Designing Real Terrorism Futures

Robin Hanson*
Department of Economics
George Mason University†

May 2005

Abstract

In July 2003, the Policy Analysis Market (PAM) was accused of being terrorism futures, and immediately killed. 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 seems to be a 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 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.

*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

But since the issue has been raised, it seems worth considering how one might design speculative markets 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 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 a positive betting stake in bad outcomes, to avoid creating incentives for people to encourage bad events to win a bet?
- *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.

The Policy Analysis Market

People have long noticed that speculative markets do a great job of aggregating relevant information. In fact, in every known field comparison with another forecasting institution, such markets have been at least as accurate. 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).

Recently, some have created new markets specifically to take advantage of this 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, and began a program on "Electronic Market-Based Decision Support," later called "FutureMAP." DARPA funded two teams, one of which developed the "Policy Analysis Market" (PAM), a market for forecasting geopolitical instability (Polk, Hanson, Ledyard, & Ishikida, 2003; Hanson, 2005b).[1] By the end of the summer of 2003, PAM had spent about $850,000, over 2/3 of its total budget, and was nearly ready for beta-testing.

PAM was going to forecast the economic growth, political stability, and military activity of eight Mideast nations each quarter of a year. Using a new combinatorial trading technology that they had developed and tested, PAM was also going to forecast how all these measures would depend on each other and on various U.S. policy choices. A few global statistics like world trade, and some to-be-determined miscellaneous items were also going to be included. Because this combinatorial technology required a subsidy, for which $50,000 was budgeted, trades were to be limited to a few tens of dollars.

Along the way, some storm clouds appeared. During 2002, the FutureMAP program was placed under the newly-appointed John Poindexter. Though ex-Admiral Poindexter had little direct effect on FutureMAP, he was infamous for his association with the Iran-Contra scandal, and many were concerned that his new "Total Information Awareness" project would invade privacy. These concerns led the Senate to cut funding for all Poindexter projects. In May 2003, a DARPA report tried to convince Congress of FutureMAP's importance using an example of predicting a bioweapons attack against Israel. In June 2003, a new PAM website appeared that used some faint sample screens as a backdrop to bold text, and a small (<2%) miscellaneous section of two such screens included the colorful examples of an Arafat assassination, a North Korea missile attack, and the king of Jordan being overthrown.

These clouds set the stage for a media storm that began on July 28, 2003. Two Senators complained that the Pentagon was letting Poindexter create a "Terror Market ... designed to predict terrorist events." The sample screens and DARPA report were their evidence. As the DARPA public relations person was out of town and unreachable that day, fifty mostly negative media articles appeared the next day, based mostly 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. Poindexter resigned the next day.

Over five hundred media articles mentioned PAM over the next few days, weeks, and years. A media analysis shows that the more informed media articles gave a more favorable impression of PAM, and eventually the typical article was favorable (Hanson, 2005c). The initial public reaction was rather negative, and probably remains negative; the more recent positive coverage hasn't reached most readers. PAM remains political poison; there seems little chance the government will do something like PAM anytime soon.

[1] An archive of information on PAM is at *http://hanson.gmu.edu/policyanalysismarket.html*.

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 look into 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," "bizzare," "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, 2005b).

Most pundits were uncomfortable simply declaring revulsion, and 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 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. (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 myriad such efforts be combined into consensus forecasts and passed up the chain of command. We have often failed badly at this task, sometimes because of political pressures to give higher-ups the conclusions they want to hear, but 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 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 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 as a test mainly of a new institution, and 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.

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 against the kind of contingent payments between government agencies that would be required to settle bets. Doing the test entirely within a large agency was 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 the value we get from our half trillion dollars a year defense budget by one part in a thousand, 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, within which agencies could combine their efforts into a consensus forecast.

Advising Terrorism Policy

Idea futures markets can directly advise policy choices via decision-conditional outcome estimates (Hanson, 1999). PAM was designed to use this *decision market* technique to forecast the aggregate effects of our largest Mideast policy choices, such as who we give money to and where we put 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 encourage terrorism, by making them mad. This could have been a valuable, if limited, contribution to terrorism policy.

Since the cost of creating a market is largely independent of the topic, the big aggregate policy questions are the obvious place to start from a cost-benefit point of view. But as the costs of creating markets fall, it will 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.

