



Combinatorial Information Market Design

Robin Hanson

Department of Economics, George Mason University, MSN 1D3,
Carow Hall, Fairfax VA 22030, USA

E-mail: rhanson@gmu.edu (<http://hanson.gmu.edu>)

Abstract. Information markets are markets created to aggregate information. Such markets usually estimate a probability distribution over the values of certain variables, via bets on those values. Combinatorial information markets would aggregate information on the entire joint probability distribution over many variables, by allowing bets on all variable value combinations. To achieve this, we want to overcome the thin market and irrational participation problems that plague standard information markets. Scoring rules avoid these problems, but instead suffer from opinion pooling problems in the thick market case. Market scoring rules avoid all these problems, by becoming automated market makers in the thick market case and simple scoring rules in the thin market case. Logarithmic versions have cost and modularity advantages. After introducing market scoring rules, we consider several design issues, including how to represent variables to support both conditional and unconditional estimates, how to avoid becoming a money pump via errors in calculating probabilities, and how to ensure that users can cover their bets, without needlessly preventing them from using previous bets as collateral for future bets.

Key Words. market maker, scoring rule, information aggregation

Introduction

Financial markets, such as stock or commodities futures markets, are usually created to allow traders to hedge risks. Once created, however, such markets also attract speculators who seek to “buy low and sell high.” And while speculators can profit if they can find any way to predict future prices from current prices, a side effect of their trades is to eliminate such profit opportunities. As a result, speculative market prices are often quite accurate estimates of future prices, aggregating a great deal of available information (Lo, 1997). For example, orange juice commodity futures markets improve on government weather forecasts (Roll, 1984), and horse races are better predicted by betting markets than by professional handicappers (Figlewski, 1979).

This ability of existing markets to aggregate information has recently inspired several “information markets,” which were created not to entertain or to hedge risks but to aggregate information on particular topics of interest. For example, when compared to concurrent major opinion polls on U.S. presidential elections, the Iowa Electronic Market forecasts were more accurate 451 out of 596 times (Berg, Nelson, and Rietz, 2001). When compared to official Hewlett-Packard forecasts of printer sales, internal corporate markets were more accurate 6 out of 8 times, even though the official forecasts were made after the markets closed and with knowledge of the market prices (Chen and Plott, 1998). Play money markets predicted the 2000 Oscar winners better than 4 out of 5 columnists who made concurrent forecasts (Pennock, Giles, and Nielsen, 2001).

Anti-gambling laws are probably now the main barrier to wider use of real-money information markets. Speculative markets are in general illegal, though regulators allow exceptions for certain purposes. For example, insurance markets allow individuals to hedge idiosyncratic risks such as fire and sickness, capital markets allow firm managers to hedge uncertainties in firm profits, commodity futures markets allow individuals and firms to hedge risks from common economic trends, and sports betting markets enhance entertainment from watching sporting events. For each of these purposes, there are regulatory bodies devoted to approving some markets for that purpose. But since there are no regulatory bodies devoted to approving markets which serve the purpose of aggregating information, information markets are now only allowed by accident.

Even when information markets are legal, however, they often fall victim to the *thin market* problem. To trade, traders must coordinate on the assets they will trade, even though there are an astronomical number of possible events for which assets might be defined.

Traders must also coordinate on when they will agree to trade, since offers that wait long before being accepted suffer adverse selection from new public information. Thus traders usually only expect substantial trading activity in a certain limited set of assets, and have little reason to offer to trade other assets. Trader coordination can be aided somewhat by call markets, where all trades happens at pre-defined moments, and by combinatorial matching markets, which search for combinations of offers that can be matched. But such aid can only go so far. Consider the case where a single person knows something about an event, and everyone else knows that they know nothing about that event. In this case, standard information markets based on that event simply cannot acquire this person's information.

Standard information markets also suffer from an *irrational participation* problem. Once rational agents have hedged their risks regarding the events covered by an information market, they should not want to trade with each other. Even if they have substantial private information not reflected in previous market prices, agents who care only about what they can buy with the assets they might gain must realize that the gains of some can only come from the losses of others (Milgrom and Stokey, 1982; Aumann, 1976). Now real markets do show surprisingly high levels of speculative activity, perhaps due to widespread irrationality, or to people using bets to signal their expertise to observers (as in "bar bets"). But it may not be possible to induce much more such activity, in order to support new information markets without taking away from other markets.

Long before there were information markets, however, there were scoring rules, designed to obtain information by eliciting probability distributions from individuals. Scoring rules do not suffer from irrational participation or thin market problems; they have no trouble inducing one person to reveal information. They instead suffer from a thick market problem, namely how to produce a single consensus estimate when different people give differing estimates. This paper considers a new technology, *market scoring rules*, which combines the advantages of information markets and scoring rules. When dealing with a single person, a market scoring rule becomes a simple scoring rule, yet when dealing with a group of people, it becomes an automated market maker supporting trades between those people. Thus a market scoring rule basically solves both the thin market and the irrational participation problems with information markets, as well as the thick market problem of scoring rules.

Given some base set of after-the-fact-verifiable variables, given patrons willing to pay a modest per-variable cost to induce information revelation on those variables, and setting aside computational limitations, market scoring rules allow people to reveal what they know by trading on *any* event defined in terms of combinations of those base variable values. So, for example, given N binary variables, people could trade on any of the 2^N possible states or 2^{2^N} possible events. A market scoring rule always has a complete consensus probability distribution over the entire state space, a distribution that anyone can change any part of at any time. And risk-neutral agents should expect to profit from such changes whenever their own beliefs differ in any way from this consensus distribution.

Of course there are in fact limits on what we can compute. So after reviewing existing information aggregation technologies and the new technology of market scoring rules, this paper will go on consider several implementation issues raised by this new technology.

