Chapter 3:
Designing Information Markets for Policy Analysis[70]
John O. Ledyard[71]

Designing Information Markets for Policy Analysis

Policy analysis, be it public or private, is all about acquiring and organizing bits and pieces of dispersed information in a way that will more clearly inform a decision that has to be made.[72] Information aggregation is the process of bringing all of the relevant information to bear on the decision. In this paper I look at an old approach to aggregating information, markets, and ask whether they might lead to improved policy analysis. I find that there are reasons to believe that information markets based on standard designs will not perform well in environments relevant to policy analysis. Two key reasons are the need for lots of contracts and the incentives to deviate from the behaviors necessary to achieve accurate aggregation of information. This means that new market designs are necessary if information markets are to be applicable to policy situations. One promising new design, a market scoring rule with conditional contracts, is described. It performs significantly better than standard markets in experimental environments that mirror difficult information aggregation situations one might expect to see in practical applications.

In "Information Markets for Prediction," I begin with an introduction to information markets and how they work to sharpen predictions. There is a growing body of theoretical, experimental, and field evidence that carefully focused markets can aggregate privately held, widely dispersed information into probability estimates that approximate the fully informed posterior. There is also a body of evidence that markets do this better than more traditional approaches such as surveys and committees. But there is also some evidence that this performance record is not perfect. In order to work well, information markets seem to require a complete set of contracts, a lot of common knowledge about the situation, and thick markets. This last is most likely to occur if there are many traders with small bits of dispersed information, who are risk-neutral and good Bayesians, and who follow straight-forward behavior.[73] In "Information Markets for Decisions," I look at what happens if one wants to use information markets as the basis for decisions. Here it is pointed out that many policy applications will lead to thin information markets. There are two basic reasons: Good policy analysis requires exponentially more contracts than prediction does, and there are more incentives to deviate from straight-forward behavior than in prediction models. One should expect neither thick markets nor straight-forward behavior in practical applications of information markets to decision analyses. Standard market designs do not work well under these conditions and there is no reason to expect them to accomplish full information aggregation. In "Knowing What Works," I ask what one might do to counteract the negative findings of "Information Markets for Decisions." Are there market designs that might perform better than standard approaches in those environments one expects to find in decision problems? Two possible design features are considered: conditional contracts and market scoring rules. The first provides more contracts for better information. The second offers a mechanism design solution to the thin market environment. I also describe a series of experiments in which two traditional designs and

two new designs are tested in fairly difficult information environments. The evidence from these experiments shows that using the market scoring rule with conditional contracts definitely sharpens the signal to noise ratio in information aggregation over traditional market designs. In "Final Thoughts," I provide some closing remarks for this paper. One of these is that there remains much room for better designs and performance increases.

Information Markets for Prediction

To get started, let us look at a simple example of a corporate policy decision. A company is significantly upgrading a piece of software that has two major components: A and B. The policy decision to be made is when to launch the marketing effort that will drive the sales of the new software. Getting the time of the launch right is worth a lot to the company. There are people in the company who have a good sense for the development of A and who can make an informed prediction on the basis of their information. They know a little about B. There are also people who are symmetrically placed with respect to B. he predictions of the two groups may well be different based on what they know. Who should be listened to? The obvious answer is both. In fact, what is desired is that the information of each group be merged or aggregated and predictions be made on the basis of the total information. If all of the individuals completely share their information, then ultimately their predictions will agree and be the same as those that would have been made by someone who had all the facts from the beginning. The dispersed information will have been aggregated. [74]

The standard methods used to collect and aggregate information, including committees, group retreats, bureaucracies, outside consultants, opinion polls and surveys are well known. Many have been around a very long time. More modern approaches include computer search engines. All have their advocates and all have drawbacks. A newer approach uses the power of carefully focused markets. These have come to be known as information or prediction markets.

What is an information market? One of the earliest examples of a market created specifically to aggregate information, the Iowa Election Market (IEM). The familiar question the market was designed to answer is: Who will win an election? In 1988, some faculty at the University of Iowa Business School securitized the Presidential election prediction and provided a trading system on the Internet. For $1.00, IEM stood ready to sell to you or buy from you a set of three securities that covered all possible outcomes of the election; Bush, Dukakis, Other. Each security paid $1 if and only if that person won. Securities were traded.

How this looks from the point of view of one trader is illustrated in Table 1. You as a trader have bought $1 worth of securities from the IEM that gives you one unit each of Bush, Dukakis, and Other. As indicated in the last column, your expected earnings if you did nothing else are $0. One of the three assets will pay you $1 and you have paid $1. Now you look at the market prices on the IEM and see that they are $0.50, $0.40 and $0.10 respectively for Bush, Dukakis, and Other. But you believe, on the basis of the information that you know, that Bush has a 70% chance of winning and Other has no chance at all. How can you make money? You buy one Bush for $0.60, sell your Dukakis for $0.35, and sell your Other for $0.05.[75] This trade costs you an additional $0.20 but you are better off. You now have 2 Bush that you believe each have a 70% probability of

paying $1. So your expected payment is $1.40. You expect to make $0.20 on your $1.20 investment.

Table 1: Trading in an Information Market

	Bush	Dukakis	Other	E(V)
You have	1	1	1	$1 -$1 = 0
Prices	$0.50	$0.40	$0.10	
You think	.7	.3	0	
So trade	1 @$0.60	-1 @ $0.35	-1 @ $0.05	- $0.20
You have	2	0	0	$1.40 - $1.20 = $.20

Other traders, *with different information from yours*, will look at these price changes, integrate the information that these price changes convey, and adjust their beliefs in response. They will also trade if their beliefs are different from the market prices. Eventually prices will settle down. Based on theoretical, experimental, and empirical evidence, there is reason to believe that, in equilibrium, the market prices for Bush, Dukakis, and Other will be an accurate evaluation of the probability each will win, based on all of the information of those participating in the market. Since 1988, the prices of the Iowa Election Market have bettered the standard polls 451 out of 596 times in predicting the election of candidates in a wide range of candidate competitions.[76]

One does not need to rely only on the evidence from the IEM, no matter how persuasive it is. There are by now many more information markets than just IEM.[77] There is also a growing body of scientific evidence to support the contention that information markets can provide significant aggregation. I turn to that evidence now.

The Field Evidence. There are many naturally occurring information markets. Stock markets, futures markets, and pari-mutual betting are just some examples. In direct comparisons with other institutions, they seem to perform very well. Racetrack odds beat track experts,[78] orange juice futures improve on weather forecasts,[79] and stock prices beat the panel of experts in the post-Challenger probe.[80]

One can also create information markets. In direct comparisons to other institutions, these designer markets, created explicitly to aggregate information, also seem to perform very well. I have already indicated the successful performance of the IEM. A market designed by Charles Plott and Kay-Yut Chen for Hewlett Packard beat the internal sales forecast process 6 out of 8 times.[81] Other evidence can be found in Wolfers and Zitzewitz.[82]

This positive evidence is all drawn from field studies. There is a lot going on in these situations that might interfere with accurate inferences, not the least of which are the

uncontrolled timing of the arrival of new information and the inability of the researcher to know for sure what the true fully aggregated information really was. I turn next to the evidence from more scientific approaches.

The Scientific Evidence. There are two basic strands of scientific work that support the hypothesis that information aggregation is possible. One relies on the intuition I gave in our first example that if information holders interact iteratively, they will eventually agree and that agreement will involve complete aggregation of their information. The second comes from the theory of rational expectations equilibrium – an equilibrium theory of information aggregation through markets rather than through personal interactions. I look at each in turn.

Iterative Aggregation. In his seminal paper, Aumann considered the situation in which two people have a common prior belief and then each privately learn something about the state of the world.[83] He proved that if their posterior beliefs become common knowledge, then they must have the same posterior, or final, beliefs.[84] Loosely stated, when one of them sees the other's beliefs, she will be able to infer something about the other person's information and add this to her own information. So will the other. This revision process continues until they have the same beliefs and these will be the full information beliefs. Geanakopolos and Polemarchakis were able to specify a precise dynamic process that implemented Aumann's insight for two people who follow the rules of Bayesian updating.[85] The process converges to full information beliefs in a finite number of steps. McKelvey and Page generalized the process to an arbitrary number of people.[86] They were also able to achieve full aggregation of information using only a public aggregate statistic to convey information from one to another. The conclusion to draw from this work is that if individuals are willing to honestly share their information and if they are able to process information as a Bayesian should, then iterative processes can lead to full information aggregation.

