourage terrorism, by making them mad. This could have been a valuable contribution to terrorism policy.

Other prediction markets could build on *red team* evaluations to make substantial contributions to terrorism policy. Unless the goal is to just give the public the illusion⁴ that “something is being done,” any sensible program to prevent terrorism must include red teams, i.e., independent groups on our side who pretend to be on the other side and attempt to penetrate our security defenses. Regarding airport security, for example, a red team might try to get dangerous items aboard a commercial airplane at a certain airport during a certain month.

If we had enough red teams attempting to penetrate our security, we could create a summary statistic regarding their rate of success, and we could use forecasts of that success rate to anchor useful decision markets. For example, decision markets could forecast red team penetration rates conditional on changes in who runs the major security agencies, or on big changes in the budgets of those agencies. Such markets could also forecast penetration rates condition on major policy changes, such as using public versus private workers for baggage screening. If red teams were independent enough to resist political pressures to make security appear better than it is, these markets could weigh in on whether we should change agency management, raise or lower budgets, or change major security policies.

Who would we want to trade in such markets? The relevant “experts” include not only DHS employees and contractors, but anyone with substantial insight into the nature of these security problems. Emergency workers, infrastructure specialists, airport employees, and even frequent airline travelers may well have useful insights to contribute to these complex forecasting problems. And as long as we stuck to broad summary statistics about penetration rates, we are unlikely to tell bad guys where our weak points are, and we are unlikely to induce our citizens to try compromise our security to win bets.

Advanced Terrorism Futures

Markets that directly advise our biggest terrorism policy choices make the most sense for intimal applications. But as experience, scale economies and innovations lower the costs of creating markets, it will come make sense to consider advising finer details of terrorism policy. Eventually, it might even make sense to try to forecast the details of individual terrorism attacks.

Let us therefore now let us consider the more extreme “terrorism futures” scenario that PAM was accused of embodying. That is, let us consider using prediction markets to forecast terrorist attacks.

One of the PAM webpage background examples that was used to accuse PAM of being a terrorism futures market was this: “Arafat assassinated during first quarter of 2004.” Note that if the price of this asset had risen, that would really do very little to help the Arafat security team prevent such an assassination. All they could really do is try to pay

⁴Many have the impression that this is in fact the goal of our security agencies. I know of an engineering company given a multi-million dollar contract who was scolded by DHS for actually building physical devices; it was only authorized to write a white paper. Similarly, many a traveler has suspected that TSA traveler inspections serve only to slow travel down, and have little or no effect on terrorism.

extra attention. The other webpage examples are similarly of only minor use for stopping terrorists.

To be of more use, the forecasts would have to be far more specific. Therefore, let us imagine markets intended to forecast specific details of upcoming terrorist attacks such as the attack time, location, target type, method of attack, and the terrorist demographics. We might further imagine markets to forecast *combinations* of these features, such as predicting the chance of an attack at a particular location and time combination.

Ideally one might hope to tempt a disaffected terrorist insider to tell what he knows about terrorist plans. But success in that goal is not required. More feasible might be to attract people who suspect their relatives, neighbors, or customers of being terrorists, and people who are close enough to the terrorist culture to know how the terrorists think. Even if one failed to attract these people, one could still collect valuable information from people who have insights into the relative vulnerability of different targets, and the effectiveness of different methods.

Such markets would not need to be highly accurate to be useful. If markets estimated that there was a one in a thousand chance of a terrorist attack in certain high risk situations, while low risk situations had only a one in a million chance, that would still greatly help us to focus attention on the larger risks.

Would such markets be feasible or wise? Let us consider the two main concerns expressed about PAM as terrorism futures, and three other concerns deserving attention.

Combinatorics

The first issue to consider is that this terrorism futures scenario is rather technically demanding. One might want to distinguish dozens of different methods of attack, dozens of possible terrorist demographics, hundreds of different times, and thousands of different locations. The number of combinations of these possibilities could easily reach into the billions or much more. And if one wanted to forecast correlations between different terrorist attacks, the number of possible combinations becomes even more vast.

The first problem to confront would be a serious user-interface problem. Traders would need effective ways to browse the current market prices for these billions of possibilities, as well as to monitor their current portfolio of bets regarding all these possibilities. They would also need effective ways to make offers and trades regarding many millions of possibilities at once.

