

Learning from Terrorism Markets

By Adam Meirowitz and Joshua A. Tucker

On July 28, 2003, Senators Ron Wyden (D-Ore.) and Byron L. Dorgan (D-N.D.) held a press conference to call attention to a DARPA (Defense Advanced Research Projects Agency) project that would allow investors to bet on future political events in the Middle East—including terrorist attacks—through a commodity-market style trading system. On July 29, newspapers responded by denouncing this “unbelievably stupid” and “grotesque” market. Later that day, the program was officially cancelled.

What can policy-makers and scholars learn from the fury generated by these events? Despite a blizzard of attention, the quality of political debate over the market was poor. Politicians rushed to deliver grandiose statements condemning it, many of which found their way into equally incensed newspaper articles. Little of what appeared in print, however, contained developed arguments about how traders would behave in the market and whether it would actually serve as a useful tool for the intelligence community.

In this essay we make the following contributions. First, we highlight the shortcomings of the journalistic policy discourse through a simple survey of print media that appeared in the weeks following disclosure of DARPA’s plans. Second, we demonstrate how a methodology that has become standard in some quarters of academic political science could have improved the policy debate. Comparative analysis of political institutions is an approach that scholars use to understand how the structure of political institutions shapes the behavior of agents and eventual policy decisions. While few published accounts developed cogent explanations for how the terrorist futures market would function, even fewer attempted to base these claims on existing academic research. We fill this gap in the discourse by highlighting relevant lessons from experimental and empirical studies of similar markets and then tracing out how various components of a futures market design would affect its ability to function as a useful tool for aggregating information. In

doing so, we highlight an example where political science research could have led to a better policy debate and possibly a better policy outcome. Finally, we make explicit comparisons between a market based approach to aggregating intelligence information and the standard existing hierarchical structures.

Media and Political Treatment of the Market

Perhaps the defining characteristic of the proposed market saga was the lack of any high-level political debate. Practically every politician that weighed in on the matter opposed it. Wyden referred to the project as “make-believe markets trading in possibilities that turn the stomach,”¹ while Dorgan labeled it “unbelievably stupid.”² Senator Hillary Clinton (D-N.Y.) was widely quoted as calling the program “a futures market in death.”³

Unlike other debates, the criticism did not stop at party lines. Senator John Warner (R-Va.) called the program “a rather egregious error of judgment,” and Senator Pat Roberts (R-Kan.) said that it “defies common sense. It’s absurd.”⁴ Even Deputy Defense Secretary Paul Wolfowitz joined in the criticism, claiming that he had only learned about the plan himself by reading the newspaper and stating that “I share your shock at this kind of program. We’ll find out about it, but it is being terminated.”⁵

The response by the press was similarly unbalanced, although not quite to the extent of the response by politicians. A search of Lexus-Nexus from the end of July until late August turned up 115 pieces on the policy analysis market (hereafter PAM).⁶ Using a very blunt coding scheme, we broke down articles into three categories: “opposed,” “neutral,” and “in favor.” We cast our “in favor” net as wide as possible, including any pieces that had something positive to say about the proposed market. We also broke down articles into news or opinion pieces, with the latter comprising editorials, op-eds, letters to the editor, and columns. News articles were assigned to the “opposed” or “in favor” category only if they had a strong tone.

As Table 1 demonstrates, only 16 articles were in favor of PAM. This stands in marked contrast to the 56 pieces that opposed the project. While there were 43 articles in the “neutral” category, the vast majority of these were news articles. Of the 67 opinion pieces, 49 were opposed to PAM, with only 14 presenting arguments in favor.

This pattern was even more pronounced across the 25 editorials, with titles such as “The Pentagon’s Sick Take on Risk

Adam Meirowitz is an Assistant Professor of Politics at Princeton University. Joshua A. Tucker is an Assistant Professor of Politics and International Affairs at Princeton University. The authors thank Ellie Powell and Tatyana Karaman for their excellent research assistance, and Gary Bass, Tom Romer, the editor of Perspectives on Politics, and two anonymous reviewers for their helpful comments and suggestions.