There has been at least one set of experiments to test whether iterative aggregation will work in practice. McKelvey and Page created an experiment in which subjects gave sequential reports in a series of rounds.[87] Between rounds there was a public report. They used a scoring rule[88] to try to provide some incentives for the correct revelation of information. They came to two conclusions: (1) observed individual behavior corresponded roughly to the predicted Bayesian updating and (2) some information aggregation occurred.[89]

Equilibrium aggregation. The second strand of work that supports the hypothesis that information aggregation is possible rests on the equilibrium analysis of markets and other institutions. At the core of the analysis is rational expectations theory.[90] Individual demands reflect their beliefs about the world, as in the example in table 1. Prices respond to those demands and so are public signals that convey information about the world. Traders incorporate that public information into their private expectations. With these new expectations come new demand functions and, therefore, new prices. In equilibrium, all is revealed and expectations are fully informed.

Putting this theory to work, one can create securities that pay $1 in a particular state. These are called a complete set of Arrow-Debreu securities[91] if there is one for every possible state.[92] If traders are risk neutral and price-taking, and if markets fully equilibrate, then in equilibrium the prices of each security will equal the common knowledge posterior probability of that state. Market equilibrium will aggregate all of the information of the traders.

There is experimental work evaluating whether the equilibrium theory accurately describes a real situation. Forsythe, Palfrey, Plott created a laboratory securities market and demonstrated convergence to a rational expectations equilibrium.[93] McKelvey and Ordeshook were able to show that information aggregation occurred in a laboratory election where individuals had private information about the candidates and polls provided the public signal.[94] This aggregation was consistent with rational expectations theory. In a really important paper in this area, Plott and Sunder created laboratory securities markets where multiple securities were traded in an uncertain world.[95] Individuals had private information about the probability of some states. One of their key findings was that if individuals have the same preferences and if there is a complete set of Arrow-Debreu securities, then there is rapid pooling of information, consistent with rational expectations equilibrium theory.

Some Storm Clouds. To this point, things look pretty good for the claim that markets can be created to aggregate information. Both iterative and equilibrium theories are positive. The experimental evidence is positive. And the field evidence is positive. But there are also some cautionary notes in the literature.

Grossman and Stiglitz point out an apparent paradox that must be confronted at some point: why pay to gather information if it is just going to be revealed to others for free?[96] The idea can be illustrated simply as follows. Suppose there is a trader who knows the state for sure. He is willing to pay up to $1 for an Arrow-Debreu security on that state and sell all other securities at any positive price. But if the others know that someone might know something,[97] then as soon as this trader bids up the price of the one security, the information will be revealed and no one will sell that security for less than $1. And no one will buy the other at any price. Prices have converged to the full information probabilities with $1 on the true state. But, and here is the potential problem, the trader with the information cannot make any money by trading, so why bother? And if they don't bother, the equilibrium is never revealed.[98]

This kind of intuition is also captured in the "no trade" theorem of Milgrom and Stokey.[99] They consider the same situation Aumann did. Two traders begin with a common prior and then each learns something privately. It can be shown that they will never trade solely on the basis of differing information. The intuition is fairly easy to understand. Suppose I have two coins: one is weighted so that heads comes up 80 percent of the time, and the other is weighted so that tails comes up 80 percent of the time. I randomly pick one of these and, for each of two people, flip the coin twice. Neither sees the results of the other's flips. I then create an information market by allowing trade in an asset that pays $1 if the coin is the one weighted towards heads. Suppose person 1 sees a head and a tail, denoted (1H, 1T). His posterior belief is then that the probability the coin is weighted heads is 50%. Suppose person 2 offers to sell two units of the asset for $0.30 each. Person 1 can immediately infer that person 2 saw,[100] (0H, 2T) and therefore knows

that the total information is (1H, 3T) giving a full-information posterior that A is weighted heads of 6 percent. So, person 1 will reject the offer (and, indeed, any offer above $0.06). From this rejection, person 2 knows that person 1 did not see (2H, 1T), and person 2 knows that the full information posterior probability the coin is weighted heads is less than or equal to 6 percent.[101] At this point, there is no basis for a trade to take place. This strongly suggests that if traders have a common prior and are risk-neutral, information markets will not work as advertised.

In another paper casting doubts on rational expectations equilibrium theory, Jordan showed that there are situations in which there are "under-revealing" rational expectations equilibria and there are situations in which the rational expectations equilibrium will not even exist.[102] The first situation can cause information aggregation to be stopped before it is complete. The second situation means that these markets won't equilibrate as expected. Both of these papers indirectly raise important questions about the dynamics of these markets—the process of price discovery. Finally, although Plott and Sunder had success in getting information aggregation under the right conditions, they also demonstrated that if the preferences of individuals differ and there is only a single asset without contingent claims, then relatively little aggregation of information happens.[103]

The iterative approach has been less extensively discussed. McKelvey and Page point out that the individual behavior they observed in the lab is not generally fully rational.[104] Bayesian updating by the subjects is incomplete—and complete updating is required to get full information aggregation.

Summary of the Scientific Evidence on Prediction. At this point, it seems fair to me to summarize the scientific evidence as cautiously optimistic. We have seen numerous papers suggesting that information aggregation is possible, either iteratively or in equilibrium. But there are storm clouds that threaten the parade. The iterative theory assumes individuals will honestly reveal their beliefs when asked, though this may not be compatible with the actual incentives faced by participants. And if not, then there is no reason to expect full or accurate aggregation. The equilibrium theory requires a complete set of securities, price-taking behavior, and adequate price discovery. In thin markets, price-taking may not be compatible with the traders' incentives and price discovery may be incomplete. That is, there can be incomplete or inaccurate aggregation. Both theories rely on each trader's ability and interest in understanding the common knowledge structure, inverting price and other signals into information about other traders' information, and using Bayesian updating when they revise their beliefs. If traders do not behave this way, then inaccurate and incomplete aggregation will occur.[105] When one tries to transform information markets into policy analysis markets, incentives, thinness, and incomplete updating will become major issues.

<div align="center">

Information Markets for Decisions

</div>

Up to now, I have been looking at information markets strictly as a method to improve information retrieval and aggregation. In policy analysis, one is trying to make a decision. Prediction is only a part of this . In committees and other processes for information aggregation, many variations in a policy are explored and evaluated before a final recommendation is made. Small variations in key sections can create significant

variations in predicted benefits and costs. For markets to duplicate these far-ranging considerations, they must require the trading of many securities simultaneously. Of course, this introduces a number of new potential problems to worry about. I consider three in this paper. There are, of course, others.

Information Markets for Decisions Require Many Securities. Consider an example taken, and modified, from Hahn and Tetlock.[106] A policy maker wants to improve the standardized scores of a school district. She is considering a policy that would delegate the right to run the schools. To estimate the benefits of such a policy, she decides to set up two markets. After the markets close, a decision will be made whether to implement the policy or not.
   - In market 1, a contract is traded that pays $x if the average score is x, conditional on the policy not being implemented. If the policy is implemented, you get your money back.[107]
   - In market 2, a contract is traded that pays $y if the average score is y, conditional on the policy being implemented. If the policy is not implemented, you get your money back.

The policy is described before the markets open.

Suppose June 30 arrives and the prices of the contracts in market 1 and 2 both equal $100. Such prices clearly suggest that the policy will have no effect, so if it costs anything, it should not be implemented. The policy had a number of key clauses that described what the firm could and could not do in running the schools. Any one of these could have caused the market valuation to be so low. The policy maker may wonder what went wrong and question how to revise the policy to improve it. But there is no way to know. If markets are to replace other aggregation procedures, they will need to be more informative than this.[108]

But suppose that the policy maker has two versions she can live with for each of, say, three clauses. Let the clauses be identified by A, B, C, and let the versions of each be either 0 or 1. If the policy maker sets up the markets so that one can buy or sell the contracts in market 1 conditional on, say, A = 1 or A = 0, and B = 1, then the prices of these conditional contracts will signal what the increase in scores will be, conditional on the particular version to the policy. Of course, if there are N clauses with K versions, then there will need to be $4x(N^K)$ securities traded, only two of which will actually pay off in the end. For the example here, that means thirty-two securities.[109]

The main point is that the number of securities to be traded increases exponentially in the fineness of the detail desired. Better analysis requires a significant number of simultaneous markets. The more markets, the more likely it is that neither price-taking nor honest revelation of information will be incentive compatible if the markets are organized in the usual way. This means that the market design is going to be crucial. New processes will be needed that can encourage the aggregation of information even though there are many securities and few traders.