Previous Technologies of Information Aggregation

The general task we consider in this paper is to induce people to acquire and reveal information relevant to estimating certain random variables. These variables may include natural outcomes like earthquakes or the weather, economic statistics like GDP, and political outcomes like elections, wars and revolutions. We may also want conditional estimates, such as the chance of war given that GDP falls, and especially decision-contingent estimates such as of the chance of war given that we elect a particular person for president, the stock price of a company given that we dump the current CEO, or a patient's lifespan given that she adopts a particular course of treatment. Such decision-contingent estimates can give us relatively direct advice about what choices to make (Hanson, 1999).

We will not assume that we know which people have which areas of expertise, but we will assume that we can someday verify the variable values with little controversy "after the fact." Ideally we would like a single complete consistent probability distribution over the entire state space of all variable value combinations, embodying everything that everyone knows, as well as everything they could learn with a modest effort. In practice, we may have to accept compromises.

Scoring Rules

To date, there have been two main approaches to the information aggregation task we have outlined, scoring rules and information markets. In order to elicit a risk-neutral person's probability distribution $\mathbf{p} = \{p_i\}_i$ over all states i (where $\sum_i p_i = 1$), one can ask him to report a distribution $\mathbf{r} = \{r_i\}_i$, and then if the true state turns out to be i give him a reward c_i according to a *proper scoring rule* $c_i = s_i(\mathbf{r})$ that satisfies an incentive compatibility constraint¹

$$\mathbf{p} = \operatorname{argmax}_{\mathbf{r}} \sum_i p_i s_i(\mathbf{r})$$

and a rational participation constraint

$$0 \leq \sum_i p_i s_i(\mathbf{p}).$$

(The reward for not participating is set here to zero.) Given these assumptions and constraints, the person will want to set $\mathbf{r} = \mathbf{p}$. Scoring rules also give agents incentives to acquire information they would not otherwise possess (Clemen, 2002). These incentives come at the expense of a *patron* who agrees to pay the reward x_i .²

In 1950 Brier introduced the quadratic scoring rule (Brier, 1950)

$$s_i = a_i + b \left(2r_i - \sum_k r_k^2 \right),$$

and in 1952 Good followed with the logarithmic scoring rule (Good, 1952)

$$s_i = a_i + b \log(r_i).$$

The logarithmic rule is the only rule that can be used both to reward an agent and to evaluate his performance according to standard statistical likelihood methods (Winkler, 1969). Scoring rules have long been used in weather forecasting (Murphy and Winkler, 1984), economic forecasting (O'Carroll, 1977), student test scoring, economics experiments, risk analysis (DeWispelare, Herren, and Clemen, 1995), and the engineering of intelligent computer systems (Druzdzal and van der Gaag, 1995).

While scoring rules in principle solve incentive problems with eliciting probabilities, many other problems have appeared in practice. For example, people are often hesitant to state probability numbers, so probability wheels and standard word menus are often used. Also, corrections are often made for human cognitive biases. In theory, risk-aversion, uncommon priors, and

state-dependent utility can also make it hard to infer the information people have from the estimates they make (Kadane and Winkler, 1988), though in practice these are not usually considered to be big problems. When they are problems, in theory risk aversion can be dealt with by paying in lottery tickets (Smith, 1965; Savage, 1971), and uncommon priors and state-dependent utility can be dealt with by having people first play a certain lottery insurance game (Hanson, 2002a).

Another problem is that, in general, the number of states is exponential in the number of variables, making it hard to elicit and compute with probability distributions with many variables. This problem is often dealt with by creating a sparse dependence network, such as a Bayes net (Pearl, 1988; Jensen, 2001; Pennock and Wellman, 2000), relating the variables. A dependence network among the variables says that each variable is, given its neighbor variables, conditionally independent of all other variables. Thus by eliciting a sparse network of dependencies, one can drastically reduce computational complexity (more on this below). While little attention is usually paid to incentive compatibility when eliciting such network structures, they often seem uncontroversial.

One big problem with scoring rules remains largely unsolved, however. When different people are asked to estimate the same random variable, they can and often do give different estimates. Yet what we really want is a single estimate that aggregates information from different people. Unfortunately, the literature on "pooling opinions," i.e., constructing a single pooled probability distribution from a set of individual distributions, is mostly discouraging (Genest and Zidek, 1986). For example, it turns out that any two of the following three apparently reasonable conditions imply a dictator, i.e., that the pooled distribution is equal to one of the individual distributions:

1. If two events are independent in each individual distribution, they are independent in the common distribution.
2. The pooling procedure commutes with updating the distributions with new information.³
3. The pooling procedure commutes with coarsening the state space (e.g., dropping a variable).

Since pooling opinions well is hard, the usual practice in building large probability distributions is to choose a single expert to specify parameters for each variable. For example, a single person might be assigned to estimate the weather in a given geographic area.

Information Markets

Information markets can overcome many of the limitations of scoring rules as information aggregation mechanisms. Like scoring rules, information markets give people incentives to be honest, because each contributor must “put his money where his mouths is.” In such markets, people typically trade assets of the form “Pays \$1 if event A happens,” and the market price for such assets is interpreted as a consensus estimate of $p(A)$. The approaches described above for dealing with number shyness, risk-aversion, and state-dependent utility can also be used with information markets.

In contrast to scoring rules, however, information markets can combine potentially diverse opinions into a single consistent probability distribution. The information pooling problem with scoring rules arises because a given probability distribution over some limited set of variables can arise from many different information sets, and so one cannot determine an agent’s information from his probability estimates. Rational agents who are repeatedly made aware of each other’s estimates, however, should converge to identical estimates, at least given common priors (Geanakoplos and Polemarchakis, 1982; Hanson, 1998, 2003). Information markets can use such repeated interaction to produce common estimates that combine available information, avoiding the opinion pooling problem. No knowledge of who is more expert on what topics is required, and corrections for cognitive bias can even be left to market participants, who can profit by making such corrections. Market traders self-select to focus on the topics where they believe they are most expert, and those who are mistaken about their areas of expertise are punished financially.