Table 1
Number of Newspaper Articles on PAM

	Opposed	Neutral	In Favor	Total
News	7	39	2	48
Opinion	49	4	14	67
Total	56	43	16	115

Strategy”;⁷ “Dr. Strangelove Caught in Bizarre Betting Parlor”;⁸ “A Brainless Scheme”;⁹ and, quoting Dorgan, “Unbelievably Stupid.”¹⁰ Indeed, only one paper—the *New York Post*—presented an editorial in favor of PAM.¹¹

The primary media attacks paralleled those made by politicians, with the system being criticized as “grotesque,”¹² “morally reprehensible”¹³ and the result of the fact that “social scientists and the policy wonks who funded them [had grown] deaf to ethics and moral responsibility.”¹⁴ Some argued that such a market could actually encourage terrorist attacks by adding a profit motive: “Faced with a large pot of money and having few scruples, Gordon Gekko types might be tempted to influence events, even in friendly countries. Alliances don’t matter when there’s a buck to be made.”¹⁵

Although in the minority, a number of pundits did point to specific concerns about the market’s ability to function as desired rather than simply condemn it wholesale. Both Bill Day of the *San Antonio Express-News*¹⁶ and Jeff Brown of the *Houston Chronicle*¹⁷ wondered whether participants in the proposed market possessed enough information concerning the likelihood of future terrorist attacks to make it an effective mechanism for predicting such events. Senator Dorgan, writing in the *Washington Post*, noted that markets have also had spectacular failures, such as the stock market’s inability to foresee the collapse of the Enron Corporation. Economist William Sharpe raised the question of whether the market would have been truly competitive due to the presence of inside information. In an interview with the *Boston Globe* he noted, “I’d worry that someone betting on when Weapons of Mass Destruction will be discovered in Iraq might in fact already know where such weapons are.”¹⁸

One of the most interesting features of the various attacks on PAM was that hardly any contained a detailed description of the history of the program. The most notable exception was an article by Bradley Graham that appeared on the front page of the August 1 edition of the *Washington Post*. To judge from the vast majority of press coverage, PAM was designed for the sole purpose of allowing people to bet on the likelihood of future terrorist attacks. Graham reported, however, that

Some “specific event” contracts also were to be written, covering the possibility of a terrorist attack, assassination or coup. But these contracts, which stirred the political outcry this week, were never meant to be the market’s focus, said Charles Polk, Net Exchange’s president. “It was never billed as a market in terrorism,” he said. “It was to be a market in the future of the Middle East.”¹⁹

Graham also identified another interesting feature of PAM. The original plan of at least one of the companies hired to build the market, Net Exchange of San Diego, had been to limit market participation to “intelligence analysts and others in the U.S. government. But the idea of setting up an internal market ran into legal prohibitions against moving money among agencies funded under separate congressional appropriations.”²⁰

It is likely that PAM will go down in history as the government program with the shortest lifespan from entering the public consciousness to complete and utter termination. Clearly, Dorgan and Wyden had struck both political and public nerves. The question remains, however, whether they had actually helped public policy by doing so.

Analyzing the Institution

So what issues would have arisen in a better debate? In this section, we present three topics for assessment: the ability of markets to predict future events; the danger of manipulation of markets to affect non-market goals by market participants; and the effect of “noise” in markets on their efficiency. While none of these topics is as sexy as the prospect of Gordon Gekko types building a nest egg by correctly predicting mass destruction, all three point to important questions for assessing the functionality of the proposed PAM and, consequently, solutions for making a hypothetical PAM in the future work better.

First, the idea of using a futures market to generate predictions is not new. Following Friedrich Hayek,²¹ the predominant view among economists is that market prices efficiently summarize the private information of individual traders. Hayek’s perspective has motivated the creation of markets that trade commodities that pay off if a particular event occurs. For example, since 1988, the Iowa market has traded commodities corresponding to U.S. elections. In their analysis of this market, Robert Forsythe and his co-authors²² conclude that “the market worked extremely well, dominating opinion polls in forecasting the outcome of the 1988 presidential election.”