Policy Maker Temptation. Potential problems in the use of information markets as decision tools may arise because of the situations in which they are embedded. If the markets are being used as the basis for decisions, then those decisions may have consequences for both the policy maker and the traders. If the consequences are large,

then each party's interest in the decision can easily outweigh their interest in the markets themselves.

Let us look at an example of the temptations faced by the policy maker. Acme Software, Inc. plans to ship its product, Zeus, next July. It is built from components A, B, C; each can be new or old. Acme management is crucially interested in getting as much information possible on the shipment date so they can correctly time the massive marketing campaign and sales force. Shipping in October could cost them $70 million in foregone sales. In January, Acme opens a market to predict the Zeus shipment date allowing trade in the following securities: {before Aug},{Aug},{Sept}, {Oct},{Nov}, and {later than Nov}. In March, the October prediction is 75 percent since the price is $0.75. Management orders a review, identifies component A as the problem, stops production of A_{new}, and retains A_{old}. Prices soon move to .03 for October and .45 for both August and September. Zeus ships in August.

So what is the problem here? Under this scenario, the traders who signaled October (assuming A_{new}) by buying at prices up to 0.75 lose.[110] If they had correctly anticipated the actions of management, they would not have revealed their information. But if they had revealed no information and the October price had stayed low, management would have lost money. Without any changes in the market setup, the traders will include in their probability estimates what they think management might do. This will significantly dilute the quality of the information being revealed by the market.

There are at least two ways to redesign the markets so that the type of problem faced above by Acme Software will not happen. One is for the policy maker to commit to not doing anything to change the outcomes on which everyone is trading. But living up to that commitment could be very costly. It would have cost Acme up to $70 million to do so in our example. There is a second, cheaper way. Let traders hedge against the risk of adverse decisions In our example that is done by creating options like {Zeus in October if A_{new}} and {Zeus in August if A_{old}}. If the condition does not occur, then the trader is off the hook. With these options, traders who know that the October delivery is likely because of the problem with A can indicate their beliefs by selling {Zeus in October if A_{new}} and buying {Zeus in August if A_{old}}. There is no need for them to worry about the effect of management's response on their market earnings.[111]

There are, of course, added costs and benefits. The added benefit is that management will not need to enter into a costly review to locate a problem. It will be easy to identify from the prices. As I argued earlier, having more markets potentially provides better information. But the added cost arises in the same way it did then. One now needs to have six dates, two features, and two states for each feature, giving us twenty-four securities. If there were four features, one would need $6 \times (4^2)$, or ninety-six, securities. The cost of hedging is an exponential increase in the number of securities to be traded. This increase leads to a thinning of the market that can lead to a loss in accurate price discovery. The trade-off faced by the market designer is between bad information revelation, for fear of future decisions, or bad information revelation, because of thin markets. New market processes that create better price discovery in thin situations would alleviate the latter.

Trader Temptation — Information Monopoly. Even if traders do not care about and cannot affect the outcomes that determine the payoffs of the securities, they may be

tempted to manipulate the market itself in order to increase their profits. If a single trader knows something no one else knows (inside information), he has a monopoly in that information and can gain, at least in standard market designs, by manipulating the price. Just as a product monopolist can create "artificial scarcity" and thus keep the price of her sales up, so can the information monopolist. The argument is simple: If the final prices reflect aggregate information, then they are a well-defined function of the various pieces of information each has. Suppose I create a very simple information market in which I want to know the probability it will snow in New York City on January 1, 2005. In a rational expectations equilibrium, if I am risk neutral, my willingness to pay is $q^i(s^i, p^*(s))$, where q^i is my belief it will snow, s^i is the private information I have, and $p^*(s)$ is the equilibrium market price given s, the signals of everyone. In a rational expectations equilibrium at the final prices under price taking competitive behavior with risk-neutral preferences, it must be true that $q^i(s^i, p(s)) = p(s)$ for all of us, otherwise I will still want to trade. Let the price that satisfies all n equations be $p^* = g(s^1, s^2, \ldots, s^n)$. If the markets accurately aggregate the information as predicted by the theory, then $g(s)$ should be the posterior based on the full information in the vector of private signals s. So, if markets are performing their predicted function of information aggregation they must be sensitive to each individual's information. But it is just this sensitivity that gives the monopolist her leverage.[112]

How might a monopolist pull off her manipulation? If everyone knows she is a monopolist, then by closely tracking her bids and offers, they can infer her position and eliminate her advantage. But she can prevent this by acting competitively (in a way that is indistinguishable to an outside observer from the price-taking behavior normally assumed) while mimicking the behavior of a trader like herself who has different information. In doing so, she can guide the market to a different equilibrium, the one that assumes different information is correct. In many cases this can be done in a way that leaves the monopolist better off. To see how this works, consider a very simple example. Suppose each agent receives a signal that is H, M, or L. If it is H, the agent knows for sure that it will snow on January 1st in New York City; if it is M, then it is likely; and if L, then it will not.[113] The more Ms, the higher the price will be in equilibrium. Suppose there are five people in this market and the information is 2 M, 2 L, and 1 H. As predicted by rational expectations equilibrium theory, when the holder of H bids the price of the security up, all others will infer H and the price will instantaneously jump to 1. The holder of H will complete few, if any, trades at a price lower than 1, and therefore, in equilibrium, prices will reflect the true probabilities but no one will have made any money. But if the holder of H pretends to be M, then they will be able to buy assets at prices reflecting the difference between p(2M, 3 L) and p(3M, 2L). The monopolist knows for sure that the security will pay off at $1 and so the monopolist makes more money with this strategy than by behaving competitively. And, the equilibrium price will not reflect the full information posterior for which p = 1.

I have purposely exaggerated this example to make the point that if information is not widely held, competitive behavior may not be incentive compatible and information may be hoarded, causing prices not to reflect the full market information. If there were a hundred people instead of five, the gains from this type of manipulation would obviously be lower and might in fact be less than those from acting like H. Also, if two people held the information H in our example above, then neither is a monopolist and competition

may drive the price to 1.[114] If information were more dispersed, then the example would not be as sharp. But if the markets are thin, there are a small number of traders, and some individuals' information is important, then standard markets will provide temptations to misrepresent that information and may lead to incomplete aggregation of information.

There is a way to design the markets so that the temptation to manipulate price is reduced or eliminated. The design does not rely on rational expectations theory and its companion assumption of competitive behavior. Instead it builds on a method known as Proper Scoring Rules. Originally suggested by statisticians such as Brier[115] and Savage,[116] who were trying to elicit probability beliefs from, weather forecasters and others, scoring rules accept the incentives of the individuals and create a mechanism to deal with it. It is very easy to see how they work. An individual is asked to report a probability r. The individual will be rewarded with $$s_i(r)$ if the actual outcome is i. Although scoring rules come in many forms, a classic example from Goode uses logarithms. For two possible outcomes, the Goode rule pays $\ln(r_1/.5)$ if outcome 1 occurs and $\ln((1-r_1)/.5)$ if outcome 2 occurs. A risk-neutral, expected value maximizing agent faced with this rule and believing the true probability of state 1 is q will want to report r to maximize $q [\ln(r/.5)] + (1-q) [\ln((1-r)/.5)]$.[117] It is easy to see that $r = q$ solves this problem uniquely.[118] So as long as one is willing to pay for it, one can elicit correct information from an information monopolist.

Can scoring rules be made to work for a group of people? The answer is yes. Hanson has provided a clever way to do this by introducing a Shared Scoring Rule.[119] A Market Maker (easily automated) stands ready to trade with any of the N traders. A trader proposes a price p and the Market Maker offers to change their assets so that they get $s_k(p) - s_k(p_{t-1})$ of the asset that pays $1 if and only if k occurs, where s(p) is one of the scoring rules and p_{t-1} is the last price the Market Maker traded at. As Hanson puts it, "Anyone can use the scoring rule if they pay off the last user." If agents follow their dominant strategy to tell the truth—to report their current posterior as their price proposal —at the time they trade, then iterative aggregation theory applies and full aggregation of information will occur.

By now I have moved a long way from standard market designs with securities that only predict and bulletin board trading systems. But that is the point. If I want to use markets to aggregate information for decisions, then there will be many frictions and temptations if I stick to traditional approaches. Good market designs will have to be guided by the principles of incentive compatible mechanism design. Two suggestions from our earlier analysis to obtain the desired conditional information and to insure against adverse behavior are (a) the use of extensive conditional contracts to better inform decisions and to protect against policy maker incentives to prematurely use the information and (b) the use of market scoring rules to protect against the information monopolist incentives to manipulate prices.

