To appear in Alan S. Gerber and Eric M. Patashnik, eds., *Promoting the General Welfare: New Perspectives on Government Performance,* 151-173, Brookings Institution Press, Nov. 2006.

Chapter 7

Decision Markets for Policy Advice

Robin Hanson

The main cause of bad policy decisions is arguably a lack of information. Decisionmakers often do not make use of relevant information about the consequences of the policies they choose. The problem, however, is not simply that public officials do not exploit readily available information. It is also that they do not take full advantage of creative mechanisms that could expand the *supply* of policy-relevant information. Among the most innovative and potentially useful information-generating mechanisms are speculative markets. Speculative markets produce public information about the perceived likelihood of future events as a natural byproduct of voluntary exchange.

Speculative markets do a remarkable job of aggregating information; in every head-to-head field comparison made so far, their forecasts have been at least as accurate as those of competing institutions, such as official government estimates. Many organizations are now trying to take advantage of this effect, experimenting with the creation of "prediction markets" or "information markets," to forecast future events such as product sales and project completion dates.

This chapter examines the uses and limitations of decision markets. Decision markets are information markets designed to inform a particular policy decision, by directly estimating relevant consequences of that decision. After reviewing the weaknesses of existing institutions, the mechanics of decision markets, and a concrete example, this chapter reviews the requirements, advantages, and disadvantages of decision markets. The chapter also takes a close look at a particular application of this tool—the controversial yet illuminating attempt to establish a "Policy Analysis Market" to forecast the consequences of major policy U.S. choices in the Middle East.

Government Information Failures

There are many ways that governments can fail to achieve good policy outcomes. Some interest groups may have too much power or want the wrong things, or leaders may pursue personal goals that do not have broad public support. Government agencies may also pursue parochial goals or fail to coordinate with each other.

Most of the government failures that we might overcome, however, are arguably due to a lack of information about the consequences of policy choices. Policymakers may have mistaken beliefs about the direct effects of a policy. They may fail to consider an indirect effect or an alternative policy. Or we might fail at a higher level to choose the right institutions to set the context for such judgments and choices. In fact, much of what economists have been doing for the past half century is showing how most social problems are at root caused by such information failures. For example, there is almost no end to the troubles caused by the principal-agent problem—that is, the fact that experts know more than their clients.[1] We rely on doctors, lawyers, teachers, and chief executive officers because they can know more than we do. Government policy similarly relies heavily on expert knowledge, expressed directly via civil servants, contractors, and official boards and commissions, and indirectly via mass media and academic reports and analyses.

We hope that such experts acquire and use their knowledge for our benefit, but we fear that they instead give us the advice that benefits them. Stereo salesmen, for example, may suggest overly expensive stereos, and surgeons may recommend surgery over drug treatments. We cannot trust such experts because of “professional ethics"; if they are trustworthy, it is because their institutional and contractual incentives make them so.

Unfortunately, agency problems in the government-citizen relationship are among the worst, because the chains of delegation are so long, the topic coverage is so broad, and citizens have so little reason to pay attention. Each election, voters, with little expertise or chance of personally influencing the outcome, form judgments about how the policy outcomes they favor relate to the candidates before them. Representatives, anticipating those judgments and any other voter preferences, oversee the agencies that

implement government policy. All the way down the long chain of delegation, each official tries somewhat to make his part of government look like what the officials above him want it to look like.

To be sure, voters can try to rely on nongovernmental advisors, such as mass media and organized political groups, but this merely replaces the problem of evaluating candidates with a similar problem of evaluating these advisors, and makes the delegation chain even longer. Those who run media firms and organized political groups are professionals, just like politicians and CEOs. They compete with others who could replace them and so mostly do whatever it takes to build their power, prestige, and profit.

Sometimes bad things happen that voters notice that can be attributed to specific actors whom voters can remember to punish in a future election. Sometimes good things happen that voters can similarly notice, identify, and remember to reward. Where voters are willing to act on such clues, government officials have an incentive to avoid these bad things and encourage these good things. Ideally, government officials could be well disciplined if a single entity ran the government over long terms and if each voter only considered whether his life had gone better or worse than expected since the last election.[2]

In fact, however, terms are short, government is divided, and voters do not just consider their own lives.[3] Voters weigh policy effects on distant others and try to distinguish policy influences, and, although these strategies may have some advantages, they can make voter information problems far worse. For example, some say the U.S. Food and Drug Administration (FDA) tries harder to avoid bad new drugs than delays in good drugs, because the former are far more noticeable.[4]

Most important, voters have many opinions about the effectiveness of policies and process, and the government naturally caters to those opinions. Policy often seems closer to what public opinion would suggest than to what relevant experts advise.[5] Even when governments apparently rely on expert opinion, they often do so to legitimate predetermined policies, rather than to gain information to help determine policy.[6]

A candidate who cares only about winning the next election should neglect what he knows about good policy and prefer the policies that voters favor.[7] Of course, this is not just a problem with government. A CEO should neglect what he knows about good long-term corporate strategy in favor of whatever will impress stock speculators when that CEO exercises his stock options.[8] This phenomenon is very general. Experts should not even bother to obtain information that their clients will not know when those clients reward the experts.[9]

This problem can, however, be much worse in government. Whereas stock speculators can have a strong personal incentive to think carefully about corporate policy, voters not only have little reason to think carefully about how to vote, but often have positive reasons to be irrational.[10] Yes, voters may need to think about how to deal with the policies that win and may want to sound knowledgeable (and compassionate and loyal and much more) when discussing politics with associates, but these are far weaker than the incentives of serious stock speculators.