As was discussed in the introduction, however, standard information markets suffer from both irrational participation and thin market problems. Once rational agents who care only about what they can buy with the assets they might win have hedged their risks, they should not want to make speculative trades with each other, even when they know things that others do not (Milgrom and Stokey, 1982; Aumann, 1976). And because traders must coordinate on the assets to trade and on when to agree to trade, traders usually only expect substantial trading activity in a certain limited set of assets, and have little reason to offer to trade other assets. For example, when one person knows something about a certain variable, and everyone else knows that

they know little about that variable, standard information markets trading assets based on that variable will not reveal what this person knows.

Market Scoring Rules

Fig. 1 illustrates some performance issues with scoring rules and information markets. When information markets are thick, their accuracy should increase with the number of traders who focus on each asset. As the number of traders per asset falls near one, however, the thin market problem should eliminate trading activity. Scoring rules, in contrast, can produce estimates in such cases, though one expects estimate accuracy to fall with the number of estimates one asks of any one person. When using scoring rules in the thick market case, however, where many people are asked to estimate the same parameter, opinion pooling problems should make them less accurate than information markets.

Market scoring rules can combine the advantages of scoring rules and standard information markets. They should thus produce an accuracy like that of information markets when many people make the same kind of estimates, and like that of scoring rules when only one person makes each estimate. Market scoring rules are, in essence, sequentially shared scoring rules. Anyone can at any time choose to use such a rule, i.e., be paid according to that rule, if they agree to pay off the last person who used the rule. So if only one person uses a market scoring rule, it in effect becomes a simple scoring rule. But if many people use that market




Fig. 1. Comparing mechanisms.

scoring rule, it in effect becomes an automated market maker facilitating trades between those people. And since each user pays off the previous user, the market scoring rule's patron need only pay off the last user.

More formally, a market scoring rule always has a current probability distribution \mathbf{p} , which is equal to the last report that anyone has made to this rule. Anyone can at any time inspect this distribution, or change any part of it by making a new report. If someone chooses to give a report \mathbf{r}^t at time t , and then the actual state eventually turns out to be i , this person will be paid a reward

$$c_i = \Delta s_i(\mathbf{r}^t, \mathbf{r}^{t-1}) = s_i(\mathbf{r}^t) - s_i(\mathbf{r}^{t-1}),$$

where $s_i(\mathbf{r})$ is some proper scoring rule, and \mathbf{r}^{t-1} is the last report made. Since this amount c_i can be negative, the person must show an ability to pay this if needed, such as by depositing or escrowing⁴ an amount— $\min_i c_i$.

Since the person giving report \mathbf{r}^t cannot change the previous report \mathbf{r}^{t-1} , he maximizes the expected value of $\Delta s_i(\mathbf{r}^t, \mathbf{r}^{t-1})$ by maximizing the expected value of $s_i(\mathbf{r}^t)$, and so wants to honestly report his beliefs here whenever he would for a simple scoring rule.⁵ Whenever someone's beliefs differ from the current distribution \mathbf{p} , he expects to profit on average from making a report, for the same reason that someone who had mistakenly made the wrong report would want to correct such a mistake before it became official.

When a person changes a market scoring rule distribution \mathbf{p} , he must raise some probability values p_i and lower others. If the true state turns out to be one where he raised the probability, he will gain, but if the true state turns out to be one where he lowered the probability, he will lose. This person is thus in essence making a bet with the market scoring rule regarding the true state. Since one could choose to change the distribution \mathbf{p} by only a tiny amount, and one could undo such a bet by reversing the change, these tiny bets are all “fair bets” at the probabilities \mathbf{p} .

Since people are free to change any combination of probabilities in \mathbf{p} , and since bets are what make those probabilities change, a market scoring rule is in essence an automated inventory-based market maker who stands ready to make *any* tiny fair bets at its current probabilities.⁶ This market maker also stands ready to make any large bets that can be constructed from tiny bets. The only catch is that the prices \mathbf{p} change as tiny bets are made.

We can summarize all this by saying that each market scoring rule in essence has a “net sales so far” vector $\mathbf{s} = \{s_i\}_i$, where each s_i says how many units have been sold of assets of the form “Pays \$1 if the state is i .” The current unit price for a tiny amount of such an asset is p_i , and these prices change according to a price function $\mathbf{p}(\mathbf{s})$, which is in essence a generalized inverse of the scoring rule function $\mathbf{s}(\mathbf{p})$. For example, for the logarithmic scoring rule $s_i(\mathbf{p}) = a_i + b \log(p_i)$, the price function is the exponential

$$p_i(\mathbf{s}) = \frac{e^{(s_i - a_i)/b}}{\sum_k e^{(s_k - a_k)/b}}.$$

To find the price for any non-tiny bundle of assets, one must integrate the price as it changes across sales of tiny bundles. If the sales history is $\mathbf{s} = \mathbf{h}(t)$, and this new sale starts at \underline{t} and ends at \bar{t} , then the total amount of this sale is $\mathbf{h}(\bar{t}) - \mathbf{h}(\underline{t})$, and the total price for this sale is

$$\int_{\underline{t}}^{\bar{t}} \sum_i p_i(\mathbf{h}(t)) h'_i(t) dt.$$

Since $\mathbf{s}(\mathbf{p})$ is a function, this integral is independent of the sales path between $\mathbf{h}(\underline{t})$ and $\mathbf{h}(\bar{t})$.

In many existing markets, all trades are made with one or a few central market makers. These actors always have public offers to buy or to sell, and update these prices in response to trades. Human market makers have been found to do as well as the standard double auction market form in aggregating information (Krahnen and Weber, 1999). Some markets, such as the Hollywood Stock Exchange (www.hsx.com), use automated market makers to fill this role. Typically, each asset has a single automated market maker which deals only in that asset. A market scoring rule is like this, except that it is a single automated market maker that deals in all of the assets associated with an entire state space. So, for example, given N binary variables, a single market scoring rule can make trades on any of the 2^N possible states, or any of the 2^{2^N} possible events (i.e., sets of states).

As with ordinary market makers, abstractions can be devised to allow users to more easily follow various standard trading scenarios. For example, a “limit order” can specify an amount and a boundary in \mathbf{p} space, and mean that one is willing to spend up to a given amount to keep \mathbf{p} from crossing this boundary.