Conditional contracts mean a large number of contracts, thinning out the process. Scoring rules mean complex formulae, causing potential behavioral problems. Equilibrium and iterative aggregation theory are unlikely to be a good predictor of how well information aggregation will actually occur. How will one really know what works?

Knowing What Works

I have found the use of experimental test beds to be of considerable help in identifying how mechanism designs may work when confronted with reality. The idea is to create an economics wind tunnel in which one can try out various designs in a controlled environment and one can measure the performance. The closer the environment is to capturing the situation in which the mechanisms are to be used, the better. To test information markets we[120] created two very difficult environments. In earlier experimental analyses like those in Plott and Sunder[121] and McKelvey and Page,[122] the primary purpose was to show that, in fairly simple environments, market aggregation occurs. Our goal was to seriously stress test the claims.[123] We succeeded.

The Experimental Test Beds. We created two environments. The first had three traders and three correlated events. An event is what one wants to estimate a probability for. An example would be "test scores in district A increase over last years scores" or "Syrian GDP declines in 2005." In our experiments, each event has an outcome that can either be 1 or 0. A set of simple markets can be set up in this situation—one for each event. But in our experiments, the events are partially correlated so that a complete set of markets would require $2^3 = 8$ markets. The second environment had six traders and eight events, requiring 256 securities for a complete set. These environments really are really thin. One measure by which to see this is the number of traders per state of the world. In simple asset experiments with standard markets, such as those in Bossaerts and Plott,[124] there are three states and forty or more traders, or 13.33 traders/state. When markets are thinner, as in an earlier study by Bossaerts and Plott,[125] there may be twelve traders and three states, meaning there are four traders per state; standard markets do not equilibrate well although call markets do.[126] In comparison, the three event environment has 0.375 traders per state and the eight event environment has 0.0325 states per trader. I believe this is a significant challenge to any mechanism, much less standard designs.[127]

In each environment we induced a common knowledge prior.[128] We described a set of possible urns of balls, each ball containing an outcome such as (A = 1, B = 0, C = 1). In the three-event environment there were six possible urns. In the eight-event environment there were 40,320 possible urns. We then selected one of these urns at random, with equal probability. To provide private information, we drew ten balls at random from the urn we had chosen. In the three-event environment, this means there are ten outcomes, like (1, 0, 1), (0, 1, 1), (1, 1, 1), and so forth, representing the outcomes for (A, B, C). We then showed part of these data to each trader. In particular, we showed the results for A and B to one trader, the results for B and C to another trader, and the results for A and C to the other trader. The feature of reality we were trying to capture was that of an "area expert." Someone may know and observe outcomes only in, say, Iraq and Iran. Another may know and observe results in Iran and Syria. And so forth. These experts' areas may overlap but are seldom identical. Everyone knew the descriptions and the selection process. From these all traders could compute a prior and an individual posterior over the possible states of the world. Also the experimentalist, knowing all of the draws, had a sample of ten and could compute the "full information posterior," which completely aggregates all of the private information. This is information that one does not have if one is using field data and is one key reason experiments can be significantly more informative.

After trading or whatever mechanism was employed, we would draw a hundred new balls and pay the subjects according to that draw. That is, if they are hold one unit of the security (1,1,1) then they are paid $1*(% of the draws that yield A = 1, B = 1, and C = 1).

Measuring Performance. In the environments we create, we can run a series of trials using different mechanisms such as opinion polls, standard markets, or new designs we might want to test. At the end of the trials we will know the probability density over states that the process produced. Because we control everything, we also know the full information aggregation probability density. We can measure the difference between the two. In the analysis below, we did this using the Kulback-Leibler measure of the difference between a process prediction, p, and the full information prediction, q.[129] The measure is $\sum_k q_k \ln(p_k/q_k)$, which equals zero when p = q.

To illustrate both how this measure works and to indicate the scope for information aggregation in our environments, I provide two figures, one for each basic environment.[130] In each figure, the horizontal axis indicates the distance to the full-information posterior where further to the left is closer. For each information draw in an environment, we computed the distance from the full information posterior of the individual priors and the individual posteriors.[131] I also include the distance that the uniform density (the completely uninformative density) is from the fully informed posterior as a lower bound on what one might expect to see in any data. I then plotted the cumulative distribution function of those distances. The height of the curve at a point on the x-axis represents the percentage of the environments that were within that distance of the full information posterior. A curve that lies to the left of another is, in some sense, more informative. In figures 1 and 2, I provide the basic data respectively for the three trader–three event environment and the six trader–eight event environment. In each case, the individual posteriors are more informative than the individual priors and those, in turn, are more informative than the uniform beliefs. If the traders are competent Bayesians, they will begin with their prior and, after seeing their individual information, believe their individual posterior. If a process is successful at fully aggregating the information available, the cumulative distribution of its predictions will be a vertical line at zero. A key point to note for later is that the scales are significantly different. All of the curves in the eight-event case lie to the right of the curves in the three-event case, indicating that eight events and 256 states is a significantly more difficult world.

Process Performance. We looked at the performance of four different information aggregation processes in each of the two environments. We looked at two standard processes: opinion surveys and bulletin board markets. For the first of these, we asked individuals to describe their beliefs and we used a scoring rule to provide some incentives for them to do this accurately. They were not involved with anyone else and received no feedback from anyone. For the second, we used Marketscape, a web-based multiple-market system developed by Charlie Plott (www.eeps.caltech.edu). There was one market for each event. We also looked at two less standard mechanisms suggested by our earlier analysis: an iterative call market with conditional contracts and a shared scoring rule with conditional contracts. For each process we ran sixty trials, where each trial involves a new draw of information and a complete run through of the process. For the three-event

environment, we ran twenty-four trials (four groups of three subjects each participated in six trials). In each trial, three individuals computed seven independent prices in twelve minutes. For the eight-event environment, we ran thirty-six trials (six groups of six subjects participated in six trials each). In the eight-event environments, at least twelve of the trials involved subjects who had already participated in six of the three-event trials and six of the six-event trials. In each trial, six individuals computed 255 independent prices in twelve minutes.

The data suggest that subjects were having problems understanding these difficult environments. In figures 3 and 4, I graph the cumulative distributions summarizing the performance of the individual scoring rule in the three-event and eight-event environments. The cumulative distribution provides the probability that a randomly selected individual will provide a probability estimate within that distance from the full information density. As one can see, in the three-state world, two-thirds of the individuals managed to beat the theoretical prior, but less than 5 percent actually provided as much information as they could have had in their individual posterior.[132] In the eight-state world performance is worse. The subjects were better informed than the uniform prior, but they were not as well informed as the theoretical priors, again providing an indication that this is a very difficult world they have to deal with.

One might wonder at this point whether the subjects really understood the environment and, if not, whether the data from various mechanisms should be taken seriously. I think the answers are no and yes. That is, the subjects had difficulty understanding the environment and some were obviously better at it than others. But each had information about the world that could be aggregated. The real question can still be asked and answered: Can any process sharpen the signal to noise ratio in a complex situation? I turn to that now.

In our first attempt to create information aggregation we used standard markets. We created two securities for each event,[133] provided subjects with an initial endowment of each, and let them trade through a bid-ask system with an open book. Surprisingly, to me, these standard markets did not help. In fact, the final predictions were even worse than those provided by individuals acting on their own. The data are graphed in figures 5 and 6. With the exception of the bottom 15 percent of the trials, the cumulative distribution for the market experiments, of the distance of the final density from the full information density, lies to the right of that for the individual data. In these very thin environments, very little trading takes place, price discovery does not occur, and equilibria are hard to find. The markets seem to garble the aggregate information rather than sharpen it. In the whirlwind we created in the lab, one can do better than markets by randomly selecting someone and asking them what they think. The implication is that unless one is convinced that there will be thick markets, one cannot achieve good information aggregation by just simply creating a collection of securities and then declaring "let there be markets." It just won't work. New designs are necessary.