What can we do about this policy information problem? We could subsidize voter education, but education is expensive, only somewhat relevant, and it is hard to keep such subsidies politically neutral. We could seek to limit the franchise to the most politically informed, but there is little political support for insulting less-informed citizens in this way.

Another approach would be to educate voters to trust and appreciate the benefits of a new policy process. Voters now accept the results of most jury trials not because they have opinions on whether the accused was guilty, but because they believe the process is likely to convict the guilty. Voters also accept most research funding allocations not because they have opinions about which grant proposals should be funded, but because they accept the peer review process in research. The public similarly accepts the results of elections, FDA drug approvals, and exams for drivers' licenses, civil service promotion, and college admissions.

All policy processes, however, do not inspire equal confidence. For example, official committee reports are routinely suspected of being biased via the selection of committee members, whereas cost-benefit (and other economic) analyses are often suspected of being biased via method choices. Statistical

analyses fare a bit better, but are still often suspect. These cases suggest that voters tend to trust simple informative mechanisms that seem difficult for interested parties to bias.

How Decision Markets Work

Decision markets are a new policy mechanism that we might hope will one day be accepted by the public as especially simple, informative, and difficult to bias. Decision markets are speculative markets that estimate policy consequences.

Speculation takes place when an actor buys or sells something today in the hope of reversing her trade later for a profit. For example, if an actor thinks housing prices are climbing, he might buy a house today in the hope of selling it next year at a higher price. But, because it costs so much to buy and sell houses, the actor could lose money even if housing prices do rise. When durable commodities are traded in markets with low transaction costs, however, speculation becomes cheap. These markets are known as *speculative markets*. For example, if you are confident that the price of gold is going up, you are pretty sure to profit from buying gold today in order to sell it tomorrow. And if you are confident that the price of gold is going down, you can profit by selling today in order to buy tomorrow. (If you do not have gold today, you can still sell today by "selling short.")

Speculative markets ultimately do an excellent job of aggregating relevant information. This is because anyone who thinks they have information about whether the price is going to go up or down expects to profit by speculating on that information, and, by speculating on information, one incorporates that information into the market price. For example, if you think the price is going up, you buy, which tends to raise the price. Speculators thus compete to be the first to profit from any new information relevant to predicting future prices. The net result is that, when speculative market prices exist, it is hard to find information that such market prices do not embody. After all, anyone who finds such neglected information can profit by trading on it and thereby eliminating this neglect.[11]

Although the information-collecting abilities of speculative markets are impressive, such markets are usually created for other reasons, such as to allow traders to entertain or insure themselves. One can,

however, create new markets specifically to try to take advantage of this information-collecting effect. For example, if one wanted to know whether it would rain here tomorrow (at place *X* and on date *T*), one might create contingent assets that declared "Pays $100 if rain at place *X* on date *T*." If one could get speculators to trade such assets, and, if their current market price was $20, one could take this as an estimate that the probability of rain tomorrow is 20%. After all, a person unconcerned about risk would be buying or selling this asset if they thought the probability of rain was anything other than 20 percent.[12]

Recently some new markets have been created specifically to take advantage of these effects. Called *prediction markets*, information markets, virtual stock markets, artificial markets, or idea futures, these markets are now used to estimate measures such as product sales, project completion dates, or election outcomes.[13] Sponsors have created both public markets, where anyone can trade and learn from the market prices, and private markets, where only a select few can trade or see prices. Sponsors have created both markets where people trade hard cash and markets where they trade "play money" for bragging rights.

Such markets (at least the hard cash versions) have so far done well in every known head-to-head field comparison with other social institutions that forecast. Orange juice futures improve on National Weather Service forecasts, horse race markets beat horse race experts, Academy Award markets beat columnist forecasts, gas demand markets beat gas demand experts, stock markets beat the official NASA panel at fingering the guilty company in the Challenger accident, election markets beat national opinion polls, and corporate sales markets beat official corporate forecasts.[14]

Although more accurate market-based forecasts can be of great value, it seems that market-based policy-conditional forecasts might be of even greater value.[15] Let us now focus on this *decision market* option.

A Concrete Example

Consider the oft-heard proposal to change U.S. health insurance into a single-payer system (SPS), such as many other developed nations use. Many claims have been made in support of such a proposal, but

certainly two of the most important claims are that a SPS would improve health and reduce costs. Although media and academic analyses and reports have weighed in on this subject for many decades, the public remains uncertain about whether a SPS would in fact have these consequences.

Let us focus on the health claim, and in particular on whether average lifespan would be increased ten years later if the United States adopted a SPS, relative to if a SPS were not adopted. More specifically, imagine the year is 2008, and people disagree about the consequences of a particular bill before the U.S. Congress, which would implement a SPS starting in 2010.

A decision market approach to this problem would create two speculative markets and then observe two prices. (We explain exactly how below.) One price would represent an estimate of U.S. lifespan in 2020 given that the SPS bill passed, while the other price would represent an estimate of U.S. lifespan in 2020 given that the SPS bill did not pass. The difference between these two estimates would say how much market speculators expected this SPS bill to raise or lower U.S. lifespan.