Costs

Since the patron of a market scoring rule must only pay the last user, his total expected loss depends only on how informative that last report is, relative to the initial report the rule was started with. This loss is otherwise independent of the number of users. To minimize his expected loss, the patron should set the initial report to his initial beliefs \mathbf{q} , in which case his worst case expected loss is

$$S = \sum_i q_i (s_i(\mathbf{1}_i) - s_i(\mathbf{q})),$$

where $(\mathbf{1}_i)_j = 1$ when $i = j$ and zero otherwise. This worst case applies if in each state i , the last report is sure that this is the true state. For the logarithmic scoring rule, this worst case expected loss is (b times) the entropy of the distribution \mathbf{q} .

How much more does it cost to have a single market maker which covers the entire state space, relative to just having a market maker for each variable? A market scoring rule that covered only a single variable with V possible values would be an automated market maker for all V assets of the form “Pays \$1 if this variable has value v .” For a logarithmic rule, the worst case loss would be (b times) the entropy of the marginal distribution $\{q_v\}_v$ over the values v . In general the sum of the entropies of the marginal distributions of each variable is no less, and usually more, than the entropy of the joint distribution over all variables.⁷ Thus (setting aside computational costs) it costs no more to fund an automated market maker to trade in the entire state space than it costs to fund automated market makers limited to each variable.

Logarithmic market scoring rules also have modularity advantages. It seems good to minimize the extent to which people making bets on some dimensions of the probability distribution regularly cause unintended changes to other unrelated dimensions of the distribution. For example, when one believes that $p(A | B) > q$, one expects to gain by giving up q units of “Pays \$1 if B” in trade for one unit of “Pays \$1 if A and B.” And by making this trade one takes no risk regarding whether the event B is true, since one trades some assets which assume B for other assets which also assume B . While it seems appropriate for such trades to change the value of $p(A | B)$ in the distribution \mathbf{p} , it does not seem appropriate for them to change the value of $p(B)$. Requiring this condition to hold for all such trades, however, is equiva-

lent to requiring the logarithmic market scoring rule (Hanson, 2002b). Given such a conditional trade, the logarithmic rule preserves not only $p(B)$, but for any event C it also preserves the conditional probabilities $p(C | \text{not } B)$, $p(C | B \text{ and } A)$, and $p(C | B \text{ and not } A)$, and the variable independence relations $I(\mathcal{A}, \mathcal{B}, \mathcal{C})$, $I(\mathcal{B}, \mathcal{A}, \mathcal{C})$, $I(\mathcal{C}, \mathcal{B}, \mathcal{A})$, and $I(\mathcal{C}, \mathcal{A}, \mathcal{B})$, for A a value of \mathcal{A} , B a value of \mathcal{B} , and C a value of \mathcal{C} . (For variables $\mathcal{A}, \mathcal{B}, \mathcal{C}$, we say $I(\mathcal{A}, \mathcal{B}, \mathcal{C})$ holds for p when $p(A_i | B_j C_k) = p(A_i | B_j)$ for all values A_i of \mathcal{A} , B_j of \mathcal{B} , and C_k of \mathcal{C} .)

Representing Variables

Market scoring rules were described above in terms of probability distributions over states i defined in terms of combinations of variable values v . But how should we define those variable values?

Imagine that we have two real-valued random variables, x and y , such as GDP change and immigration. One user might want to express his opinion on how much he expects GDP to change, while another user might want to express his opinion that immigration will be high if GDP rises somewhat. If we want to let both these users express their opinions, but we also want to describe GDP in terms of a small number of assets, to limit computational complexity, what should those assets be?

A user who wants to use GDP as a condition, as in “if GDP rises somewhat,” might prefer “digital option” assets defined in terms of whether x is in a given region, such as “Pays \$1 if $x < x_1$,” “Pays \$1 if $x_1 \leq x < x_2$,” and “Pays \$1 if $x_2 \leq x$.” A user who wants to express his opinion on the expected value of GDP change, however, might prefer “linear” assets such as “Pays $\$ \hat{x}$,” and “Pays $\$(1 - \hat{x})$,” where we have defined a rescaled variable

$$\hat{x} = \max \left(0, \min \left(1, \frac{x - \underline{x}}{\bar{x} - \underline{x}} \right) \right),$$

which is zero up to $x = \underline{x}$, is one above $x = \bar{x}$, and moves linearly with x in between.

While scoring rule payoffs were defined above only when a given state i is found to be true with certainty, we can easily reinterpret that framework to include such linear assets. Scoring rules can instead payoff according to any final assignment $\mathbf{a} = \{a_i\}_i$ where $\sum_i a_i = 1$. In essence, each asset of the form “Pays \$1 if the ‘state’ is i ” pays $\$a_i$. In this case, people should want

to report their expectations for \mathbf{a} , as in $r_i = E_{\mathbf{p}}[a_i]$. When individual variables are expressed linearly, such as with value assignments \hat{x} and $1 - \hat{x}$, value assignments for variable combinations can be products, as in $\hat{x}\hat{y}$, $(1 - \hat{x})\hat{y}$, $\hat{x}(1 - \hat{y})$, and $(1 - \hat{x})(1 - \hat{y})$. That is, when the real-valued random variable j has value z_j , let a corresponding discrete variable value v have assignment $a_{jv}(z_j)$, so that $\sum_v a_{jv}(z) = 1$ for all z . We can then let the assignment of state i be

$$a_i(\mathbf{z}) = \prod_j a_{jv_j(i)}(z_j),$$

where $v_j(i)$ is the value of discrete variable j in state i , and $\mathbf{z} = \{z_i\}_i$.