A natural conjecture based on the work of Bossaerts, Fine, and Ledyard[134] is that the process should work better in thin situations if one includes conditional contracts and uses a call market.[135] With conditional contracts, a trader can buy or sell securities such as (A = 1 if B = 0) which pays \$1 if both A = 1 and B = 0, pays \$0 if A = 0 and B = 0, and returns one's money if B = 1. In this way, traders can display and trade on partial information they may feel more comfortable with. Further, the process should also work

better if we use a call market to "thicken up" the trading by clearing a larger number of contracts at one time. Following this intuition, we created and ran a Call Market with Conditional Contracts. The Call Market was an iterative process where, in a series of rounds, subjects submitted bids on the complete set of contingent contracts and then an algorithm provisionally cleared the market and set a price for each state. Then a new round began for up to five rounds. Bids that were provisionally cleared in one round were automatically resubmitted to the next round. The data from our experiments are provided in figures 7 and 8. In the three-event environments, this design provides an improvement over both individual observations and standard markets. The predictions of the Call Market are closer to the full information prior. But this is only a qualified success at best. Even though there is a sharpening of the information, the subjects still don't get to the distribution of theoretical individual posteriors. Further, in the eight-event environment, there is actually no improvement over standard mechanisms. There is no difference in top 50% of the distributions for any of the three processes: individual observations, standard markets, or call markets with conditional contracts. Further design is clearly necessary.

From earlier discussions, one might conjecture that one of the inhibiting factors to the full aggregation of information could be the fact that subjects are not "small" in these environments. In thin markets, informational monopoly power exists and can lead to incomplete bidding. It is not incentive compatible for traders to behave as is required for markets to fully aggregate information. To mitigate such frictions, we turned to a Shared Scoring Rule with Conditional Contracts. The data from those trials are summarized in figures 9 and 10. Here there is some success. There is a significant improvement in the eight-event world and some improvement in the three-event trials. The shared scoring rule does better the traditional approaches to market design. But it is also true that the subjects still did not get to the distribution of theoretical individual posteriors. We have found a process that does sharpen the signal to noise ratio through some type of aggregation process. We have not found a process that fully aggregates all information.

Final Thoughts

Information markets for predictions are possible and desirable. There is a growing body of scientific evidence that supports the ability of carefully focused markets to aggregate privately held, widely dispersed information into probability estimates that approximate the full information posterior. There is also a growing body of applications that support the hypothesis that markets improve on other more traditional methods like polls, committees, and other forms of expert opinion. Markets seem to work best in thick situations where there are a lot of traders with small bits of information, where traders behave competitively as price-takers, are risk-neutral, and are good at Bayesian updating, and where the traders do not have a direct stake in the outcomes. The increasing use of these markets will improve our ability to recognize and deal with uncertainty.

The evidence for the use of information markets to support policy decisions is not as straightforward. Many policy applications will be in thin situations. Better policy analysis requires more contracts, which exponentially explode in the number of alternatives considered. There are often a small number of traders with exclusive and important pieces of information whom I have called informational monopolists. And there are temptations for both the policy analyst operating the market and the traders in the market to take actions that will cause information not to be revealed. Traditional

market designs do not work well under these conditions and there is no reason to expect them to create a full aggregation of information.

New market designs will be needed. Designs must be able to operate effectively in thin situations. They must be sensitive to issues of incentive compatibility, computational complexity, and immediacy. We looked at two possible alternatives to traditional approaches: call markets and shared scoring rules. We found that the Shared Scoring Rule design definitely sharpens the signal/noise ratio in thin and ultra-thin markets over that achieved by traditional markets. But we did not find a process that did even as well as the theoretical individual posteriors. It is not clear one exists, but it is highly likely there are other designs that can do better.[136]
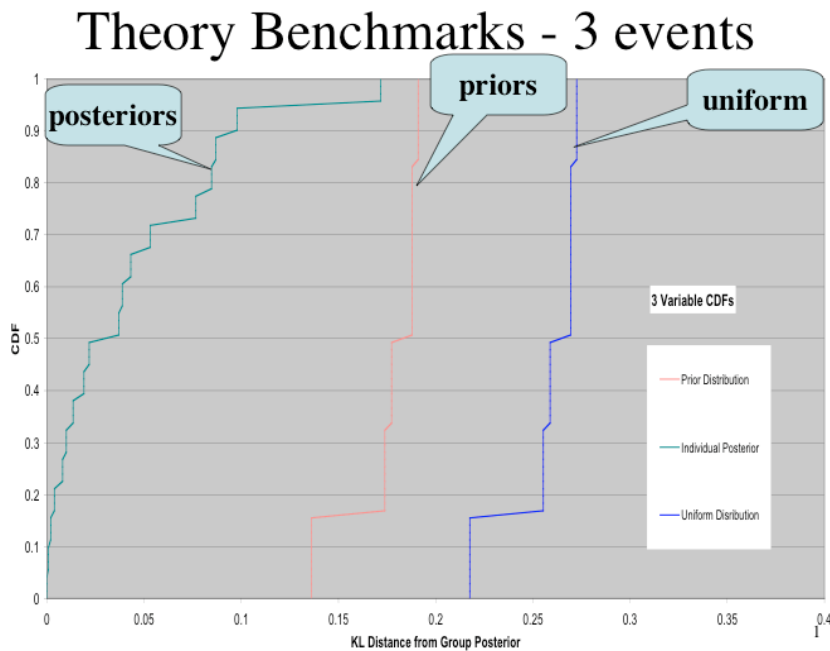
Figure 1
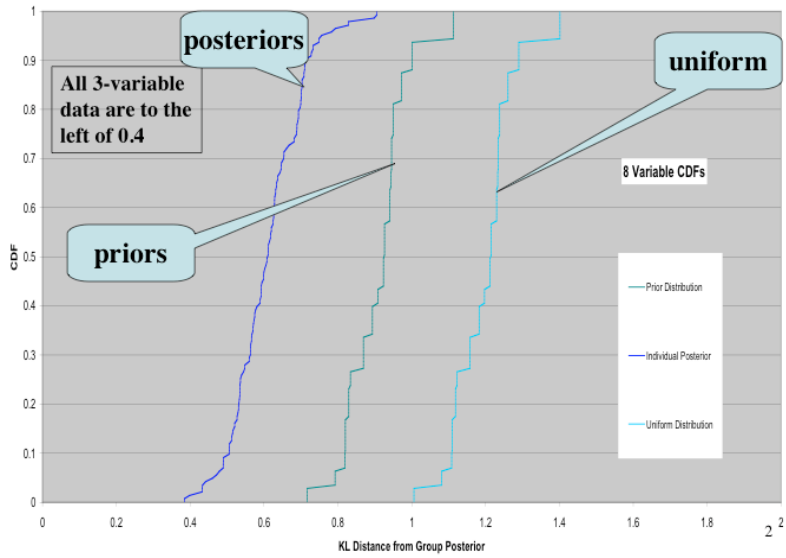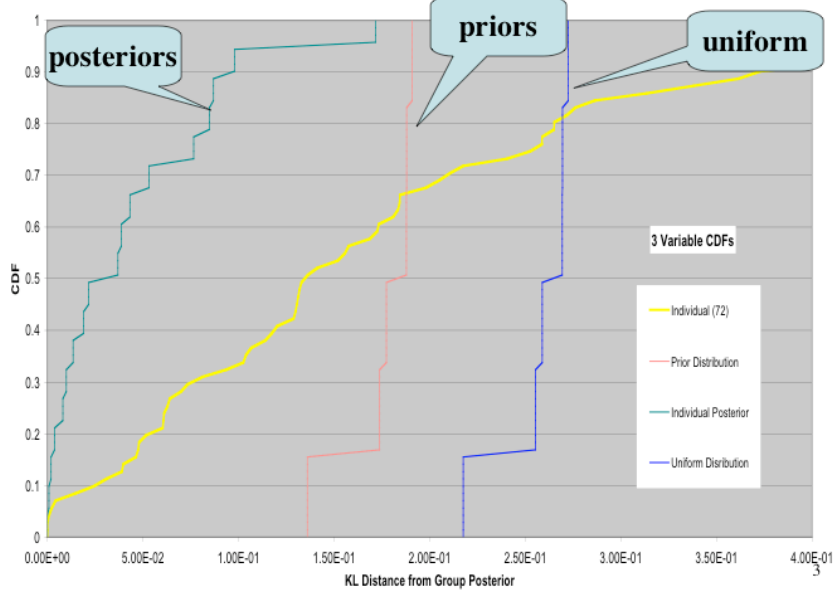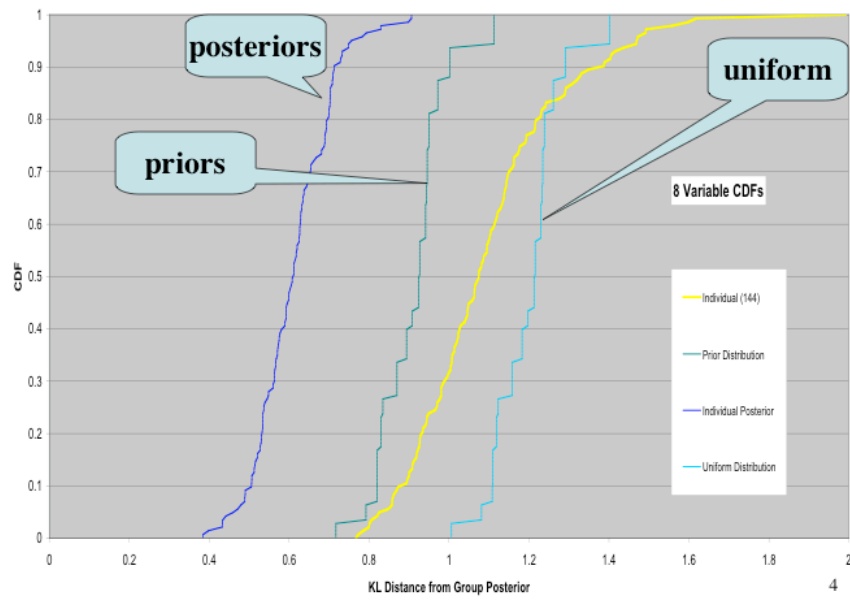
Figure 2
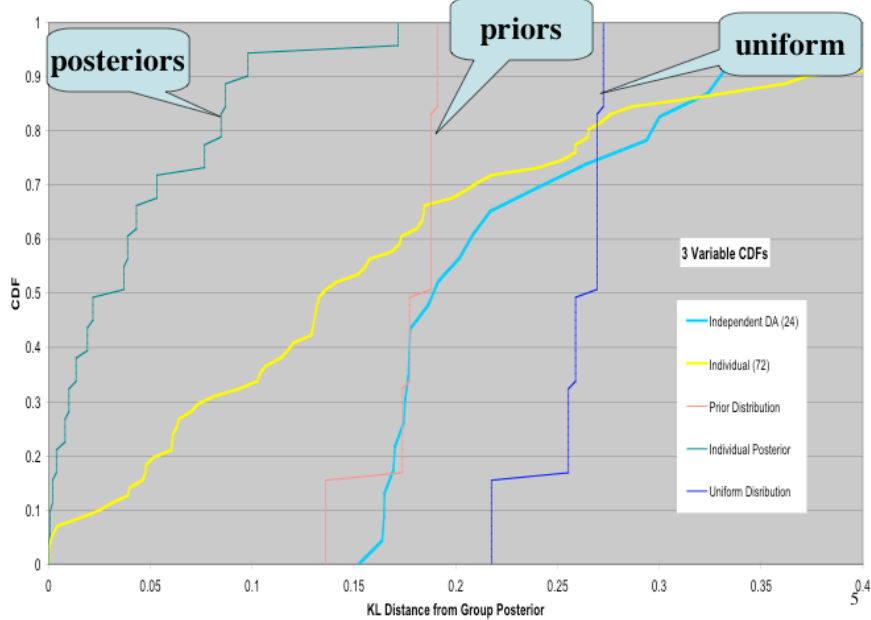
Figure 3