As some of the participants in the debate over PAM noted, most field experiments of this form have been quite satisfying.²³ However, political markets do not always perform well, and it is instructive to see which ones perform poorly. Ben Jacobsen and his co-authors study a political stock market in the Netherlands and find the market’s predictions to be much less accurate in multiparty races than in American elections.²⁴ The authors rule out many potential reasons for the poor performance and forward one explanation, which speaks directly to the proposed market on terrorism: If traders overweigh private information, then small parties will be overvalued and larger parties will be undervalued. In proportional representation systems with many parties, this effect is more pervasive than in U.S. elections. A market with futures contracts that pay off if terrorist activities occur is closer to a typical European proportional representation election than an American two party election. So while the few proponents of PAM that offer supporting evidence from the Iowa market are hitting on an important point, closer inspection of these field

experiments suggests that success in the United States may overstate the predictive value of markets in terrorist settings.

While the media tended to draw connections with the Iowa political markets, potentially more relevant comparisons existed in the experimental markets at Hewlett Packard (which, incidentally, were created by some of the same scholars who were involved in the DARPA project). The H.P. markets traded futures related to the sales of the company's products and its overall profitability. In contrast to the Iowa markets, which were open to the public, the H.P. markets were expressly designed to involve a small number of carefully chosen traders. "They were selected specifically from different parts of the business operation because they were thought to have different patterns of information about the targeted event. These patterns of information, including market intelligence, specific information about big clients, and pricing strategies, were in need of aggregation."²⁵ Traders could also track the trades made by individuals (although their real identities were suppressed). Thus, if a particular trader was repeatedly buying an asset, the other traders could observe this. Kay-Yut Chen and Charles Platt²⁶ demonstrate that the HP markets tended to outperform official forecasts and, in some cases, were remarkably accurate.

Early discussion of the terrorism market also focused on the consequences of having open registration for participation, which would potentially allow terrorists to profit from their activities. Senator Barbara Boxer (D-Calif.), in particular, feared this scenario because "knowing they were planning an attack, [terrorists] could have bet on the attack and collected a lot of money."²⁷ Some suspected that the real goal of the market was to gain information by luring knowledgeable collaborators out of the fold with the temptation of earning a quick profit, with one editorial claiming that "It was supposed to be a sting operation: Presumably, terrorists would know where they'd be attacking and what leaders they'd be gunning for, and would tip their hand on the Internet in hopes of making a killing on the killings."²⁸

But on closer examination, concerns about terrorists making money seem far less pressing than the concern that they would use the market to foil the intelligence community. As Dorgan noted, "the market invited manipulation."²⁹ Given the willingness of terrorists to expend resources—including lives—on their missions, it is not far-fetched to imagine that they would be willing to take a financial loss to mislead the Pentagon (Sowing havoc among the infidels: PRICELESS). At the least, this could lead to valuable resources being wasted; in the worst case scenario, it could increase the likelihood of a successful attack.

Finally, a related issue pertains to the relationship between the number of traders, the number of informed traders, and the amount of money available to traders. Information aggregation requires that individuals have information. The markets that predict electoral outcomes have been successful precisely because there is information (polls, news reports, et cetera) upon which participants can draw. The H.P. markets were successful because the traders possessed relevant information. But only a small handful of people are likely to have any advance knowledge of a terrorist attack. Consequently, in an

open market, most traders will be doing little more than guessing. Even worse, the "noise" produced by these bets will dilute whatever useful information might have been generated by the few traders with relevant knowledge. While markets tend to be less susceptible than elections or polls to this dilution effect, the market force to close arbitrage opportunities will not be sufficient to offset the dilution effect if the cap on the amount of money an investor can stake is too small.

Consider as an example the game show *Who Wants to Be a Millionaire?* Contestants on the program are allowed to poll the audience for help in choosing their answer to a trivia question. Imagine if the host of the show posed a math problem for which only three mathematicians knew the correct answer. If the contestant could question just these mathematicians, she would learn the appropriate response. However, if the contestant surveyed the entire audience—even if it contained these mathematicians—then the opinions of those who did not know the answer would overwhelm the few knowledgeable experts, and the aggregate results of the survey would be practically useless. Terrorist attacks that depend on the element of surprise are, by their very nature, similar to questions whose answers are only known by a select few.