Given linear variable value assignments, a user could in a sense bet that immigration will be high given that GDP change is high, i.e., that y will be high, given that x is high, by giving up q units of “Pays $\$x$ ” for one unit of “Pays $\$x\hat{y}$ ” whenever he believed that $q > E_{\mathbf{p}}[\hat{x}\hat{y}]/E[\hat{x}]$. For most people, however, this seems an unnatural way to describe their conditional expectations. To have a single asset representation that can let people express opinions on the expected change in GDP, and let people express opinions which use “if GDP rises somewhat” as a condition, we might use “sawtooth” assets, such as illustrated in Fig. 2. Each of a set of $K + 1$ points x_k would have a matching asset with a final assignment a_k , so that $\sum_k a_k = 1$ no matter what happens. When $x < x_0$, then $a_0 = 1$ and the other a_i are zero, and if $x > x_K$ then $a_K = 1$ and the others are zero. If $x = x_k$ for any k then that $a_k = 1$ and the others are zero. Otherwise the assignment varies linearly between these values. That is, for any x there is some k where $x_k < x < x_{k+1}$, and we set $a_j = (x_{k+1} - x)/(x_{k+1} - x_k)$, $a_{k+1} = 1 - a_k$, and all others to zero.

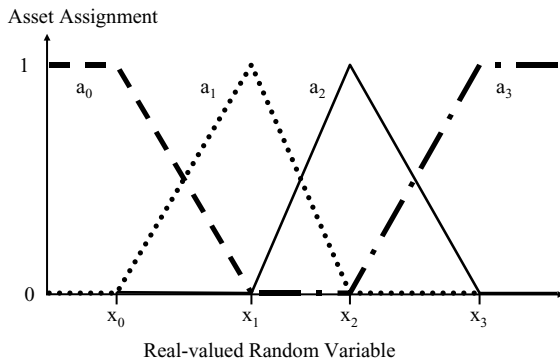


Fig. 2. Sawtooth asset representation.

A sawtooth asset only has value when the real variable x is “near” a particular value x_k . Thus such assets can be used as a condition, as in “if GDP rises somewhat.” Yet linear combinations of sawtooth assets can also reproduce assets like “Pays $\$x$,” and they can do this for \underline{x} and \bar{x} equal to any pair of values x_k . Thus if the GDP change and immigration variables are both represented by sawtooth assets, users can naturally bet on both the expected value of GDP change, and on the expected value of immigration, given that GDP change is somewhat high.

Note that the state space for a variable can be refined on the fly when real variables are represented either by sawtooth assets or digital options. Both representations rely on a set of cutpoints x_k , and new cutpoints can be inserted at anytime, as it is straightforward to translate assets defined in terms of the old cutpoints into assets defined in terms of the new cutpoints. However, since adding cutpoints makes it possible for the last report to be more informative about the actual state of the world, adding cutpoints can increase the expected cost to the patron to pay for a market scoring rule. New variables can also be easily added on the fly, if a patron is willing to pay for the extra cost.

Computational Issues

When there are only a dozen or so variables, each with a handful of values, one can exactly compute a market scoring rule. A central data structure can store the current probability p_i explicitly for every state, and probabilities of events can be directly calculated via sums, as in $p(A) = \sum_{i \in A} p_i$. A similar data structure per user can describe his assets, by giving his payoff in each state. That is, for each user one stores a number c_i for each state, which describes an asset of the form “Pays $\$c$ if state i .” If a user acquires an asset “Pays $\$c$ if event A happens,” this asset can be stored by increasing c_i by the amount c for each state $i \in A$. This representation of assets allows users to maximally reuse assets they acquired in previous bets to support future bets.

When value v of variable α has probability p_v , and then is verified to have assignment a_v , the entire distribution can be updated to $p'_i = p_i a_v / p_v$ for each state i where $v_\alpha(i) = v$. When all the values of this variable are assigned, the state space can be coarsened to eliminate this variable, by merging all states which differ only in their value of this variable. In this merging

process, the probability of the new state is the sum of the probabilities of the old states, and the assets holdings for each person in the new state is an average of their asset holdings in the old states, weighted by the assignment a_v .

Consider a logarithmic market scoring rule based on $s_i = \beta \log_2(p_i)$, with β an integer. The patron of this rule can avoid roundoff error losses by storing the amounts c_i each user holds of asset “Pays \$1 if state i ” as exponentials $\theta_i = 2^{c_i/\beta}$. So when a user makes a “cash” deposit d that is a multiple of β , we can set all his initial holdings to the integer $\theta_i = 2^{d/\beta}$. If probabilities p_i are always rational numbers, then whenever a user starts by holding θ_i , and then changes a probability from p_i to p'_i , we can exactly compute his new holding to be the rational $\theta'_i = \theta_i(p'_i/p_i)$. And if we never allow a user to make a change that would produce any $\theta_i < 1$ (in a state where $p_i > 0$), then users can never go bankrupt. We can also allow a user to withdraw cash by specifying a $w \leq \min_i \theta_i$ that is a multiple of β . We just set $\theta_i = \theta_i/w$, and give him cash $\beta \log_2 w$.⁸ (A sample implementation, in CommonLisp, is available at <http://hanson.gmu.edu/mktscore-prototype.html>)

Beyond a dozen or so variables, each with a handful of values, the above approach quickly becomes infeasible. After all, when N variables each have V possible values, there are V^N possible states, and 2^{V^N} possible events. When V^N is more than a billion or so (e.g., thirty binary variables), the above approach becomes infeasible on today’s computers. In this case some other approaches must be used to maintain a consistent set of current market prices, to maintain a description of the assets each person holds, and to determine when a user’s assets can allow them to back up a given new bet.

Limiting the Probability Distribution

One approach to dealing with enormous state spaces is to choose a particular family of probability distributions, and limit users to choosing distributions within that family. Bets must change the probability distribution from one member of this family to another, and user assets changes must be consistent with that constraint. For example, the family of Bayes linear distributions has some attractive computational advantages (Goldstein, 1987). We will here elaborate on Bayes nets

(Pearl, 1988; Jensen, 2001), however, because of their popularity.