Worse, there could be a serious *thin market* problem here. If a simple double auction market mechanism was used, for example, random offers would so rarely match each other that traders would learn to congregate at a few very standardized assets, ignoring most of the billions of possible variations.

Fortunately, a combinatorial betting mechanism was developed and tested by the PAM team that can overcome these problems. For many years before PAM, Net Exchange had specialized in combinatorial markets, where buyers and sellers can exchange complex packages of items. So from the start, the team's plan was to see how far we could go in developing

combinatorial prediction markets. During phase I of our DARPA funding, we put together a combinatorial market similar to their previous combinatorial markets, and at the end of Phase I we ran a complex simulation where a dozen students traded over a few days for real money. Unfortunately, only about a dozen trades occurred, a serious failure.

During the interim phase of our DARPA funding, we prepared for and ran lab experiments comparing two new combinatorial trading mechanisms we developed with each other, with the traditional double auction mechanism, and with a method for obtaining individual forecasts. These experiments, where six traders set 255 independent prices in three minutes, found that a combinatorial market maker was the most accurate (Ledyard, Hanson, & Ishikida, 2005). DARPA funding phase II was mostly being spent implementing a scalable production version of this market maker.⁵

This *market scoring rule* mechanism requires a subsidy per base combination, and limits the size of trades to match the subsidy level.⁶ When the number of combinations becomes too large to store explicitly on a computer, special computational mechanisms, still under development, can be used.

Effective combinatorial markets also make it easier to deal with the problem of self-defeating prophecies. For example, if markets predicted an attack on a certain airport at a certain time, officials might close down that airport near that time. While this action might thwart the attack, it would also punish the traders who warned of it.

To deal with this problem, we can let people trade on the chances of an attack *conditional* on the level of response. Traders could then say that there is a high chance of attack given that the airport stays open, but that this chance is much less if the airport is closed. Traders would then in effect recommend closing the airport. Allowing for these responses clearly makes the set of possible combinations even larger.

Let us imagine then that it is technically feasible to create markets to forecast the details of terrorist attacks, and consider what other kinds of problems such markets might face.

Manipulation

A widely-expressed concern about PAM as terrorism futures was that bad guys might “manipulate” prices. Now many kinds of activities are called “manipulation.” For example, in the recent accounting scandals corporate executives “manipulated” prices by basically lying to investors about their financial situation. Trying to deceive people by just directly lying is possible in pretty much any forecasting institution. If prediction markets are not more vulnerable to this problem than other institutions, the possibility of this problem does not argue specifically against prediction markets. No matter how individual opinions are combined into a consensus forecast, people can lie to try to influence those opinions and hence that forecast (Hanson, 2005a).

⁵I designed this mechanism (Hanson, 2003).

⁶The subsidy required goes as the log of the number of different times distinguished plus the log of the number of locations distinguished, and so on. PAM had budgeted \$50,000 for this subsidy, and individual bets were to be limited to a few tens of dollars.

Another kind of price “manipulation” is when a trader temporarily trades contrary to his information, in order control the rate at which his information is revealed. Because his profits come from trading against ignorant “noise traders,” who may appear at random times, a trader who is sure that no one else has his information prefers to wait and profit from a sequence of small trades instead of one big trade. And so as to not tip his hand via a sequence of recognizably similar trades, he needs to “mix it up” along the way (Chakrabortya & Ylmaz, 2004). This kind of “manipulation” seems relatively harmless. The best way to induce traders to reveal information quickly is to have them each fear that other traders may soon stumble onto the same information.

A third kind of manipulation is where traders make trades contrary to their information without later reversing those trades, and so lose money in order to try to fool those who make use of the prices. Now even if such manipulation were possible, there would be two straightforward ways to deal with it. First, one could limit who can trade in the market, in order to eliminate people likely to want to manipulate in this way. Whether to let someone trade would then be a tradeoff between the information they might add to increase accuracy and the noise they might add to decrease accuracy.

Second, even if manipulation resulted in added error, prediction market estimates would still be useful as long as we had calibrated that error. Every forecasting method has error. As long as we know roughly what that error level is, we can choose how much to rely on those forecasts. It is possible that errors intended to mislead may be especially large just when such errors might cause the most harm. But if we have a track record of both forecasts and reality, for situations both important and mundane, we should be able to calibrate the relevant error rates. Research projects like PAM can help such calibration efforts.