Once the public accepted such market estimates as good policy estimates, the main remaining tasks for the public, media, and lobbyists would be to monitor the consistency between prices and policy; the policies chosen should seem reasonable given market-specified beliefs.[16] For example, a single-payer health insurance system should not be approved if markets estimate that it would both raise costs and reduce lifespan, at least when these are the main arguments for such a change. Alternatively, we might approve new institutions that more directly put market estimates in charge of decisions.[17]

So how would this work exactly? To get decision market estimates on the SPS question we would first need to define our measures well enough to settle bets. So we might define U.S. average lifespan in 2020 using the official statistics of the U.S. National Vital Statistics System in 2022. We might define SPS adoption to mean that the particular SPS bill, however amended, was signed into U.S. law during 2009. And we might define the assets being bet to be U.S. currency paid in 2022, corrected for officially measured inflation from 2009. Judges would have to be appointed to decide about any remaining unanticipated ambiguities.

Once these measures were defined, one or more financial institutions could create assets to represent the relevant bets. One asset would be "Pays $100 if *S*," which in 2020 pays 100 inflation-adjusted U.S. dollars if *S*—that is, if the SPS bill is passed (and nothing otherwise). A second asset would be "Pays $100 if not *S*." A financial institution would take no net risk by allowing $100 (of inflation-adjusted U.S. dollars in 2020) to be exchanged at anytime for the pair "Pays $100 if *S*" and "Pays $100 if not *S*."

There would also be no net risk in allowing $100 to be exchanged for the pair "Pays $*L*" and "Pays $(100 – *L*)," where *L* is U.S. average lifespan in 2020 (or 100 in the unlikely case that lifespan were in excess of 100). There would further be no risk in allowing "Pays $*L*" to be exchanged for the pair "Pays $*L* if *S*" and "Pays $*L* if not *S*."

Once such assets were available, we could create the two decision markets to trade them. There are many possible market structures. For example, the standard double auction structure supports a single kind of standard trade, such as gold for cash. In the double auction, anyone can at anytime post an offer to buy or sell at their chosen price, cancel any of their previous offers, or accept one of the best current offers.

Another option is a simple automated market maker. Such a market maker always has an offer to buy a standard quantity at some price and a similar offer to sell at some higher price. The market maker sets these prices using a monotonic function of the assets it holds; whenever someone accepts the market maker's sell offer it raises both prices, and whenever someone accepts its buy offer the market maker lowers such prices. As information is revealed, the market maker should lose on average, but a strict bound can be placed on how much it can ever lose, and the gains are given only to those who move the market price toward its ultimate resolution.[18]

Another market structure is a call auction, where offers collect to be matched together at the same price. Call auctions have lower price noise, at the cost of a delay in getting prices. Combinatorial versions of all these structures can support the combined trading of many different related assets. Markets

structures also include mechanisms for ensuring that traders can and do make good on their offers, and they include choices about who can see what offers and trades.

Whatever the details of the exchange mechanisms, the important thing would be the appearance of two *market prices*, or trading rates of exchange. In the first decision market we need, people would exchange one asset "Pays $*L* if *S*" for some chosen fraction of the asset "Pays $100 if *S*." In the second market people would exchange one "Pays $*L* if not *S*" for some fraction of "Pays $100 if not *S*." These fractions could be taken as market estimates of the conditional expected values E[*L* | *S*] and E[*L* | not *S*], respectively. If the first fraction were higher than the second, market speculators would be saying that they expect U.S. lifespan to be higher if the SPS bill is approved than if not. Such estimates are what we would want to advice SPS policy.

Requirements for Decision Markets

We have good reasons (including formal theory, lab experiments and field data) to think that decision markets can give accurate estimates. But decision markets are not cure-alls; they require:

Important-enough Claims

There are fixed costs of using this decision market process. You have to set up markets, get traders to pay attention, and so on. Yet, if other forecasting institutions are working well, market estimates might only be slightly more accurate than estimates from other institutions. The policy question asked thus needs to be important enough for this added accuracy to be worth the added fixed costs.

Careful thought should be given to what question one wants to pose to market speculators. If you ask the wrong question, the answer you get may not be very useful. For example, if reducing lifespan inequality is as important a benefit of SPS as increasing average lifespan, one might want to ask markets about inequality as well. If one cannot afford to ask about both, one might ask about a weighted average of the two.

Enough Influence

Prices in real markets can contain noise, because of random influences on trader incentives and behavior. So the decision options considered need to have a large enough influence on the outcomes measured for random price fluctuations not to obscure the relevant outcome differences. Decision markets will probably not show how a single act of capital punishment affects the murder rate, for example. Averaging prices over a short time may help, and simple statistical tests can suggest whether a price difference is statistically significant.

Distinct Options

There must be a way after the fact to determine whether a particular decision option was chosen or not. One could not ask markets about the consequences of "progressive reform" or "using quality standards" unless one could define relatively clearly what would count as such a policy. As with the legal system, ambiguities might be dealt with using defaults created by precedents and by using judges who are trusted to be neutral.

Measured Outcomes

There must be a way to create measures of some of the important outcomes of interest after the fact and to identify what those measures will be before the fact. Large measurement errors are not much of a problem, nor are other influences on the measured outcomes. Instead the problems are correlations between errors and the decision and correlations between errors and efforts to influence the measurement process.

One does *not* need to be able to tell after the fact how much each decision influenced the outcome. Standard decision theory is clear that, when you are faced with some options A, and you want to achieve the outcome B, you should choose the option A associated with the maximum value of your expectation E[B | A]. These summary expectations are all that should matter, not the details that produce them.