Figure 4

Figure 5
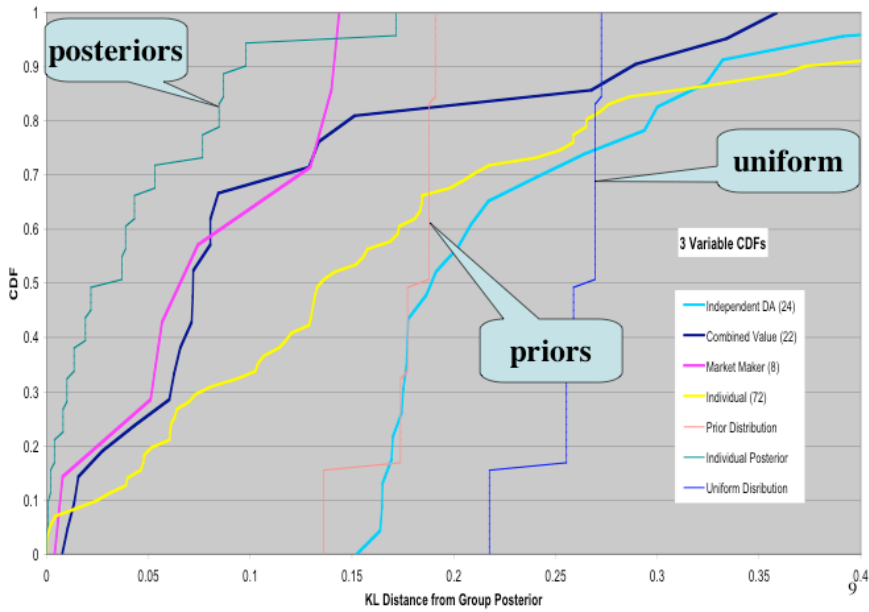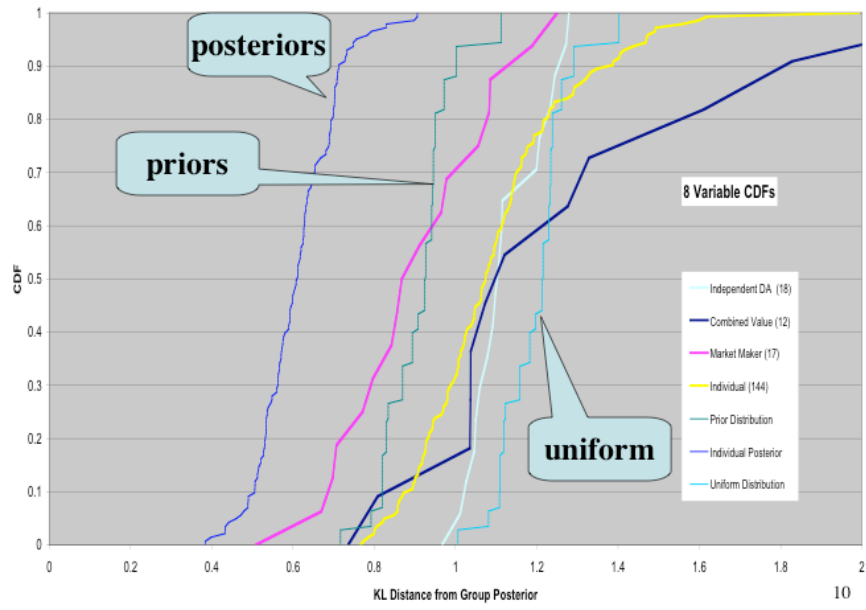