It turns out, however, that unless the informed 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. The key thing to understand is that all known speculative markets have a lot of “noise trading,” and a manipulator is mainly just another kind of noise trader. In theory, perfectly rational traders could use a subsidized market to aggregate their information and exactly reveal their common estimate (Hanson, 2003). Real markets, however, are full of fools, mental mistakes, and people trading for non-information reasons. In fact, in most markets the opportunity to win against such noise traders is the main profit incentive that attracts informed traders.

If we hold other trading behavior constant, adding more noise trading must increase price errors. But when other traders expect more noise trading, they change their behavior in two important ways.

First, informed traders eagerly scale up the size of their trades for any given amount of information they might hold, as this increases their expected profits. In the limit where the amounts traded are small compared to traders’ aggregate risk tolerance, this should fully compensate for the increased noise, leaving the price error exactly the same. That is, as long as there are a few participants with deep enough pockets, or enough participants with shallow pockets, there will be enough people willing to accept the noise traders’ on average losing bets. Of course formal limits on how much each person can trade may cause problems.

The Iowa Electronic Markets, for example, limits each person to trading from a \$500 deposit.

Now it may well be that because of risk-aversion, financial market traders do not fully correct for increases in aggregate noise trading in the world economy, at least along the handful of dimensions that command risk premia. Irrational traders who underestimate the risk they are taking on can create aggregate risks that rational traders cannot afford to eliminate (De Long, Shleifer, Summers, & Waldmann, 1990). But this doesn't seem very relevant for prediction markets, which do not estimate aggregate risks.

The second way that informed traders change their behavior in response to more noise trading is to increase their efforts to acquire relevant information. After all, the more noise trading there is the more profit there is to be made from informed trading. So, on net, more noise trading should *increase* price accuracy (Kyle, 1989; Spiegel & Subrahmanyam, 1992). And in fact, empirically it seems that financial and information markets with more noise trading, and hence a larger trading volume, tend to be more accurate, all else equal (Berg, Forsythe, & Rietz, 1996).

Models of financial market microstructure have considered several types of noise traders, including fools who act randomly, traders with immediate liquidity needs, traders who seek to manipulate a closing price in order to raise their futures market settlement (Kumar & Seppi, 1992; Hillion & Suominen, 2004), and more generally traders with quadratic preferences over the market price (Hanson & Oprea, 2004).

These models verify that more noise trading means more accurate prices, and that manipulators are just another kind of noise trader. A manipulator has hidden information about his bias, i.e., how much and in what direction he wants to bias the price. (This includes the possibility of zero bias, i.e., of not being a manipulator.) Other traders can respond only to the average expected bias. When the hidden bias is exactly equal to the average bias, it is as if there were no manipulator. When the bias is higher or lower than expected, the price will be higher or lower than expected. But competition between speculators ensures that on average the price is right, and that price accuracy does not decrease with more noise trading.

The data seems to confirm these theoretical predictions that manipulators do not decrease average price accuracy. While we know of at least one apparently successful manipulation attempt (Hansen, Schmidt, & Strobel, 2004), many have reported that manipulators have failed to reduce price accuracy, historically (Strumpf & Rhode, 2004), in the field (Camerer, 1998), and in the laboratory (Hanson et al., 2007). A recent review article concludes that “none of these attempts at manipulation had much of a discernible effect on prices, except during a short transition phase” (Wolfers & Zitzewitz, 2004).

Now even if manipulators reduce the size of price errors on average, there is still a logical possibility that they might also increase the harm from price errors. Imagine that the harm from a price error depended not just on the magnitude of the error, but also on some additional state that was positively correlated with the hidden manipulator bias. For example, in a market estimating the chance of a terrorist attack, terrorists might perhaps arrange for the size of the attack to be correlated with the forecast error. The market might then become more accurate in estimating whether an attack would occur, but miss the big attacks more often. In such a case the expected harm from price errors could increase with

more manipulation, even as the expected error decreased.