However, if instead of polling the audience the contestant was offered an alternative lifeline—*create a market and watch the trading*—the result would be different. If one of the mathematicians was in the audience, she would continue to buy shares of the option that represented the correct answer until the price reflected her subjective confidence that she knew the answer (or until she hit the cap on trading). If the cap on trading were small, then the mathematician's knowledge would not produce a noticeable spike in prices. If the cap were large, however, the mathematician's knowledge would be reflected in the price of the commodity. Alternatively, with a small cap and more mathematicians, the market would also do better. This analogy suggests that for hard problems in which a very small percentage of the population has real information, ensuring that a large percentage of market traders are potentially informed or ensuring that spending caps are not too small are important issues to consider.

From this brief analysis of the incentives, we can draw several conclusions. First, a PAM limited to those with expert opinions would likely have done a better job anticipating future developments than one which invited broad public participation. Second, some form of entry requirements designed to exclude likely terrorists would minimize the manipulation concerns. Third, allowing traders to cast rich enough bids so that one confident person can influence the market will help reduce the dilution effect. Fourth, contracts on a small number of events may limit the market's exposure to problems that surfaced in the multi-party proportional representation field experiments.

Based on this analysis, the PAM proposed by DARPA—which would have had open participation; the potential for large numbers of contracts on very specific events; and, most likely, low caps on trading—clearly had serious flaws that would have inhibited its value as a predictive market. But by thinking

about how the structure/rules of the market might influence the behavior of traders and, ultimately, the market's effectiveness, institutional analysis yields several amendments to the proposal. With these changes, it seems possible that a PAM might provide the intelligence community with useful information. However, the value of such an amended PAM should not be assessed in a vacuum. Instead, we should contrast it with alternative means for filtering information.

Comparing the Market to an Alternative

Earlier, we touted the potential relevance of comparative institutional analysis. One of the striking features of the brief debate over the PAM was that this comparative dimension was almost completely absent. Since no one viewed the market as a replacement for the current intelligence community, the relevant question to consider is whether it is likely that PAM might offer information that would be missed by current intelligence gathering procedures.

Bruce Berkowitz and Allan Goodman present a lucid description of the status quo intelligence gathering institutions:

Despite popular conceptions, the defining character of the modern intelligence community is not . . . James Bond. It is Max Weber, the German sociologist best known as the inventor of the concept of bureaucracy. The intelligence community is a classic bureaucracy, characterized by central planning, routinized operations, and a hierarchical chain of command.³⁰

Most scholars would agree that hierarchies like those found in the intelligence community are capable of missing available information. Typically, information only makes its way up the chain if, at each level of the hierarchy, the supervisor chooses to pass it on. Accordingly, an unpopular speculation can easily be blocked by one of the many links in the chain. Moreover, even if the relevant information reaches the top of the hierarchy, the actual evaluations and policy prescriptions are made by a small set of individuals.

The recent Congressional inquiry into the 9/11 terrorist attacks provides an apt example of the potential shortcomings of such methods of information aggregation. The report notes that, as a whole, the intelligence community “did have information that was clearly relevant to the September 11 attacks, particularly when considered for its collective significance.”³¹ In particular, various parts of the intelligence community had received information suggesting that aircraft might be used as weapons, that the possibility of terrorist attacks in the United States was increasing, and that Osama bin Laden was sending students to the United States for aviation training. Despite these warning signs, the report concludes that “Some significant pieces of information in the vast stream of data being collected were overlooked, [and] some were not recognized as potentially significant at the time and therefore not disseminated. . . . For all these reasons, the [i]ntelligence [c]ommunity failed to fully capitalize on available, and potentially important, information.”³² By contrast, markets allow individuals the ability to act on relevant information unfettered by the constraints of whether someone higher up the chain of

command thinks the information is important. As *The Economist* noted when discussing the H.P. markets, “markets of this sort offer a much better way of gleaning valuable information. That is because traders have a strong incentive to get their predictions right (rather than say what they think their managers want to hear).”³³

As noted above, as long as caps are high enough, even one individual can significantly affect the price of a commodity. In the existing hierarchical world, agents at both FBI headquarters and the New York field office chose to ignore the Phoenix Electronic Communication warning of the “inordinate number of individuals in investigative interest” seeking flight training in Arizona³⁴ and, consequently, the information went nowhere. Had a market been in place, however, the authors of the Phoenix memo might have been able to affect the price of a related futures contract, something that would have been noticeable throughout the intelligence community, including to those agents who had been warning about the possibility of aircraft being used in terrorist attacks since 1994.