Decision-insider Traders

Standard decision theory is also clear, however, that the relevant expectations E[B | A] must be calculated relative to the beliefs of the agent who will actually make the decision. Thus, if the people who

will actually make the decision have different beliefs from market speculators, a *decision selection* problem can arise.[19] If speculators think that, when the decision is made, decisionmakers will know something important that speculators do not now know, speculators will use the actual decision as a clue about this important hidden information. The market price will then reflect speculators' best guess about how the world would be *given* that decision, rather than their best guess about how the world will be *because* of that decision.

For example, speculators might estimate that on the whole SPS would reduce lifespan, but they may also foresee a small chance that the decisionmakers will have surprising information favoring SPS. If this scenario is the main way that the decisionmakers might approve SPS, speculators should estimate that conditional on seeing SPS be approved, SPS would raise lifespan. To avoid this decision selection problem, persons with access to decisionmaker information should be permitted to trade in these markets.[20] Also, the timing of the key decisions should be clearly announced just before such decisions are made so that speculators trading then need not fear the decision will be based on future information.

Enough Informed Traders

Decision markets aggregate information available to their participants, so no minimum number of participants is needed, nor is there a minimum level of participant information. With a subsidized market maker, even a single trader can profit by making the market price express what little he knows. At the other extreme, by inducing more traders to put in more effort, there is often no obvious limit on how accurate prices might become.

Thus, it only makes sense to ask if there are "enough" informed traders relative to some standard of comparison. If anything below a given accuracy is useless, you need traders with enough information to produce that accuracy. If decision markets are being compared with a competing institution, some of the people with the sort of information available to that competing institution must be allowed to trade in the decision markets.

Although the accuracy of some forecasting institutions, such as polls, depends on the information level of the average participant, the accuracy of speculative markets depends more on the most informed

participants; the typical person able to trade need not be informed or even rational. Instead, a small group whose members know they are better informed can dominate prices, as long as they have enough resources and if market rules do not artificially limit how much they can trade. Those who know they are less informed tend to back away, and those who do not know that they do not know tend to lose and then back away.

Enough Trader Incentives

Even when people with access to relevant information are able to trade, such people need a reason to bother to trade, relative to the other ways they could spend their time and money. They might trade for insurance, for fun and sport, or because of an irrational overconfidence in their opinions; however, when people expect few others to trade in a market, they may all stay away, and so a market price may fail to exist.

Fortunately, one can ensure market prices exist either via a large enough financial subsidy to fund an automated market maker, or via a large enough chance market prices would influence important decisions. The larger the subsidy or influenced decision, the more incentives traders acquire, and the more accurate prices should be. If individuals vary how much they enjoy trading, or in their cost of trading or acquiring relevant information, prices should also be more accurate when more people are allowed to participate.

There is also the question of whether the assets risked are ordinary highly convertible "real money" assets or "play-money" assets of limited convertibility, where winners mainly win bragging rights. Most of the successes of speculative markets have been for open unlimited mostly anonymous trading of highly convertible assets. But legal barriers are much lower for play-money, and some research has found play-money markets can do as well real-money markets, at least on topics where many people already acquire relevant information for other reasons.[21]

Trader Anonymity

Even when the benefits from trading are large enough to overcome the costs of bothering, threats of retribution could deter participation. Leaders of a government agency, for example, might retaliate

against employees who contradicted their official party line in decision markets. Such retaliation can be prevented, however, if people can trade and gain their trading rewards anonymously.

Aggregate-enough Outcomes

Even if prices are accurate reflections of reality, would that accuracy come at the cost of people changing reality, by hurting others to win bets? Such acts of harm might come if people were betting on events small enough for individuals or small groups to influence. Fortunately, it is very difficult for individuals or small groups to have much influence over typical policy aggregates, such as average U.S. lifespan. Furthermore, we do not really see much in the way of attempts to harm in order to profit from trades on the stocks of amusement parks, hospitals, or other companies that individuals or small groups might plausibly effect.[22]

It seems that the sort of people who are willing and able to cause harm do not have access to the capital required to implement such trading strategies; they usually prefer simple extortion. And decision markets should cause far less temptation because they involve far less money than ordinary financial markets. Thus, decision markets should be safe as long as they stay away from estimating parameters that are too easily influenced by individual criminals.[23]

Linear-enough Outcomes

If policymakers look to decision market prices as a guide to policy, others may be tempted to manipulate those prices in order to manipulate policy. Fortunately, the addition of manipulators should *increase* price accuracy. Manipulators are in essence noise traders, because their trades are not correlated with asset value information, and markets with more noise traders generally have more accurate prices, because more informed traders are attracted to profit from the noise traders.[24] This predicted inability of manipulators to hurt price accuracy has been confirmed by lab experiments and in the field.[25]

It is theoretically possible, however, for manipulation to increase the harm from price errors even as it increases price accuracy. For example, a market estimating the probability of some disaster might better estimate the probability of disaster, but do worse at catching the worse disasters. A simple way to avoid this possibility is to trade assets whose payoffs are relatively linear in the relevant harm.

Conditional-enough Outcomes

One important aspect of asking the right question is avoiding self-fulfilling prophecies. For example, imagine that the market estimated that a new ambitious government computer system would be delivered late, and in response the contract was renegotiated to make the system less ambitious so it could be delivered on time. This would punish the market speculators for giving their valuable forecast. A better approach is to ask about outcomes conditional on the relevant policy responses. In this example, one could ask the market about the system delivery date conditional on not changing the contract and then conditional on some change in the contract. This approach would reward speculators for giving good advice.