To avoid this problem, one should choose the parameters that markets estimate to be as close as possible to the actual decision parameters of interest. This makes it less likely that there are hidden states which modulate the magnitude of the harm from estimation errors, and which are correlated with some manipulator's bias. For example, it is probably better for a terrorist attack market to estimate the harm caused by an attack, such as lives or dollars lost, and not just whether an attack will occur.

The fact that adding more noise traders generally increases market accuracy suggests that there is little need to worry that allowing some group of people to trade will decrease price accuracy. Even if this were a group of fools, as long as the non-fools had deep enough pockets to trade against them, the net effect should be to increase average accuracy.

Moral Hazard

The other big concern expressed about PAM as terrorism futures was that it might actually fund terrorism. While the manipulation concern was that bad guys would give up money in order to reduce price accuracy, this moral hazard concern is that bad guys might be willing to increase price accuracy in order to gain money.

Some suspected that the September 11, 2001 terrorist attacks on the New York World Trade Center 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 (Eldor & Melnick, 2004).

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 (Kean, Hamilton, Ben-Veniste, Kerrey, Fielding, Lehman, Gorelick, Roemer, Gorton, & Thompson, 2004). Nor were any trades linked to the Tylenol poisonings. The closest example I can find is the case of Roger Duronio, a PaineWebber employee who in 2002 set off a logic bomb in one thousand company computers after investing \$20,000 in options, betting that the stock price would fall. There was \$3 million in damage, but system redundancy prevented any loss of data, the stock price did not fall, and Duronio was soon caught (Geller, 2002).

We do know of examples of murder to gain life insurance, where the insurance was purchased with this plan in mind. Thus speculation on sabotage is possible when one person can acquire a large enough stake in an asset whose value they can directly enough influence. We also have examples of extortion of large corporations, by people who first demonstrate their ability to cause large amounts of damage. Compared to speculating on sabotage, the extortion strategy runs a greater risk of detection, but requires less capital to implement.

This suggests that the need for secrecy makes it very hard for skilled labor and willing capital to find each other to implement the strategy of speculating on sabotage. Since relevant prices usually move for other reasons, one needs a large portfolio of sabotage acts to be reasonably confident of a net profit. But those who are well positioned to commit a single

act of sabotage are usually not well positioned to commit a stream of such acts. A willing source of capital would thus have to find many skilled saboteurs, and would risk detection with each new potential saboteur contacted.

Terrorism futures would have two key differences from most financial markets. First, prediction markets are typically very thin compared with most financial markets, with relatively little money changing hands. All else equal, this makes them very poor places to speculate on sabotage. Second, financial markets are typically tied to large economic aggregates, which are hard for individuals to reliably influence. Terrorism futures markets, on the other hand, would trade on events that small groups of people could substantially influence.

If terrorism futures markets remain very thin, we should not worry about funding terrorism. Police regularly pay informants small amounts to tell about crimes, even though informants are typically involved themselves in other crimes. While this provides a small subsidy to the criminal class, the information gained helps crime fighting and more than makes up for the small subsidy. We would happily pay a bank robber \$10 to tell us where he will strike next. Paying a similar price for terrorism information could also be a good deal.

If trading volume in terrorism futures should for some reason become large, then we might want to do more to deal with the moral hazard problem. I can imagine four different approaches. One approach is to limit what someone can do with their winnings. If they only win bragging rights, or if the winnings must go to their favorite charity, there would clearly not be much of a moral hazard problem. This solution might come at a high cost, perhaps greatly reducing trader incentives to get it right.

A second approach is to limit participation. This is similar to the way that insurance regulations now only allow those with an “insurable interest” to bet that bad events will happen. One could limit traders to those who are already trusted in some way, such as police, or allow people to become certified as trusted enough to trade large sums in terrorism futures. This approach is relatively crude, however, and may forgo most of the information held by those who are excluded.

A third approach is to limit individual trading positions. We might, for example, limit each not-especially-trusted person to gaining no more than \$20 from each terrorist attack. After all, no one is going to help a terrorist attack for a mere \$20. This approach would also require some loss in anonymity, to deal with the possibility of a person trading via multiple accounts. This approach might also make it harder for informed traders to overwhelm possibly numerous fools.