James Surowiecki³⁵ makes a similar comparison:

PAM might also have been effective because traders in a market have no incentive other than making the right prediction—that is, there are no bureaucratic or political factors influencing their decisions—so they eliminate many of the hurdles that limit the flow of information within organizations. That's especially important in the case of the intelligence community because we know that, for example, in the case of 9/11 there was lots of valuable and relevant information available before the attack took place. What was missing was a mechanism for aggregating that information in a single place.

Clearly, U.S. intelligence gathering capabilities in recent years have not been beyond reproach. Would a better designed PAM have helped supplement these capabilities? Due to the policy debate that emerged in July of 2003, we'll probably never know the answer to that question.

Conclusion

Our review of the press treatment of the proposed PAM leads to the conclusion that comparative institutional analysis had little influence in the decision process that killed the market. More troubling, even if a literature containing exemplary analyses of potential markets and detailed accounts of the extant field and lab experiments surfaced on July 29, it is likely that the political debate would have ignored the information. The program was not scrapped because intelligent people believed it would be ineffective. Rather, we suspect that PAM was cancelled because calling the program “offensive” and “morally wrong” was easy, while making the claim that prevention of future attacks might involve unsavory types of institutions and explaining how one weighs these issues was hard.

For all the moral outrage that surfaced in the press, it is hard to find a clear standard by which one concludes that PAM is *wrong* and current intelligent practices are *right*. Several authors, in fact, made the case that PAM is no more objectionable than current intelligence practices. It is true that under PAM some individuals will make money by pondering the unthinkable.

However, intelligence operatives are already paid to ponder the unthinkable and sometimes promoted and receive bonuses when they are right. Why the hierarchy mechanism for aggregating information is less objectionable than the market mechanism is unclear.

We have surely only scratched the surface of the issues that ought to be discussed in evaluating a proposal like PAM and, as such, we offer no definitive policy verdict. However, we hope that this brief analysis demonstrates that there are real issues regarding PAM that can be discussed in a systematic manner. By studying the ways that institutions structure the behavior of individuals (in this case, traders or intelligence community members), scholars can help inform the debate about institutional design and change.

Notes

- 1 Bowman 2003, A1.
- 2 *St. Louis Post Dispatch* 2003, C12.
- 3 Hudson 2003, A1.
- 4 Guggenheim 2003.
- 5 Gosselin 2003, E5.
- 6 We searched by the keywords *terror*, *futures*, and *market*, and eliminated any articles that were off topic.
- 7 *Pittsburgh Post-Gazette* 2003, A10.
- 8 *Santa Fe New Mexican* 2003, A7.
- 9 *The Augusta (Ga.) Chronicle* 2003, A4.
- 10 *St. Louis Post-Dispatch* 2003, C12.
- 11 One other editorial attacked PAM, but urged restraint in attacking DARPA more broadly.
- 12 *Portland Press Herald* 2003, 4C.
- 13 *The Virginian-Pilot* 2003, B10.
- 14 *San Mateo County Times* 2003, 9.
- 15 *The Virginian-Pilot* 2003, B10.
- 16 Day 2003.
- 17 Brown 2003.
- 18 Glenn 2003, E5.
- 19 Graham 2003, A1.
- 20 Ibid.
- 21 Hayek 1948.
- 22 Forsythe et al. 1992.
- 23 See *ibid.* for a review of the field experiments.
- 24 Jacobsen et al. 2000.
- 25 Chen and Plott 2002, 5.
- 26 Ibid.
- 27 Graham and Loeb 2003, A1.
- 28 *The Augusta (Ga.) Chronicle* 2003, A4. The editors, however, appeared dubious of the likely success of such a sting operation, noting that “sometimes one wonders what they use for brains in the government. Any bets on that?”
- 29 Dorgan 2003, A20.
- 30 Berkowitz and Goodman 2000, 67.
- 31 Senate Report No. 107-351, 2002, xi; House Report No. 107-792, 2002.
- 32 Ibid., xi.
- 33 *The Economist* 1999, 67–8.
- 34 Senate Report No. 107-351, 2002, xiii; House Report No. 107-792, 2002.
- 35 Surowiecki 2003.