Figure 6

Figure 7

53
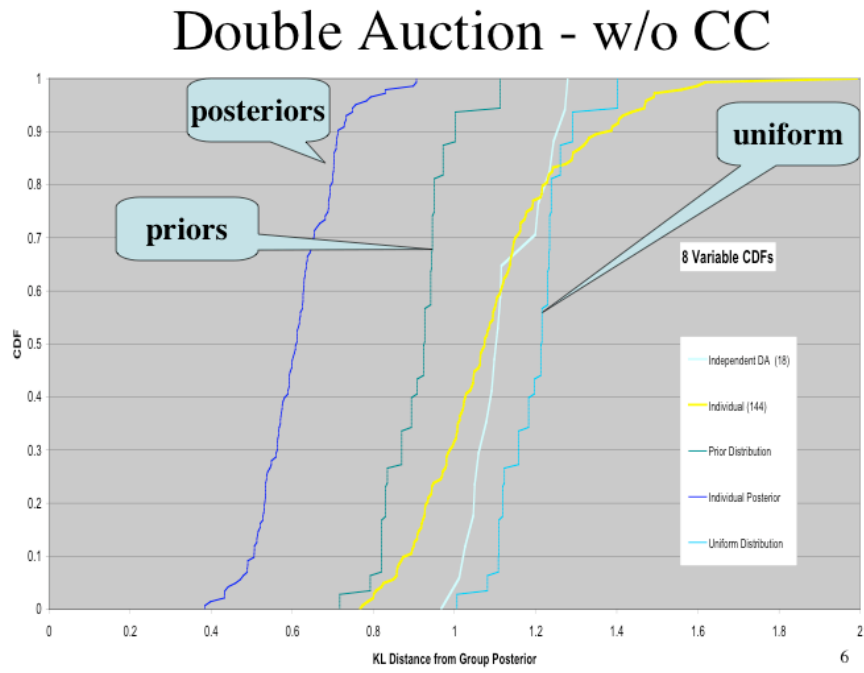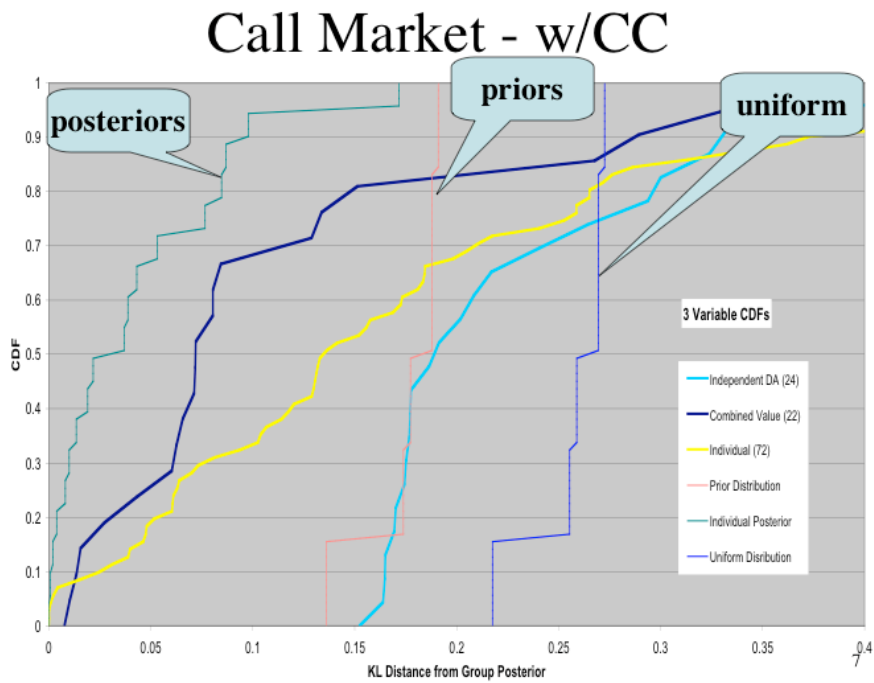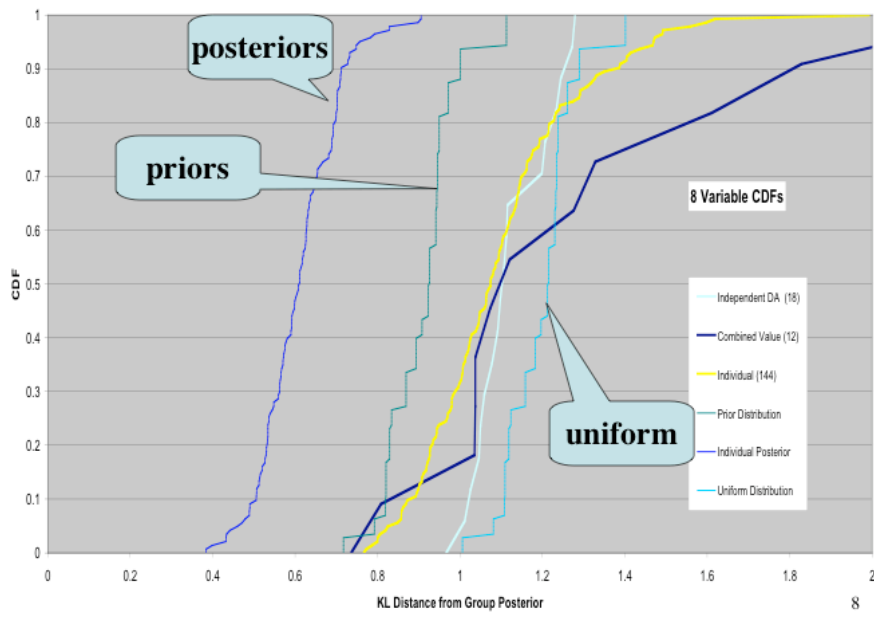
Figure 8

Figure 9

Figure 10

[70] Much of this paper is based on research I did together with Robin Hanson and Takashi Ishikida, particularly the experimental work reported in Section 3. They are not responsible for the content. I thank the participants at the AEI-Brookings conference on "Information Markets: A New Way of Making Decisions in the Public and Private Sector," Washington D.C., December 2004 for their helpful comments, particularly Robert Hahn.

[71] Contact information: Professor of Economics. 1200 E. California Blvd., MC228-77, Pasadena, CA 91125. Email: jledyard@caltech.edu.

[72] Policy analysis can, of course, also be about persuasion and obfuscation and not about clearing up the issues. But most groups, even those that intend ultimately to cloud issues, privately want to know as much as they can about the true situation before doing so.

[73] By straight-forward behavior I mean price-taking behavior in the marketplace or honest reporting of beliefs when asked. This is most likely to occur when all traders consider themselves to be part of a team, such as a jury, that really wants to find out what the answer is. It is less likely when the traders in the market have differing beliefs as to what the outcomes based on the information should be as will generally be the case in a trading situation where one person's gain is the other's loss.

[74] All beliefs will be the same and will be equal to the full information posterior probability distribution function.

[75] Actually, as long as the price of Bush is less that $0.70, the price of Dukakis is larger than $0.30 and the price of Other is larger than $0, you will be better off after the trade. I am, of course, assuming throughout this example that you are risk-neutral.

[76] Joyce E. Berg, Robert Forsythe, Tom A. Rietz, "The Iowa Electronic Market" in Dean Paxson and Douglas Wood, eds., *Blackwell Encyclopedic Dictionary of Finance* (Oxford: Blackwell Press,1997).

[77] Examples can be found at the Web sites of Common Knowledge Markets, HedgeStreet, Goldman Sachs and Deutsche Bank, NewsFutures, and TradeSports.

[78] Stephen Figlewski, "Subjective Information and Market Efficiency in a Betting Model," *Journal of Political Economy* 87 (1979): 75–88.

[79] Richard Roll, "Orange Juice and Weather," *The American Economic Review* 74, no. 5 (1984): 861–80.

[80] Michael 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* 9, no. 4 (September 2003): 453–79.

[81] Charles Plott and Kay-Yut Chen, "Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem" (working paper 1131, Caltech Social Science, Pasadena, CA, 2002).

[82] Justin Wolfers and Eric Zitzewitz, "Prediction Markets," *The Journal of Economic Perspectives* 18,2 (Spring 2004): 107-126

[83] Robert Aumann, "Agreeing to Disagree," *The Annals of Statistics* 4 (1976): 1236–1239.

[84] Common knowledge means that each knows the other's beliefs, each knows the other knows that each knows the other's beliefs, and so on. This will be true if, for example, they each announce their beliefs when they are physically in the presence of each other.

[85] John Geanakopolis and Heracles Polemarchakis, "We Can't Disagree Forever," *Journal of Economic Theory* 28 (1982): 192–200.

[86] Richard McKelvey and Talbot Page, "Common Knowledge, Consensus and Aggregate Information," *Econometrica* 54, no. 1 (1986): 109–127.

[87] Richard McKelvey and Talbot Page, "Common Knowledge, Consensus and Aggregate Information," *Econometrica* 54, no. 1 (1986): 109–127.

[88] Scoring rules are of real interest in this paper and will be covered in detail below. They are one way to provide monetary incentives to someone to accurately reveal their probability beliefs.

[89] Each of these conclusions was somewhat qualified, as we will see in a little while.

[90] See, for example, Jerry Green, "Temporary General Equilibrium in a Sequential Trading Model with Spot and Future Transactions," *Econometrica* 41 (1973): 1103–23; Sandford Grossman, "The Existence of Futures Markets, Noisy Rational Expectations, and Informational Externalities," *The Review of Economic Studies* 44 (1977): 431–49; Robert Lucas, "Expectations and the Neutrality of Money," *Journal of Economic Theory* 4 (1972): 103–24; Beth Allen, "General Equilibrium with Rational Expectations," in Werner Hildenbrand and Andreu Mas-Colell, editors, *Contributions to Mathematical Economics in Honor of Gerard Debreu*. North-Holland, Amsterdam: North-Holland. 1986

[91] See, for example, Kenneth Arrow, "The Role of Securities in the Optimal Allocation of Risk Bearing," *The Review of Economic Studies* 31 (1964): 91–96; Gerard Debreu, *Theory of Value* (New York: John Wiley and Sons, 1959).

[92] In the IEM example above, there are 3 mutually exclusive states: Bush wins, Dukakis wins, and Other wins. It is a complete set of securities.

[93] Robert Forsythe, Thomas Palfrey, and Charles Plott, *"Asset Valuation in an Experimental Market,"* *Econometrica* 50 (1982): 537–68.

[94] Richard D. McKelvey and Peter C. Ordeshook, "Sequential Elections with Limited Information," *American Journal of Political Science* 29 (1985): 480–512.