If it seems unhealthy or unwise to have large numbers of people who stand ready to gain even small sums from nearby terrorist attacks, one could make sure that everyone near an attack stood to lose on net. For example, one might impose a \$22 tax on everyone who lived or worked within 100 miles of a terrorist attack, and put the tax revenue into a pool with at least ten times as many taxpayers. If each trader was then limited to gaining no more than \$20 via betting that the attack would happen, then no local would stand to gain anything, and no non-local could gain more than \$20, from a terror attack.

A fourth approach is to have less than complete trading anonymity. Trading in financial markets today is mostly anonymous, but trading records are available on a limited basis to

those who are investigating crimes such as insider trading. Similarly, one could allow those investigating particular acts of terrorism to have limited access to records of trades regarding those acts (Hogg & Huberman, 2002). Traders might be asked to give up some aspects of their privacy in order to be certified as trusted traders. This approach is probably good enough to deal with most cases.

It seems unlikely that terrorism futures would become popular enough to make moral hazard a substantial concern. If trading volume were surprisingly high, however, we would have many ways to deal with the problem. Less than complete anonymity is probably the best general approach.

Hiding Prices

If we succeed in aggregating what we know about terrorist attacks into publicly-available terrorism futures prices, we might cause two related problems. First, we might unduly alarm and terrorize the public, who might over-react when there is a rise in the risk of an attack soon nearby. While rational agents should not be harmed on average by getting better information about events of concern to them, the public may not be rational in this way.

Second, we would in effect be constantly giving the terrorists something close to our best estimate of what they are likely to do. This would certainly help their planning, and perhaps this harm would outweigh the benefit of our having more information about what they are likely to do. These problems can be reduced, however, if the most problematic prices are not visible to the public. I can see two ways to make this possible.

The first approach is to drastically limit who can trade in terrorism futures. If publicizing market prices were illegal, then a small enough group of traders would be unlikely to leak the prices to the public or terrorists. Of course this solution might come at a large cost in terms of information not obtained.

The second approach is to limit how much traders can see about certain current market prices. It turns out that prices in ordinary financial markets change fast enough that most traders have to deal with the fact that they do not know what the exact market prices will be when their trading orders arrive at the market. So we might be able to hide more price information from traders by adapting and extending the techniques traders now use to deal with uncertain prices.

When prices change infrequently, it is enough for a trade order to say in effect “cancel this order if the price has changed more than this much from the number I saw.” When prices change more often, traders often put in “market” orders that say “trade this quantity at whatever the current price is” or “limit” orders that say “trade this quantity if the price is at least as good for me as this price.” The common strategy here is for each order to specify how that order should change in response to changing prices.

Imagine that everyone could see all prices about possible attacks to occur a week or more in the future, but that prices were hidden to most people regarding attacks less than a week away. Those who wanted to trade on attacks to occur within the next week would usually only know the last publicly-visible price. A person’s order to trade would then have to

specify how it responds to later price changes. Such a trader would not be told what trades had resulted from his order until after the attack date had passed, nor would he know the exact condition of his portfolio. In this way neither the public nor the terrorists would be alarmed or informed by changing estimates about attacks to happen within the next week.

In principle, one could allow traders to submit computer code that is given access to the entire hidden history of prices and trades, code that can use any such information in making its trading decision. One might let some set of trusted traders directly see the hidden prices and trades, but one would have to consider whether giving these traders this advantage would unduly discourage other traders from participating during this period.

Decision Selection Bias

There is one further complication that seems worth discussing. This complication is that market prices can give a misleading impression of what speculators know about decision-contingent estimates when decision makers know more than speculators.

Standard decision theory is clear that the best choice is the one that gives the maximum expected benefits (minus costs), conditional on making that choice, and given the decision maker's information. So when speculators know everything that decision makers know, the estimate of market speculators regarding some consequence, conditional on a certain decision, is exactly the sort of advice that decision makers should attend to. But having decision makers know more than speculators can cause problems.

Consider the example of estimating the chance of a terrorist attack on a certain building conditional on a certain change in security policy. Imagine that this policy change is one that some people expect to reduce the chance of a successful attack, but where reasonable arguments also suggest that this measure might actually increase that chance.