Intermediate-enough Estimates

Speculative markets sometimes seem to show "long-shot" biases when estimates are extreme. For example, when estimating the probability of an unlikely event, those who want to raise the probability of that event have to spend a lot less in each trade than people who want to lower the probability. This can result in a bias toward less-extreme prices, at least when transaction costs are high enough to discourage attempts by other speculators to profit by correcting such biases.

Low transaction costs can reduce this problem, as can allowing combinatorial trading. With combinatorial trading, high probability outcomes can be broken down into lots of low-probability components, so all components traded have a similar probability.

Can Show Prices

It is more difficult (though hardly impossible) for traders to inform prices when they cannot observe recent market prices, but such prices reveal to observers something close to our best estimates at that point in time of the consequences of our choices. If we do not want potential "adversaries" to know these estimates, we may want to limit who is allowed to trade or how recent are the prices they can observe. Such limitations are possible, but probably cost something in the way of price accuracy.

Legal Permission

The cost of following the legal process to gain permission to open a few markets to evaluate a policy must be smaller than the benefits that an interest group might gain from advocating that policy. Otherwise no one will bother to set up markets in the hopes of convincing the rest of us of the value of a proposal. Unfortunately, legal fixed costs are far too high for this condition to hold today, at least for real money markets.

Public Credibility

The policy audience must perceive decision market estimates to be relatively accurate and difficult for interested parties to bias. In a democratic government, this means the public must accept such market processes similar to they way they now accept jury trials as a reliable neutral forum for policy debate. This condition does not yet hold today.

Advantages of Decision Markets

Decision markets can directly advise our important policy decisions, by giving us more accurate estimates of the aggregate consequences of those decisions. Decision market price estimates should be numerical, precise, respond quickly to new information, be self-consistent across a wide range of issues, and be at least as accurate as other publicly available estimates. They are also relatively democratic in the sense of allowing anyone to participate and yet are relatively difficult for interested parties to bias.

Estimates based on cost-benefit or statistical analyses are often criticized as out being of date, being based on bad modeling assumptions or data sets, misjudging causality directions, or failing to control for relevant parameters. Market estimates are mostly immune to these criticisms. Although the analyses that individual traders use to choose their trades may well have these problems, it is up to traders to make their best judgments to fix them. Anyone who could identify a bias in any aspect of the current market estimates would in effect be offered a financial reward to correct that bias.

The cost of creating decision markets does not vary greatly from topic to topic. Most of the variation is probably based on whether information is spread out among many people and on how many people know any given relevant piece of information. The benefit of creating decision markets varies

enormously, however, with the importance of the decision being advised. Thus, the most benefit relative to cost will come from decision markets advising the largest most important decisions, where existing forecasting institutions have serious failings.

Early decision markets attempts, however, tend naturally to be efforts to validate the technology. For validation efforts it makes more sense to look for policy decision contexts in which a status quo institution gives concrete enough forecasts to allow a numerical comparison, where that status quo institution is suspected of serious failings, and where many similar decisions have consequences that will be known within a few years. This is where decision markets have their best chance of clearly demonstrating a superior forecasting accuracy.

Disadvantages of Decision Markets

For now, decision markets have two overwhelming disadvantages. First, decision markets have not yet gained enough credibility in the public mind to be an attractive forum in which to argue for particular policy choices. And second, there are enormous legal barriers preventing the creation of public markets trading hard currency.

If these problems can be overcome, other disadvantages will remain. For example, the public may not always want more neutral and accurate forecasts of the consequences of policy decisions. The public may have self-flattering beliefs, ideologies, and entrenched opinions about policy consequences, and they may not appreciate being contradicted by decision markets; sometimes people do prefer to shoot the messenger. Also, the public may have nonoutcome preferences over the processes that produce their policy estimates and over the people who influence it. Accuracy may not be the only issue the public cares about.

Even when the public truly *does* want more accurate forecasts about policy consequences, more disadvantages will remain. Decision markets have nontrivial costs, such as to carefully define relevant options and outcomes and to induce informed traders to participate. And decision markets have nontrivial

limitations, such as needing distinct anticipated options which have enough influence over measured aggregate-enough, linear-enough, and conditional-enough outcomes.

Policy Analysis Market: A Case Study

An interesting and well-publicized decision market was the now-defunct. Policy Analysis Market, also known as “terrorism futures” or “terror betting.”[26] In 2000, Michael Foster, who ran 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 information 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 first winner was Neoteric Technologies, subcontracting to Martek and professors at the University of Iowa. The second winner was Net Exchange, founded by a Caltech professor and subcontracting to two George Mason University professors, and later also to the Economist Intelligence Unit. The Net Exchange project was later named the "Policy Analysis Market" (PAM).

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

teams—at Penn State, Metron, Institute for Counter-Terrorism, George Mason University, Sparta, and BBN Technologies —were awarded $100,000 each.

Neotek sponsored an end of phase I conference in June 2002, and showed a few demonstration markets, using their preexisting 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 two-thirds of its phase II funding, and the new small projects had spent little of their funding. Michael Foster had asked for, but not received, $8,000,000 more in FutureMAP funding 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. Also from the start, they planned to forecast military and political instability around the world, how U.S. policies would affect such instability, and how such instability would influence U.S. and global aggregates of interest.[27] The reasoning behind this choice was that the cost to create markets does not depend much on the topic, but the value of estimates varies enormously with the topic. Thus, the greatest benefit relative to cost would come from the highest value estimates. And what could be more valuable than to inform the largest defense policy decisions?