References

- The Augusta (Ga.) Chronicle*. 2003. A brainless scheme. 4 August, A4.
- Berkowitz, Bruce D., and Allan E. Goodman. 2000. *Best Truth: Intelligence in the Information Age*. New Haven: Yale University Press.
- Bowman, Tom. 2003. Want to bet on terrorism? Pentagon wants to let you. *The Baltimore Sun*, 29 July, A1.
- Brown, Jeff. 2003. Lack of intelligence true demise of terrorism futures. *The Houston Chronicle*, 3 August, 1.
- Chen, Kay-Yut, and Charles R. Plott. 2002. Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Social Science Working Paper 1131, California Institute of Technology.
- Day, Bill. 2003. Terror market is still a bad idea. *The San Antonio Express News*, 16 August, 2H.
- Dorgan, Byron L. 2003. The Pentagon's ill-conceived market. *The Washington Post*, 7 August, A20.
- The Economist*. 1999. The heyday of the auction. 24 July, 67–8.
- Forsythe, Robert, Forrest Nelson, George R. Neumann, and Jack Wright. 1992. Anatomy of an experimental political stock market. *The American Economic Review* 82:5, 1, 142–61.
- Glenn, Joshua. 2003. Betting on trouble. *The Boston Globe*, 3 August, E5.
- Gosselin, Peter G. 2003. Plan for terrorism futures market killed. *Los Angeles Times*, 30 July, 11.
- Graham, Bradley. 2003. Poindexter to leave Pentagon research job. *The Washington Post*, 1 August, A1.
- Graham, Bradley, and Vernon Loeb. 2003. All bets off as Pentagon scraps “invest in terror” plan. *The Record (Bergen County, N.J.)*, 30 July, A1.
- Guggenheim, Ken. 2003. Warner says Pentagon threat-bet program to be canceled. *The Associated Press*. 29 July.
- Hayek, Friedrich A. 1948. *Individualism and Economic Order*. Chicago: University of Chicago Press.
- Hudson, Audrey. 2003. Hill hits gambling program on terror. *The Washington Times*, 29 July, A1.
- Jacobsen, Ben, Jan Potters, Arthur Schram, Frans van Winden, and Jorgen Wit. 2000. (In)accuracy of a European political stock market: The influence of common value structures. *European Economic Review* 44:2, 205–30.
- The New York Post*. 2003. All bets are off. 31 July, 28.
- Pittsburgh Post-Gazette*. 2003. Futures shock: The Pentagon's sick take on risk strategy. 2 August, A10.
- Portland Press Herald*. 2003. Weigh security with privacy rights carefully. 3 August, 4C.
- San Mateo County Times*. 2003. Futures market, Poindexter Pentagon's bad bets. 1 August.
- Santa Fe New Mexican*. 2003. Dr. Strangelove caught in bizarre betting parlor. 31 July, A7.

- St. Louis Post-Dispatch*. 2003. "Unbelievably stupid." 30 July, C12.
- Surowiecki, James. 2003. Damn the slam PAM plan! *Slate Online*. 30 July. Available at slate.msn.com/id/2086427. Accessed 5 January 2004.
- U.S. Congress, Senate, Select Committee on Intelligence, and U.S. Congress, House, Permanent Select Committee on Intelligence. 2002. Joint Inquiry into Intelligence Community Activities Before and After the Terrorist Attacks of September 11, 2001. 107th Congress, 2d Session, Senate Report 107-351, House Report 107-792.
- The Virginian Pilot*. 2003. Pentagon drowns in its own "dead pool." 31 July, B10.
- Wolfers, Justin, and Eric Zitzewitz. 2003. The furor over "terrorism futures." *The Washington Post*, 31 July, A19.