Figure 1 shows an example of how a decision selection bias can arise in such a context. There is a two-dimensional space of possible states, with each state describing the chance of an attack given that the security policy is changed, and given that the policy is not changed. Assume that the chief of security of the building knows the exact state, the point in this space, but that market speculators know only that the state is somewhere in the oval region shown. Also assume that the security chief will only change the security policy when doing so will lower the expected chance of an attack. Thus the diagonal line in Figure 1 separates the cases where the decision will be to change policy from the cases where the decision will be to not change policy.

Given where the oval happens to be in Figure 1, the best estimate of market speculators is the point labeled "true center", which suggests that on average it is better to not change the security policy. However, since speculators know that the actual decision will be made according to the decision line shown, speculators should condition their trading estimates on being on one side or the other of that line. The oval conditional on changing policy is shown in crosshatch, while the oval conditional on keeping the old policy is shown in stripes. The centers of these areas are marked, and combine to produce the point marked "apparent center," which gives the rational speculator prices in this case.

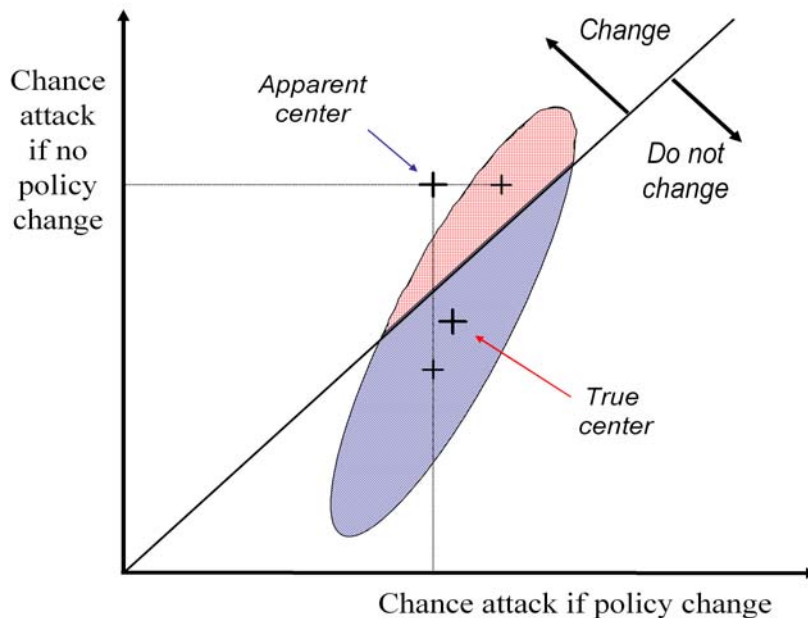


Figure 1: Decision Selection Bias Example

Thus a naive interpretation of market prices would suggest that market speculators think that the security policy should be changed, even though their best information is that the policy not should be changed. Thus decision-contingent prices can be misleading when speculators suspect that the decision maker may have relevant information that they do not. And this possibility might give decision makers an excuse to ignore market advice, even when decision makers do not in fact have better information than market speculators.

Furthermore, a similar problem can arise even when only market prices are used to make a decision. If speculators think that the decision will probably be made at a later date, and if they think that they may know more later, at that decision time, then the same decision selection bias framework applies. So if the decision is being made now based on market prices, but speculators think the decision will probably be made later, the same problem can occur.

To avoid this decision selection bias problem, one can either directly reveal decision maker information to speculators, or allow people with access to decision maker information to trade in these markets. One should also clearly inform speculators that a decision will probably be made soon, so that they do not fear that the decision will be made later when they know more. When speculators know that all decision maker information is in market prices, because the decision is being made now and insiders can trade, there can be no decision selection bias problem.

Another general approach to this problem is to make assumptions about causality relations and then use probabilistic representations of causality (Pearl, 2000) in order to describe the expected consequences *caused* by a given decision. For example, imagine that the se-

curity chief committed to rolling two dice and then if snake eyes came up, flipping a coin to change policy or not. In this case, it would suffice to look at prices conditional on snake eyes and on changing policy or not changing policy. When snake eyes did not come up, the security chief could look at the snake eyes market prices to inform his decision. This approach, however, depends heavily on either making correct causality assumptions, or on allowing the decision to be made randomly some substantial percentage of the time.