Charles Polk, Net Exchange president, named this the Policy Analysis Market (PAM). The 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, U.S. military activity, and U.S. financial involvement. In addition traders would predict U.S. GDP, world trade, U.S. military casualties, and western terrorist casualties, and a few to-be-determined miscellaneous items. (The miscellaneous items were to be decided at the last minute based on suggestions from traders, with approval by the Economist Intelligence Unit.) This would add up to a few hundred base parameters to estimate.

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

For many years before PAM, Net Exchange had specialized in combinatorial markets, where buyers and sellers could exchange complex packages of items. So from the start, the plan was to see how far they could go in developing combinatorial prediction markets. In phase I, Net Exchange put together a combinatorial market similar to their previous combinatorial markets, and at the end of phase I they ran a complex simulation in which a dozen students traded over a few days for real money. Unfortunately, only about a dozen trades occurred, a serious failure.

In the interim phase, the Net Exchange team prepared for and ran lab experiments comparing two new combinatorial trading mechanisms with each other and with a traditional mechanism. These experiments had three traders set seven independent prices in three minutes, and then had six traders set 255 independent prices in three minutes. They found that a combinatorial market maker was the most accurate.[28] Phase II was mostly being spent implementing a scalable production version of this market maker. Because this mechanism requires a net subsidy to traders, they had budgeted $50,000 for this subsidy, and individual bets were to be limited to a few tens of dollars.

The PAM team was concerned that they might not attract enough traders in the final phase to achieve a meaningful test. Team members had considered running markets within government agencies, but found strong legal barriers to conditional transfers of money between agencies. In the absence of a single agency strongly interested in collaborating with them, they choose to create public markets.

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 their Web page and to give talks to drum up interest. Charles Polk created the PAM Web site, wherein as a backdrop to bold text, there were faint background sample screens. In a small (less than 2 percent)

section of two such screens, he included as colorful examples of possible miscellaneous items 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 support, because of privacy concerns with another IAO project, "Total Information Awareness." Because of this funding uncertainty, the PAM plans then were to start on September 1 with 100 testers, who would each be given $100 to trade. Online registration of people interested in being testers was to open August 1, and public trading of as many as 1,000 initial traders was to begin January 1, 2004. The fact that the PAM team was acting as an agent of the Pentagon was going to give them legal protection against violating antigambling laws.

On July 28, 2003, Senators Ron Wyden (Democrat, Oregon) and Byron Dorgan (Democrat, North Dakota) released an open letter to John Poindexter in which they complained that the U.S. Department of Defense was planning a "terror market" for people to bet on terrorist attacks.[29] They emphasized PAM's association with John Poindexter, who actually had little involvement with PAM, and described PAM as being "designed to predict terrorist events," when in fact it was focused on geopolitical trends. Their main evidence was those miscellaneous items in the PAM website background screen.

Wyden and Dorgan mainly complained that "terrorists themselves could drive up the market for an event they are planning and profit from an attack, or even make false bets to mislead intelligence authorities." Yet a few tens of dollars would hardly pay for an attack, and the PAM team had already told DARPA about their lab experiments showing manipulators do not hurt price accuracy.

A media storm immediately ensued. Although five mostly positive media articles mentioning PAM had appeared in the previous few months, fifty mostly negative articles appeared on July 29. DARPA public relations person was out of town and unreachable when the story broke. So DARPA was silent, and initial media reports were based mostly on the senators' complaint.

On July 29, Deputy Secretary of Defense Paul Wolfowitz told the Senate Foreign Relations Committee that he first heard of PAM reading the newspaper a few hours before and that the project was being terminated.[30] During that crucial previous day, no one from the government asked the PAM team

whether the accusations were correct or whether the more offensive aspects could be cut from the project. Many slow wheels of government had begun to turn. For example, the president's Council of Economic Advisors was considering coming out in favor of PAM, but the political decision to abandon PAM was made quickly, at a high level, and based on little information or lower-level input.

On July 30, seventy-eight media articles on PAM appeared, even more negative. That day, Poindexter reportedly resigned, and two months later all IAO research was ended. Over the following days, weeks, months, and years, more than 600 more media articles have mentioned PAM, many at first, and then gradually fading in frequency.

An analysis of over 500 articles on PAM found that coverage gradually became more positive, and the most recent fifty articles on average give readers a positive impression of PAM. One can collect eleven indicators of how informative were the articles about PAM. These indicators include publication date, citing someone with firsthand knowledge, article length, a news or an editorial style, author anonymity, and the awards, circulation, frequency, and topic specialties of the periodical. All eleven indicators individually predict that more informed articles give readers a more favorable impression of PAM. In a multiple regression model using additional six control variables, including media types, political leaning, and author gender, all six of the clearly significant information indicators predict that more informed articles favor PAM more. The more informed articles were more favorable, and eventually the average article was favorable, but the political decision to cancel PAM seems unlikely to be reversed anytime soon.

Psychological research by Philip Tetlock and others on "taboo tradeoffs" can help us to make sense of this unusual political event.[31] Tetlock and others study how people react upon learning that they, or others, have crossed a moral boundary and traded the sacred for the secular. They find that people are not only outraged at anyone who would cross such a moral boundary, but are almost as outraged at someone who would even think about crossing such a boundary.

For example, people might be given a description of a hospital administrator who is considering cutting costs at the possible risk of patient lives. Such people are then outraged at an administrator who

chooses to cut costs, regardless of the financial situation of the hospital or the ratio of money saved to lives risked. Furthermore, they are almost as outraged at an administrator who takes several days to think about the decision, even if he ends up not cutting costs. Apparently someone who would even think of doing such a thing is considered nearly as morally bankrupt as someone who actually does it. Bush had long defended Poindexter against attacks on his "Total Information Awareness" project, widely criticized as an attempt to collect and integrate databases on the public, but defending PAM seemed to have been beyond the pale.