But before considering that extreme, it is worth noting that more modest idea futures markets might still make substantial contributions to terrorism policy. For example, many suspect that most money spent by the Department of Homeland Security (DHS) is a waste, and is spent mainly to show voters that we are "doing something." I have heard 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.

Decision markets could be created to forecast the results of changing the DHS budget, or the DHS administration, in terms of the actual aggregate terrorism harm we suffer. Regarding airport security, such markets might forecast the percentage of "red team" agents who will successfully penetrate an airport's security system, conditional on changes in airport security policy. The relevant "experts" one might want in such markets include not only DHS employees and contractors, but anyone with substantial insight into the nature of these security problems. Emergency workers, infrastructure specialists, and even frequent airline travelers may well have useful insights to contribute to these complex forecasting problems.

Though such markets might be more reasonable and cost-effective, let us return to considering the extreme "terrorism futures" scenario than PAM was accused of embodying. That is, let us consider using idea futures 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 extra attention. The other webpage examples are similarly of little use for stopping terrorists.

To be of much use, the forecasts would have to be much more specific. Therefore, let 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 and hideout. We might further imagine markets to forecast the *combinations* of these features, such as that an attack at a particular location is particularly likely at a certain time.

Ideally one might hope to tempt a disaffected terrorist insider to tell what he knows about terrorist plans. But success on that count is not required. More commonly, one might hope to attract people who suspect their relatives, neighbors, or customers of being terrorists, people are close enough to the terrorist culture to know how the terrorists think, and people with insight into the relative vulnerability of different methods or targets.

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 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 get up into the billions or much more. And if one wanted to forecast the correlations between different terrorist attacks, the number of possible combinations becomes even more vast.

Now first there 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 their current portfolio of bets regarding all these possibilities. They would also want effect ways to make offers and trades regarding many millions of possibilities at once.

Worse, there could be a very 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 much better combinatorial betting mechanism was developed and tested by the PAM team. While this *market scoring rule* mechanism requires a subsidy per base combination,[2] it allows trading to continue smoothly with any number of combinations and traders (Hanson, 2003). In laboratory tests of difficult combinatorial problems, this mechanism was more accurate than several other mechanisms tested (Ledyard, Hanson, & Ishikida, 2005). 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

[2] E.g., the log of the number of different times distinguished plus the log of the number of locations distinguished, and so on.

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 idea futures are not more vulnerable to this problem than other institution, this is not an argument against idea futures. No matter what way individual opinions are combined into a consensus forecast, people can lie to try to influence those opinions and hence that forecast (Hanson, 2005a).

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 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 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 make sure traders reveal information is to for each trader with information to 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 with added error, idea futures 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. Of course errors intended to mislead may be especially large just when errors can 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 are very limited in how much they can trade, manipulation efforts should on average make idea futures *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 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 amount they trade 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 trader's 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 idea futures 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 information. 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 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.

The data seems to confirm the claim 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 failed attempts to manipulate speculative market prices with trades, historically (Strumpf & Rhode, 2004), in the field (Camerer, 1998), and in the laboratory (Hanson, Oprea, & Porter, 2005). 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 increase price accuracy on average, there is still a logical possibility that they might 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 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 can 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 manipulators bias. For example, it could be better for a terrorist attack market to estimate the harm caused by the attack, such as lives or dollars lost, and not just whether an attack occurs.

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 have deep enough pockets to trade against them, the net effect should be to increase average accuracy.

Moral Hazard

The other big concern that people expressed about PAM as terrorism futures was that it might 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 Roger 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, idea futures 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 enough to more than make up for this loss. 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 slightly reduce the anonymity of traders. Trading in financial markets today is mostly anonymous, but trading records are available on a limited based to those who are investigate crimes such as insider trading. Similarly, one could allow those investigating particular acts of terrorism to have a limited access 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.

A fourth 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. 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.

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, we could all know than no local would stand to gain anything, and no non-local could gain more than $20, from a terror attack.

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.

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 solved if the 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 an 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 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 to for a trade order to say in effect "cancel this order if the price has changed from this last number." When prices change more rapidly, traders often put in "market" orders that say "trade this quantity at whatever the current price is" and "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, and he would not be told what trades had resulted from his order until after the attack date had passed. 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, and can use any of that information in making its trading decision. If there were some set of trusted people who were allowed trade after seeing the hidden prices and trades, 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. 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 when decision makers know more than speculators, that can cause a problem.