In a Bayes net, variables are organized by a directed graph in which each variable α has a set of “parent” variables \mathcal{P}_α . The probability of any state i can then be written as

$$p_i = \prod_{\alpha} p(v_{\alpha}(i) | \{v_k(i)\}_{k \in \mathcal{P}_{\alpha}}),$$

where $v_{\alpha}(i)$ is the value of variable α in state i . If each variable α has V_{α} possible values, this distribution can then be stored in tables, with one table of size $V_{\alpha} \prod_{k \in \mathcal{P}_{\alpha}} V_k$ for each variable α . Thus when the network is *sparse*, meaning that each variable has only a few parents, there are only a few parameters to specify and store per variable, it is straightforward to compute the probability of any given state, and it is easy to calculate or change the probability of a variable value, conditional on the values of the parents of that variable. A market scoring rule which represented its probabilities in terms of a Bayes net could thus easily support bets which change these particular conditional probabilities.

But what about supporting bets on other conditional probabilities, or on the unconditional probabilities of variable values? In order to support simple bets on variable values, one wants to be able to calculate the unconditional probability distribution over the values of a variable given the current network, and one wants to know how to change that network to be consistent with a new desired distribution. It turns out that all this and more is possible if the Bayes net happens to be singly-connected, so that there is at most one path connecting any two variables. In this case, one can change the unconditional distribution of any variable, and then propagate those changes across the network, to exactly update the unconditional distributions of other variables. Conditional distributions of one variable given another can be exactly calculated, such as by provisionally changing one variable and propagating that change through to the other variable. Also, since one can merge a set S of variables, each with V_{α} values, into a single variable with $\prod_{\alpha \in S} V_{\alpha}$ values, one can also apply this approach to Bayes nets that are “nearly” but not exactly singly-connected.⁹

Thus a market scoring rule which represented its probabilities in terms of a nearly singly-connected Bayes net could support bets which change any unconditional variable probabilities, and any conditional probabilities between pairs of variables. Such

a market scoring rule could even support bets which change the network structure, as long as those changes keep the network nearly singly-connected. Unfortunately, while the Bayes nets that seem plausible to people are usually sparse, with only a few connections per variable, they are *not* usually nearly singly-connected. So if a Bayes-net-based market scoring rule allowed users to create any sparse Bayes net they wished, it would typically be infeasible to exactly calculate most unconditional variable probabilities, and most conditional probabilities relating variable pairs.

There are of course many ways to compute approximations to these probabilities. So one might be tempted to endow a market scoring rule with such an approximation algorithm, and then let it accept fair bets at the approximate probabilities determined by that algorithm. This approach, however, might turn the market scoring rule patron into a money pump. This is because the accuracy of these approximations varies in complex ways with the context. If any user could, for any part of the probability distribution, find any systematic pattern in when the official approximation made over-estimates vs. under-estimates, he could in principle use that pattern to make money via arbitrage. To arbitrage, he would make simple bets one way at the approximate probabilities, and make some combination of other bets at the more basic probabilities. Of course it may be hard to find such patterns and bet combinations that exploit them, and the amounts gained may be small compared to the effort required. Also, a market scoring rule with a smart manager might limit losses by detecting such activity. Potential market scoring rule patrons might not find these considerations sufficiently reassuring, however.

In order to guarantee that a market scoring rule patron loses no more than a given amount of money while supporting an indefinite number of trades, one might adopt a policy of only allowing bets on probabilities that one can exactly compute. This policy, however, could prevent most of the trades that users would be most interested in.

Overlapping Market Makers

Another approach to dealing with enormous potential state spaces is to have *several* market scoring rules, each of which dealt exactly with some part of the same

total state space. For example, a market scoring rule that represented its probabilities via a general sparse Bayes net might be combined with market scoring rules that each dealt only in the unconditional probabilities of a single variable. Then if the unconditional probabilities of a variable were not consistent with the Bayes net probabilities, users who discovered this inconsistency could profit thereby, but their profit would be bounded. After all, if each market scoring rule is exactly computed, we can exactly bound the loss of each rule, and hence bound the total loss.

Perhaps the simplest approach of this sort would consist of a set of market scoring rules m that each used a completely general probability distribution over some subset S_m of the variables. Users could then directly make any bets and change any probability estimates, as long as the set of variables S_b used in that bet or probability estimate b was a subset of the set S_m of some market maker m . Such bets or changes would just be implemented by dealing simultaneously and directly with all of the market makers that overlap that set, i.e., by trading with $M(b) = \{m : S_b \subset S_m\}$. Between user-initiated trades of this sort, the system could bring the market makers into greater consistency with each other by seeking out arbitrage opportunities, i.e., sets of trades which produce a pure cash net profit.

Let us call two market makers m and n neighbors if they have variables in common, i.e., if $S_m \cap S_n$ is non-empty. Arbitrage opportunities would exist between any two neighboring market makers if they had different probability distributions over the set of variables they shared. (For logarithmic market makers, arbitraging this difference would change the probability distributions over these shared variables to be a normalized geometric mean of the distributions of the market makers.) Thus all the market makers could be made consistent with each other via waves of arbitrage passing through a network of market makers. Neighboring market makers would be arbitrated, and if this produced a large enough change, neighbors of those neighbors would also be arbitrated, and so on. Once this process stopped, users would in essence be invited to profit by finding and correcting any remaining inconsistencies. This profit might be large if the remaining inconsistencies were large, but it would be strictly bounded by the total subsidy offered by all the market makers.

Under this sort of approach, if one could anticipate the sorts of variables that users would be likely to want

to relate to one another, one might construct an appropriate set of market makers to support those users. Each rule would be based on a type of probability distribution whose consistency across the entire state space could be guaranteed, and each rule would only accept bets on probabilities that it could exactly compute from this distribution. There could be rules that covered single variables, rules for all value combinations for certain sets of a dozen variables, rules for nearly-singly-connected Bayes nets, rules for general sparse Bayes nets, rules for sparse linear covariance matrices, and rules for logistic regressions. Market scoring rule patrons might want the first crack at arbitraging any inconsistencies between probabilities from different rules, if those patrons are willing to take on the risks required. But after this, users would be invited to profit by fixing any remaining inconsistencies.