PAM seems to have been accused of crossing a moral boundary, which can be paraphrased roughly as "none of us should intend to benefit when some of them hurt some of us." (Yes, many of us do in fact benefit from terrorist attacks; but we can plausibly argue that we did not intend to do so.) So, by the taboo tradeoff effect, it was morally unacceptable for anyone in Congress or the administration to take a few days to think about the accusation. The moral calculus required an immediate response.

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

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

Conclusion

Problems of acquiring and aggregating information on policy consequences lie at the core of most political failures, and these problems are often quite severe. Decision markets are a new policy information process intended to help overcome these problems. If our experience with similar markets is any guide, decision markets would be at least as accurate as any coexisting policy information process. Such markets are also precise, consistent, responsive, difficult to bias, and equalitarian in the sense of allowing broad participation.

Decision markets are not cure-alls, of course. They will not function effectively without the ability to create distinct options that have a measurable influence on relevant aggregate outcomes, and the problem has to be important enough to be worth paying to give traders, some of them insiders, enough incentive to trade. The legal costs of market creation also must be low enough, and the public must see market estimates as credible, accurate, and neutral.

It is these last two conditions that remain the biggest problem for decision markets. Without lower legal costs we will not see many trials, and without successful trials the approach cannot gain enough public credibility. Yet without that public credibility there is not enough political support to pass a bill to lower those legal costs to allow more trials.

The history of the aborted Policy Analysis Market shows that, absent such credibility, decision markets can be misrepresented and turned into political poison. This vulnerability of a market mechanism to such an attack should not be too surprising. After all, most familiar financial products, including stocks, insurance, futures, and options markets, were once prohibited by laws against gambling. It took a long time for the relevant industries to convince the public to see each of these products as not "just gambling." Of course, all of these products *are* ways to gamble, but, because they serve useful social functions, they have become politically and morally acceptable. Similarly, decision markets can also serve an important social function, especially when applied to government policy. Unfortunately, there is no large industry yet with an interest in lobbying the public to see decision markets as more than gambling. In light of elected officials' fear of making the same political mistake twice, the PAM experience is unlikely to be repeated anytime soon. Nevertheless, decision markets remain a potentially

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

---

1. John W. Pratt and Richard J. Zeckhauser, *Principals and Agents: The Structure of Business* (Harvard Business School Press, 1991).

2. Jeffrey S. Banks and Rangarajan K. Sundaram, "Adverse Selection and Moral Hazard in a Repeated Elections Model," in *Institutions, Competition and Representation*, edited by William Barnett, Norman Schofield, and Melvin Hinich (Cambridge University Press, 1993).

3. David Sears and Carolyn Funk, "Self-Interest in Americans' Political Opinions," in *Beyond Self-Interest*, edited by J. Mansbridge (University of Chicago Press, 1990), pp. 147–70.

4. Samuel Peltzman, "An Evaluation of Consumer Protection Legislation: The 1962 Drug Amendments," *Journal of Political Economy*, vol. 81, no. 5 (1973): 1049–91.

5. Avinash K. Dixit, *The Making of Economic Policy: A Transaction-Cost Perspective* (MIT Press, 1997); William Ascher, *Why Governments Waste Natural Resources: Policy Failures in Developing Countries* (Johns Hopkins University Press, 1999); Joseph Stiglitz, "The Private Uses of Public Interests: Incentives and Institutions," *Journal of Economic Perspectives*, vol. 12, no. 2 (1998): 3–22.

6. Anthony Barker and B. Guy Peters, *The Politics of Expert Advice* (University of Pittsburgh Press, 1993).

7. Paul Heidhues and Johan Lagerlof, "Hiding Information in Electoral Competition," *Games and Economic Behavior*, vol. 42 (2003): 48–74; Eric Maskin and Jean Tirole, "The Politician and the Judge: Accountability in Government," *American Economic Review*, vol. 94, no. 4 (2004): 1034–54.

8. Adam Brandenburger and Ben Polak, "When Managers Cover Their Posteriors: Making the Decisions the Market Wants to See," *RAND Journal of Economics*, vol. 27, no. 3 (1996): 523–41.

9. Robin Hanson, "He Who Pays the Piper Must Know the Tune," George Mason University, Department of Economics, May 2003 (http://hanson.gmu.edu/expert.pdf, last accessed July 2006).

10. Bryan Caplan, "Rational Irrationality and the Microfoundations of Political Failure," *Public Choice*, vol. 107, nos. 3/4 (2001): 311–31.

11. Andrew Lo, "Finance: A Selective Survey," *Journal of the American Statistical Association*, vol. 95, no. 45 (2000): 629-635. Koleman S. Strumpf and Paul W. Rhode, "Historical Presidential Betting Markets. *Journal of Economic Perspectives*, vol. 18, no. 2 (2004): 127–41.

12. One also has to assume that this agent's utility does not otherwise depend on which outcome is realized.

13. Justin Wolfers and Eric Zitzewitz, "Prediction Markets," *Journal of Economic Perspectives*, vol. 18, no. 2 (2004): 107–26; Martin Spann and Bernd Skiera, "Internet-Based Virtual Stock Markets for Business Forecasting," *Management Science*, vol. 49, no. 10 (2003): 1310–26; David M. Pennock, C.L. Giles, and F.A. Nielsen, "The Real Power of Artificial Markets," *Science*, vol. 291 (2001): 987–88. Robin Hanson, "Market-Based Foresight: A Proposal," *Foresight Update* (1990): 1–4; Robin Hanson, "Could Gambling Save Science? Encouraging an Honest Consensus," *Social Epistemology*, vol. 9, no. 1 (1995): 3–33; Robin Hanson, "Idea Futures," *Wired*, vol. 3, no. 9 (1995): 125.