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 from the cases where the decision will be to not change.

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.

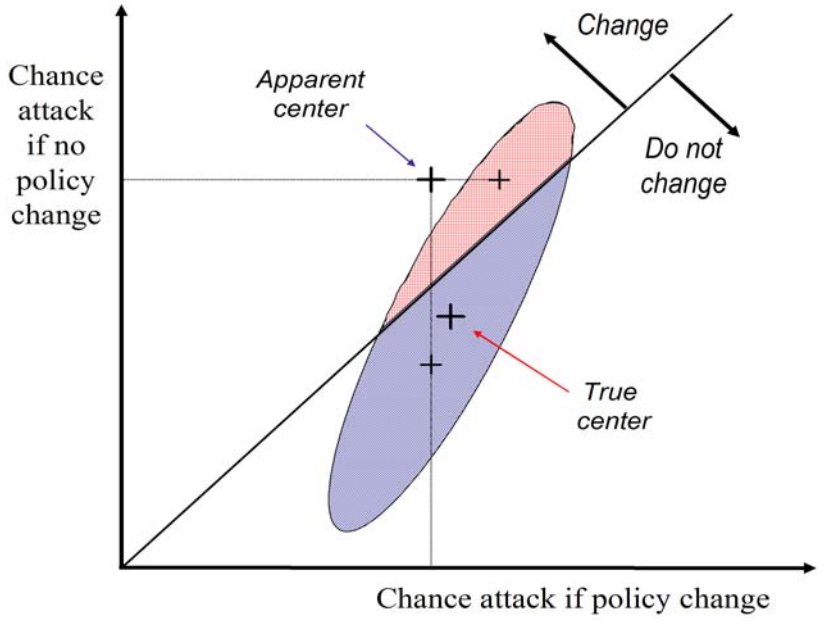

Figure 1: Decision Selection Bias Example

Thus a naive interpretation of market prices would suggest that market speculators think 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 they 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 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 is about to be made, 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 causal 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 the security chief committed rolling two dice and then if snake eyes came up, flipping a coin to change policy

or not. In this case it would sufficient to look at prices conditional on snake eyes and on changing policy or not. When snake eyes did not come up, the security chief could look at the 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 the resulting outrage. Many expressed simple incredulity and repugnance, but attempted to identify more specific concerns. While some of these concerns seem to be based on misunderstandings, others are more reasonable and deserve attention if one were actually going to create a terrorism futures market.

In this paper I have discussed five big concerns, two of which were raised often in the publicity surrounding PAM: manipulation and moral hazard. 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. 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 were: 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 an attractive technical solution.

Thus in the end none of these problems seem insurmountable. Were there a political will to pursue this concept, terrorism futures has a reasonable chance of helping us to deal with terrorism. No such will seems forthcoming anytime soon. 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, however, this situation might change. A desperate public might be more willing to try something new, 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., Gth, W., Hammerstein, P., Moldovanu, B., & van Damme, E. (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 Racetrack 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). 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. (2005c). The Informed Press Favored the Policy Analysis Market.. http://hanson.gmu.edu/PAMpress.pdf.

Hanson, R., & Oprea, R. (2004). Manipulators Increase Information Market Accuracy. http://hanson.gmu.edu/biashelp.pdf.

Hanson, R., Oprea, R., & Porter, D. (2005). Information Aggregation and Manipulation in an Experimental Market. *Journal of Economic Behavior and Organization*.

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.

Jorgensen, M., Teigen, K. H., & Molokken, K. (2004). Better Sure than Safe? Overconfidence in Judgment Based Software Development Effort Prediction. *Journal of System and Software*.

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., & Nielsen, F. (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.

Roll, R. (1984). Orange Juice and Weather. *The American Economic Review*, *74*(5), 861–880.

Santos-Pinto, L. (2003). Positive Self-Image and Incentives in Organizations. Tech. rep., University of California, San Diego.

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*.

Strumpf, K. S., & Rhode, P. W. (2004). Historical Presidential Betting Markets. *Journal of Economic Perspectives*, *18*(2), 127–141.

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.