Conclusion

The Policy Analysis Market was accused of being a terrorism futures market, and was quickly canceled due to the resulting outrage. Many expressed simple incredulity and repugnance, but some attempted to identify more specific concerns. While some of these concerns seem to be based on misunderstandings, others are more reasonable and would deserve attention if one were actually going to create a terrorism futures market.

In this paper I have discussed five big concerns. Two of these concerns, manipulation and moral hazard, were mentioned often in the publicity surrounding PAM. Manipulation does not seem to be a real concern, nor do noisy traders more generally, at least when informed traders are not too severely limited in how much they can trade. Moral hazard is potentially a problem for high-volume markets, and several methods were identified for dealing with this problem, especially slightly reduced anonymity. But in fact terrorism futures is likely to be low-volume, making moral hazard a non-issue.

The three issues discussed here that were not mentioned much in the publicity surrounding PAM are: combinatorics, hiding prices, and decision selection bias. The combinatorics and decision selection bias problems are relatively technical, and once understood have relatively technical solutions. The problem of hiding prices, so as to not unduly alarm the public or inform the terrorists, also turns out to have a reasonable technical solution.

Thus in the end none of these problems seem insurmountable. If there were a political will to pursue this concept, terrorism futures markets would have a reasonable chance of helping us to deal with terrorism. No such political will seems forthcoming anytime soon, however. Let me suggest that this is because the U.S. public is not actually very afraid of terrorism at the moment. Were another large terrorist attack to occur on the U.S. mainland, this situation might change. As desperate situations call for desperate measures, a desperate public might be more willing to try a promising new approach like terrorism futures.

References

- Berg, J., Forsythe, R., & Rietz, T. (1996). What Makes Markets Predict Well? Evidence from the Iowa Electronic Markets. In Albers, W., Guth, W., Hammerstein, P., Moldovanu, B., van Damme, E., Strobel, M., & Selten, R. (Eds.), *Understanding Strategic Interaction: Essays in Honor of Reinhard Selten*, pp. 444–463. Springer, New York.

- Berg, J., Nelson, F., & Rietz, T. (2001). Accuracy and Forecast Standard Error of Prediction Markets. Tech. rep., University of Iowa, College of Business Administration.
- Camerer, C. (1998). Can Asset Markets Be Manipulated? A Field Experiment with Race-track Betting. *Journal of Political Economy*, 106, 457–482.
- Chakrabortya, A., & Ylmaz, B. (2004). Manipulation in market order models. *Journal of Financial Markets*, 7(2), 187–206.
- Chen, K.-Y., & Plott, C. R. (1998). Prediction markets and information aggregation mechanism: Experiments and application. Tech. rep., California Institute of Technology.
- De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Noise Trader Risk in Financial Markets. *Journal of Political Economy*, 98(4), 703–738.
- Eldor, R., & Melnick, R. (2004). Financial markets and terrorism. *European Journal of Political Economy*, 20(2), 367–386.
- Figlewski, S. (1979). Subjective Information and Market Efficiency in a Betting Market. *Journal of Political Economy*, 87(1), 75–88.
- Geller, A. (2002). ‘Pained’ Webber: Geek Tried To Sink Stock With Cyber ‘Bomb’. *New York Post*.
- Hansen, J., Schmidt, C., & Strobel, M. (2004). Manipulation in Political Stock Markets - Preconditions and Evidence. *Applied Economics Letters*, 11(7), 459–63.
- Hanson, R. (1990). Market-Based Foresight - A Proposal. *Foresight Update*, pp. 1,3,4.
- Hanson, R. (1995a). Could Gambling Save Science? Encouraging an Honest Consensus. *Social Epistemology*, 9(1), 3–33.
- Hanson, R. (1995b). Idea Futures. *Wired*, 3.09, 125. (excerpt appears in WSJ 30Aug95, A14).
- Hanson, R. (1999). Decision Markets. *IEEE Intelligent Systems*, 14(3), 16–19.
- Hanson, R. (2003). Combinatorial Information Market Design. *Information Systems Frontiers*, 5(1), 105–119.
- Hanson, R. (2005a). Foul Play in Information Markets. In Hahn, B., & Tetlock, P. (Eds.), *Information Markets: A New Way of Making Decisions in the Public and Private Sectors*. AEI-Brookings Press.
- Hanson, R. (2005b). The Informed Press Favored the Policy Analysis Market.. <http://hanson.gmu.edu/PAMpress.pdf>.