14. Richard Roll, "Orange Juice and Weather," *American Economic Review*, vol. 74, no. 5 (1984): 861–80; Stephen Figlewski, "Subjective Information and Market Efficiency in a Betting Market," *Journal of Political Economy*, vol. 87, no. 1 (1979): 75–88; Pennock and others, "The Real Power of Artificial Markets"; Spencer Jakab, "New ICAP-Nymex Derivatives Have U.S. Gas Market's Number," *Wall Street Journal*, August 4, 2004, p. 1; Michael T. Maloney and J. Harold Mulherin, "The Complexity of Price Discovery in an Efficient Market: The Stock Market Reaction to the Challenger Crash," *Journal of Corporate Finance*, vol. 9, no. 4 (2003): 453–79. Joyce E. Berg and Thomas A. Rietz, "Prediction Markets as Decision Support Systems," *Information Systems Frontiers*, vol. 5, no. 1 (2003): 79–93; Kay-Yut Chen and Charles R. Plott, "Prediction Markets and Information Aggregation Mechanism: Experiments and Application," California Institute of Technology, Social Science, 1998.

15. Robin Hanson, "Decision Markets," *IEEE Intelligent Systems*, vol. 14, no. 3 (1999): 16–19; Marc Stiegler, *Earthweb* (Baen Books, 1999); Berg and Rietz, "Prediction Markets as Decision Support Systems." Michael Abramowicz, "Information Markets, Administrative Decisionmaking, and Predictive Cost-Benefit Analysis," *University of Chicago Law Review*, vol. 71, no. 3 (2004). Robert W. Hahn and Paul C. Tetlock, "How Information Markets Could Change the Policy World," AEI-Brookings Joint Center for Regulatory Studies, 2004.

16. The public would also need to monitor and punish any corruption in parameter judging.

17. Robin Hanson, "Shall We Vote on Values, But Bet on Beliefs?," George Mason University, Department of Economics, September, 2000 (http://hanson.gmu.edu/futarchy.pdf, last accessed July 2006). Robert W. Hahn and Paul C. Tetlock, "How Information Markets Could Change the Policy World."

18. Robin Hanson, "Combinatorial Information Market Design," *Information Systems Frontiers*, vol. 5, no. 1 (2003): 105–19.

19. Robin Hanson, "Designing Real Terrorism Futures" *Public Choice* (2006).

20. Such insider trades might be subject to disclosure and other rules similar to those used to govern insiders' participation in the stock market.

21. Emile Servan-Schreiber and others, "Prediction Markets: Does Money Matter?" *Electronic Markets*, vol. 14, no. 3 (September 2004): 243 - 251.

22. No evidence has ever been found connecting the September 11, 2001, terror attacks or the 1982 Tylenol poisoning to such trading. In fact, I have found no record of any such connection between a large harmful act and a price change.

23. Robin Hanson, "Foul Play in Information Markets," in *Information Markets: A New Way of Making Decisions in the Public and Private Sectors*, edited by Bob Hahn and Philip Tetlock (AEI Press, 2006).

24. Robin Hanson and Ryan Oprea, "Manipulators Increase Information Market Accuracy," George Mason University, Department of Economics, 2004 (http://hanson.gmu.edu/biashelp.pdf, last accessed July 2006).

25. Robin Hanson, Ryan Oprea, and David Porter, "Information Aggregation and Manipulation in an Experimental Market," *Journal of Economic Behavior and Organization* vol. 60 (2006): 449-459; Colin Camerer, "Can Asset Markets Be Manipulated? A Field Experiment with Racetrack Betting," *Journal of Political Economy*, vol. 106 (1998): 457–82; Wolfers and Zitzewitz, "Prediction Markets."

26. Charles Polk and others, "The Policy Analysis Market: An Electronic Commerce Application of a Combinatorial Information Market," in *Proceedings of the 4th ACM Conference on Electronic Commerce* (2003), pp. 272–73. An archive of information on PAM can be found at http://hanson.gmu.edu/policyanalysismarket.html, last accessed July 2006.

27. It was my job to survey possible application areas and recommend one. I recommended this one.

28. This mechanism is described in "Combinatorial Information Market Design." The experiments are described in Robin Hanson, John Ledyard, and Takashi Ishikida, "An Experimental Test of Combinatorial Information Markets,"

George Mason University, Department of Economics, February, 2005, http://hanson.gmu.edu/testcomb.pdf, last accessed July 2006

29. Ron Wyden and Byron Dorgan "Wyden, Dorgan Call for Immediate Halt to Tax-Funded `Terror Market' Scheme," July 28, 2003 (http://wyden.senate.gov/media/2003/07282003_terrormarket.html, last accessed July 2006).

30. Michele Norris and David Welna, "Reaction to Pentagon's Proposed Terror Futures Market," *All Things Considered*, National Public Radio, (9 p.m., July 29, 2003).

31. Philip E. Tetlock and others, "The Psychology of the Unthinkable: Taboo Trade-offs, Forbidden Base Rates, and Heretical Counterfactuals," *Journal of Personality and Social Psychology*, vol. 78 (2000): 853–70.