[95] Charles R. Plott and Shyam Sunder, "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," *Econometrica* 56, no. 5 (1998): 1085–118.

[96] Sanford J. Grossman and Joseph Stiglitz, "Information and Competitive Price Systems," *The American Economic Review* 66, no. 2 (1976): 246–253.

[97] If they don't know this then there is no way for them to make inferences from the prices.

[98] Robin Hanson and Ryan Oprea, "Manipulators Increase Information Market Accuracy" (working paper, Department of Economics, George Mason University, Fairfax, Va., 2004) provides, on the other hand, a financial trading model in which traders actually acquire more information when a potential manipulator is present.

[99] Paul Milgrom and Nancy Stokey, "Information, Trade, and Common Knowledge," *Journal of Economic Theory* 26 (1982): 17–27.

[100] If 1 saw (H,H) or (H,T) 1 would not be willing to part with their assets for any price less than $0.50.

[101] It would be less than if person 1 had seen (0H, 2T).

[102] James S. Jordan, "The Generic Existence of Rational Expectations Equilibrium in the Higher Dimensional Case," *Journal of Economic Theory* 26 (1982): 224–243.

[103] Charles R. Plott and Shyam Sunder, "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," *Econometrica* 56, no. 5 (1998): 1085–118.

[104] Richard McKelvey and Talbot Page, "Common Knowledge, Consensus and Aggregate Information," *Econometrica* 54, no. 1 (1986): 109–127.

[105] Incomplete updating is a problem for all information aggregation procedures including polls and committees.

[106] Robert W. Hahn and Paul C. Tetlock, "Using Information Markets to Improve Policy", Harvard Law and Public Policy, Forthcoming http://ssm.com/abstract=598882

[107] IEM has a particularly clever way to implement this market, as they did for the Google IPO. Two securities are issued: one that pays $0 if the policy is not implemented, and $(ax) if the policy is implemented and the average score is x (where a(max x) = $1 and max x is the highest score possible. The other pays $1 if the policy is not implemented and pays $(1-ax) if the policy is implemented. Note that no matter what happens the portfolio value of 1 of each security is $1. The market maker can offer to buy or sell this portfolio at the price of $1 risk-free.

[108] It is also true that if the price of security 2 were .10 higher than the price of security1 so that a predicted increase of .1a were being made, the policy maker will never know whether more could have been done for less.

[109] As an indication of how quickly this can get out of hand in real applications consider the Policy Analysis Market (PAM) proposed by Net Exchange for DARPA. It would have involved 5 quarterly indices for each or 8 countries (plus 5 other indices) for 4 quarters. Each market was whether the index went up or down. There were to be initially 1000 traders growing eventually to 10,000. A complete set of contracts would have required $2^{180} = 1.54E+50$ contracts.

[110] Even if management did not reopen the markets after their inquiry, these folk lose since the Aug securities end up out of the money. Alternatively, if management gives everybody their money back, the traders will be not very interested the next time information markets are tried. In all cases, there is a serious problem.

[111] This does not eliminate another trader problem that can arise from management reaction to the information in the market – loss of one's job. If the individual who signals that A_{new} is in trouble works exclusively on A_{new} then cancellation of that can be worse for them than the gains in the information market. They will suppress any information that is bad for them – especially if they are the unique holder of the information.

[112] It is a well-known theorem in mechanism design (see Leo Hurwicz "On informationally decentralized systems" in Decision and Organization: A volume in honor of Jacob Marschak, ed. Roy Radner and C. Bart McGuire 1972 pp.297-336, Amsterdam: North-Holland Publishing Co.) \as long as there are a finite number of traders, each trader – whether an informational monopolist or not – has some market power and can, through misrepresentation of preferences, cause the outcome to move in a direction she would prefer. Of course the incentive to do this usually decreases as the number of traders increases.

[113] Don't ask me how they would know this; this is just an example. We could change the example prediction to "a whaler will land in Boston on 12/25/1850" and allow this person to have a very powerful telescope.

[114] This is a form of Bertrand competition. Charlie Plott reports this type of phenomenon in many of his experiments where the two insiders compete with each other in the order queue until they drive the price to 1 and no one gains anything. They act as if they are competitive even if they aren't. He argues that this is an important mechanism in getting private information out into the market place.

[115] Glenn W. Brier, "Verification of Forecasts Expressed in Terms of Probability," *Monthly Weather Review* 75 (1950): 1–3.

[116] Leonard J. Savage, "Elicitation of Personal Probabilities and Expectations," *Journal of the American Statistical Association* 46 (1971): 783–801.

[117] I. J. Good, "Rational Decisions", Journal of the Royal Statistical Society, 14 (1952): 107-114.

[118] Risk-averters will report a probability density that is flatter than the truth.

[119] Robin Hanson, "Combinatorial Information Market Design," *Information Systems Frontiers* 5, no. 1 (2003): 105–119.

[120] I am now reporting on material that is in Robin Hanson, Takashi Ishikida, and J Ohn Ledyard, "An Experimental Test of Combinatorial Information Markets" (manuscript, Caltech, Pasadena, CA, 2005). Hanson was the prime designer of the experimental environments. Ishikida was the prime implementer.

[121] Charles R. Plott and Shyam Sunder, "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," *Econometrica* 56, no. 5 (1998): 1085–118.

[122] Richard McKelvey and Talbot Page, "Public and Private Information: An Experimental Study of Information Pooling," *Econometrica* 58, no. 6 (1990): 1321–39.

[123] We were analyzing designs for the Policy Analysis Market. That is a really thin environment and we did not expect standard market designs to do well.

[124] Peter Bossaerts and Charles Plott, "Basic Principles of Asset Pricing Theory: Evidence from Large-Scale Experimental Financial Markets," *Review of Finance* 8 (2004): 1–35.

[125] Peter Bossaerts and Charles Plott, "The CAPM in Thin Experimental Financial Markets," *Journal of Economics Dynamics and Control* 26 (2002): 1093–1112.

[126] Peter Bossaerts, Leslie Fine, and John Ledyard, "Inducing Liquidity in Thin Financial Markets Through Combined-Value Trading Mechanisms," *European Economic Review* 46 (2002): 1671–95.

[127] For purposes of comparison, the Policy Analysis Market was intended to have 320 basic assets: one for each combination of 8 countries, 5 indices (economic growth, political activity, military activity, US aid \$, US military activity) each with 3 outcomes (up, down, constant), and 4 time periods (the next 4 quarters of the year). If one were to create a complete set of contracts one would need 2^{320} or $2*(10^{96})$ contracts.

[128] It would be perhaps more appropriate to say "we tried to induce a common prior". We handed out instructions, available on request, with much detail. We repeated the markets within each session, and we brought back subjects for further sessions. But one never knows for sure what the subjects really have in their heads.

[129] There are many possible measures one might use. This measure is reasonably standard in the literature so we used it here. It does have a problem if probabilities are close to zero but we ignore that here.

[130] See figures 1 and 2 at the end of the paper.

[131] In the 3 event environment experiments there were 6 different priors and full information posteriors. And, since there were 3 subjects in each there were 18 different individual posteriors. In the 8 event environment there were also 6 priors and full information posteriors. Since there were 6 subjects in each there were 36 different individual posteriors.

[132] The subject pool was Caltech undergrads. I do not know what would happen with a different pool.

[133] One paid $1 if and only if the event = 1. The other paid $1 if and only if the event = 0.

[134] Peter Bossaerts, Leslie Fine, and John Ledyard, "Inducing Liquidity in Thin Financial Markets Through Combined-Value Trading Mechanisms," *European Economic Review* 46 (2002): 1671–95.

[135] In a call market, as opposed to say Marketscape where bids, offers, and acceptances occur asynchronously, all bids and offers are collected over some pre-specified period of time. At the end of that period, the bids and offers are matched in a way that maximizes the total gains from trade. This allows time to accumulate trades so that, at the call, the market is thicker.

[136] There is a lot of research going on in this area, significantly aided by experimental methods. See Charles Plott, Jorgen Wit, and Winston Yang, "Pari-mutuel Betting Markets as Information Aggregation Devices: Experimental Results," *Economic Theory* 22 (2003): 311–51; Kay-Yut Chen, Leslie Fine, and Bernardo Huberman, "Eliminating Public Knowledge Biases in Information-Aggregation Mechanisms," *Management Science* 7 (2004): 983–94; David Pennock "A Dynamic pari-mutuel market for hedging, wagering, and information aggregation" ACM Conference on Electronic Commerce, May 2004