- Hanson, R. (2006). Impolite Innovation - The Technology and Politics of 'Terrorism Futures' and Other Decision Markets. In Patashnik, E., & Gerber, A. (Eds.), *Promoting the General Welfare: American Democracy and the Political Economy of Government Performance*. Brookings Institution Press.
- Hanson, R., & Oprea, R. (2004). Manipulators Increase Information Market Accuracy. <http://hanson.gmu.edu/biashelp.pdf>.
- Hanson, R., Oprea, R., & Porter, D. (2007). Information Aggregation and Manipulation in an Experimental Market. *Journal of Economic Behavior and Organization*. to appear.
- Hillion, P., & Suominen, M. (2004). The Manipulation of Closing Prices. *Journal of Financial Markets*, 7(4), 351–375.
- Hogg, T., & Huberman, B. A. (2002). Avoiding Moral Hazards in Organizational Forecasting. <http://www.hpl.hp.com/research/idl/papers/moral/moral.pdf>.
- Kean, T. H., Hamilton, L. H., Ben-Veniste, R., Kerrey, B., Fielding, F. F., Lehman, J. F., Gorelick, J. S., Roemer, T. J., Gorton, S., & Thompson, J. R. (2004). The 9-11 Commission Report.. <http://www.9-11commission.gov/report/911Report.pdf>.
- Kumar, P., & Seppi, D. J. (1992). Futures Manipulation with 'Cash Settlement'. *Journal of Finance*, 47(4), 1485–1502.
- Kyle, A. S. (1989). Informed Speculation with Imperfect Competition. *The Review of Economic Studies*, 56(3), 317–355.
- Ledyard, J., Hanson, R., & Ishikida, T. (2005). An Experimental Test of Combinatorial Information Markets.. <http://hanson.gmu.edu/testcomb.pdf>.
- Maloney, M. T., & Mulherin, J. H. (2003). The complexity of price discovery in an efficient market: the stock market reaction to the Challenger crash. *Journal of Corporate Finance*, 9(4), 453–479.
- Pearl, J. (2000). *Causality : models, reasoning, and inference*. Cambridge University Press, New York.
- Pearlstein, S. (2003). Misplacing Trust in the Markets. *Washington Post*.
- Pennock, D. M., Giles, C. L., & Nielsen, F. A. (2001). The Real Power of Artificial Markets. *Science*, 291, 987–988.
- Polk, C., Hanson, R., Ledyard, J., & Ishikida, T. (2003). The policy analysis market: an electronic commerce application of a combinatorial information market. In *ACM Conference on Electronic Commerce*, pp. 272–273 New York. Association for Computing Machinery.

- Roll, R. (1984). Orange Juice and Weather. *American Economic Review*, 74(5), 861–880.
- Spann, M., & Skiera, B. (2003). Internet-Based Virtual Stock Markets for Business Forecasting. *Management Science*, 49(10), 1310–1326.
- Spencer, J. (2004). New ICAP-Nymex Derivatives Have U.S. Gas Market’s Number. *Wall Street Journal*.
- Spiegel, M., & Subrahmanyam, A. (1992). Informed Speculation and Hedging in a Noncompetitive Securities Market. *The Review of Financial Studies*, 5(2), 307–329.
- Stiglitz, J. E. (2003). Terrorism: There’s No Futures in It. *Los Angeles Times*, B13.
- Strumpf, K. S., & Rhode, P. W. (2004). Historical Presidential Betting Markets. *Journal of Economic Perspectives*, 18(2), 127–141.
- Tetlock, P. E., Kristel, O. V., Elson, S. B., Green, M. C., & Lerner, J. S. (2000). The psychology of the unthinkable: Taboo trade-offs, forbidden base rates, and heretical counterfactuals. *Journal of Personality and Social Psychology*, 78, 853–870.
- Wolfers, J., & Zitzewitz, E. (2004). Prediction Markets. *Journal of Economic Perspectives*, 18(2), 107–126.
- Wyden, R., & Dorgan, B. (2003). Wyden, Dorgan Call For Immediate Halt to Tax-Funded ‘Terror Market’ Scheme. Press Release.