Avoiding Bankruptcy

Even when probability distributions are exactly consistent, users might still turn a market scoring rule patron into a money pump by playing “heads I win, tails I’m bankrupt.” That is, users might make bets for which they are paid when they win, but which they cannot pay when they lose. To avoid this, a market scoring rule should compare the price changes that a user proposes to the assets that he offers as collateral to cover his proposed bet. Specifically, a logarithmic rule based on $s_i = \beta \log_2(p_i)$ could treat any set of collateral assets as in principle describable in terms of amounts c_i held of each state asset “Pays \$1 if state i .” If a market scoring rule can determine that in all states

$$1 < 2^{c_i/\beta} \left(\frac{p'_i}{p_i} \right),$$

then it can safely approve a change from \mathbf{p} to \mathbf{p}' .

Of course the best way to make such a determination will likely depend on the type of distribution \mathbf{p} handled by this market scoring rule, and on the type of assets offered as collateral. This determination can be very easy when the collateral offered is “cash,” i.e., assets with state-independent value. On the other hand, this determination may get very hard when the collateral is the result of previous bets made with respect to rather different distributions or types of distributions. For example, if the structure of a Bayes net changed substantially over time, then a user who once made a

bet with respect to an old version of the net may find it hard to use that bet as collateral to support a new bet.

While it might be tempting to require cash as collateral for all bets, such a policy would greatly discourage users from making small changes, i.e., changes to the probability distribution which covered only small parts of the state space. This is because small changes would have a similar opportunity cost to large changes, in terms of assets which can no longer be used to support bets, even though the potential payoff from such bets is much smaller. Since it is often not hard to verify that other kinds of assets can serve as collateral, insisting on cash seems much too severe a constraint. Market scoring rules thus need a non-trivial policy regarding how much computation they are willing to undergo to determine that any given collateral will cover any given bet. This policy can of course depend on the kind of changes proposed and collateral offered. Since it is often easier to check proofs than to create them, one possible policy is that in more complex cases the user must provide a “proof” of coverage, a proof that the market scoring rule need only check.

The fact that changes to probability distributions can make it harder for old bets to be used as collateral to cover new bets means that users who make bets can cause a substantial externality on previous users. For example, in a singly-connected Bayes net market scoring rule, changing the unconditional distribution of a variable will have little impact on previous users, but changing the structure of the network can have a large impact. A user who made a bet about the conditional probabilities associated with a given link might easily later undo that bet if that link still remained in the current network, but if the link was gone they might need to first add it back in. And if there were binding complexity limits on the network, this may require the elimination of some other links in the network. Thus previous users can be harmed by taking away links that they have made bets on, and can be benefited by replacing such links. Similar effects come if users can add or delete cutpoints from a digital option or sawtooth representation of a real-valued random variable.

The law and economics perspective seems an appropriate framework to consider this externality (Cooter and Ulen, 2000). If it were easy for each new user to negotiate with effected previous users, then we could just clearly assign a property right for them to trade. This seems difficult, however, both because there will often be many previous users, and because of the possibility of strategically threatening to become a new user.

When negotiations are difficult, we must guess at what rights the parties would assign, if they could negotiate. If the parties could negotiate, would they agree to tax or subsidize certain probability changes, because of the net externality due to those changes?

For topics where we expect too little information aggregation, such as due to free-riding problems among potential patrons, we expect users to produce a net positive information aggregation externality, on average. In contrast, for topics where we expect too much information aggregation, such due to wasteful signaling, we expect a net negative information aggregation externality, on average. With some way to distinguish particular uses from others, in terms of the relative magnitudes of benefits and costs, we should want to subsidize users in the first case and tax them in the second case. While structural changes do seem to have more potential for negative impact on previous users, when there is a net positive information aggregation externality such changes also seem to have more potential for a positive impact via producing a better probability model.

In the absence of a more detailed analysis which can better pinpoint which sorts of changes should be discouraged, relative to other changes, the simplest policy seems best. Let users make whatever structural changes they like, if they can afford the bets and the resulting structural complexity is acceptable, and leave it up to previous users to mitigate the externalities such changes produce on them. Users can mitigate such externalities by not waiting too long before undoing bets, avoiding betting on structural changes they do not expect to last, and seeking synergies among the bets which make to make it easier to prove that earlier bets can serve as collateral for later bets.

In general in the world, the introduction of new products benefits those who use that new product, and others associated with those beneficiaries, but this act can also produce negative externalities on those who are tied to the old products. While this is a shame, on net it still seems better to usually encourage new products. Until we can tell which new products are net harms, we should just embrace them all.

Conclusion

Information markets seem to do well at aggregating information in the thick market case, where many traders

all estimate the same value. But to have combinatorial information markets, which estimate a probability distribution over all combinations of values of many variables, we want to better deal with the thin market and irrational participation problems. Simple scoring rules avoid these problems, but suffer from opinion pool problems in the thick market case. Market scoring rules avoid all these problems.

After reviewing information markets and scoring rules, and presenting market scoring rules, this paper has considered several implementation issues with using market scoring rules as combinatorial information market makers. We saw that sawtooth assets can support both unconditional estimates of expected values, and conditioning on a variable being “near” some value. We saw that by using several market scoring rules on the same state space, each internally consistent but with inconsistencies possible between the rules, one can avoid becoming a money pump due to the difficulty of computing consistent probabilities, while still allowing bets on most probabilities of interest. We saw that market scoring rule managers should want to make some efforts to allow past bets to be used as collateral for future bets, and that it remains an open question just how much effort is appropriate. Finally, we saw that while bets that produce structural changes can have a negative externalities on those who have made bets using older structures, on net it seems better to usually leave the mitigation of this externality to those with old bets.

A summary of the proposed design is as follows:

1. There are several logarithmic market scoring rules, each of which maintains a distribution \mathbf{p} of a certain type, such as a nearly singly-connected Bayes net.
2. Patrons choose a set of variables to patronize. Variables can be added and refined on the fly, if further funding is available.
3. Each real-valued variable z is represented by a discrete variable α , whose values v are assigned payoffs according to sawtooth functions $a_{\alpha v}(z)$ described earlier.
4. Patrons choose an initial distribution \mathbf{q} for each market scoring rule.
5. Patrons choose a subsidy level β for each market scoring rule. The expected loss of patrons has a guaranteed bound, and with uniform initial distributions \mathbf{q} , the exact loss has a guaranteed bound.
6. Any user can at any time make a change to any of these distribution, by changing distribution \mathbf{p}

into another distribution \mathbf{p}' of the same type. Such changes are calculated exactly, without roundoff error.

7. In between user changes to individual market scoring rules, arbitrage opportunities between such rules are found and exploited, reducing the inconsistencies between their probability distributions.
8. Consider a user who has so far on net deposited "cash" d , and has so far made a set of changes $\mathbf{b}(\mathbf{p}', \mathbf{p})$ to various scoring rules, where we define $b_i = p'_i/p_i$. For a logarithmic market scoring rule, such a user can make one more change, and withdraw cash w , if he can prove that, considering all past changes \mathbf{b}_n and this proposed new change, for each state where now $p_i > 0$ we have

$$1 < 2^{(d-w)/\beta} \prod_n b_{in}.$$

9. Whenever a variable's values are assigned, state probabilities are reweighted according to this assignment, and the state space is coarsened to eliminate this variable. Asset holdings become a weighted average of asset holdings in the old states.

The main category of issues not dealt with in this paper is that of user interfaces. Users need to be able to browse a probability distribution over the set of all variable value combinations, to identify the estimates, such as probabilities and expected values, which they think are in error. User then need to be able to choose values for those estimates, and see whether they have enough collateral to cover the bets required to make these changes. Users also need to be able to see how much risk they have already acquired along these dimensions, so they can judge how much more risk is appropriate. While reasonable approaches are known in many cases, it is far from clear what approaches are most effective and robust across a wide range of the most interesting cases.

Acknowledgments

For their comments, I thank participants of the 2002 conference on Information Dynamics in the Networked Society, participants of the 2002 Informal Seminar on Progress in Idea Futures Technology, Hassan Masum, Ravi Pandya, and especially David Pennock.

Notes

1. For a scoring rule that is not proper, this argmax contains other distributions besides \mathbf{p} . In this paper, we consider only proper scoring rules.
2. For rewards that are in the form of "bragging rights," the patron is in essence whatever audience is impressed by such brags.
3. Procedures "commute" when the result is independent of the order of applying them.
4. Note also that if much time passes before states will be verified and payments made, trades are in essence of future assets. Thus if interest rates are positive, fewer real resources need be deposited now to guarantee such an ability to pay. For example, if the future assets traded were shares in a stock index fund, one could deposit or escrow the same number of shares now.
5. This analysis ignores the strategic consideration that one's report may influence the reports that others make later, which could then influence one's profits from making yet more reports after that. But if this process ever reaches an equilibrium where people don't expect further changes, the equilibrium market scoring rule distribution should be what each person would report to a simple scoring rule.
6. It has long been known that in one dimension, an agent using a scoring rule is equivalent to his choosing a quantity from a continuous offer demand schedule (Savage, 1971). This equivalence also holds in higher dimensions.
7. This is the same fact that in physics ensures that the total entropy of two correlated systems is lower than the sum of the entropies of each system considered separately.
8. Users could also withdraw non- β multiple amounts, if they were willing to suffer a rounding down to cover the worst possible roundoff error in calculating the logarithm.
9. Technically, the directed network of a Bayes net can be translated into an undirected network by adding "moral" links between the parents of each variable. The "cliques" of this undirected network are then the largest sets of nodes that are all linked to each other. By adding more links to this network, these cliques can be organized into a "join tree," where for every pair of cliques, their intersection is a subset of every clique on the tree path between these two cliques. The cost to consistently update this network is roughly proportional to the number of variables, times a factor exponential in the size of the cliques. (The computational complexity of finding the cliques of a networks is NP-complete in the worst case, however (Cooper, 1990).)

References

- Aumann R. Agreeing to disagree. *The Annals of Statistics* 1976; 4(6):1236–1239.
- Berg J, Nelson F, Rietz T. Accuracy and forecast standard error of prediction markets. Tech. rep., University of Iowa, College of Business Administration, 2001.
- Brier GW. Verification of forecasts expressed in terms of probability. *Monthly Weather Review* 1950;78:1–3.

- Chen K-Y, Plott CR. Information aggregation mechanism: Concept, design, and implementation for a sales forecasting problem. Tech. rep., California Institute of Technology, 1998.
- Clemen RT. Incentive contracts and strictly proper scoring rules. *Test* 2002;11:195–217.
- Cooper GF. The computational complexity of probabilistic inference using Bayes belief networks. *Artificial Intelligence* 1990;42:393–405.
- Cooter R, Ulen T. *Law and Economics* (3rd edn.), New York: Addison-Wesley, 2000.
- DeWispelare AR, Herren LT, Clemen RT. The use of probability elicitation in the high-level nuclear waste regulation program. *International Journal of Forecasting* 1995;11(1):5–24.
- Druzdel MJ, van der Gaag LC. Elicitation of probabilities for belief networks: Combining qualitative and quantitative information. In: *Uncertainty in Artificial Intelligence*, 1995:141–148.
- Figlewski S. Subjective information and market efficiency in a betting market. *Journal of Political Economy* 1979;87(1):75–88.
- Geanakoplos J, Polemarchakis H. We can't disagree forever. *Journal of Economic Theory* 1982;28:192–200.
- Genest C, Zidek J. Combining probability distributions: A critique and annotated bibliography. *Statistical Science* 1986;1(1):114–148.
- Goldstein M. Systematic analysis of limited belief specifications. *Statistician* 1987;36(2/3):191–199.
- Good IJ. Rational decisions. *Journal of the Royal Statistical Society. Series B (Methodological)* 1952;14(1):107–114.
- Hanson R. Consensus by identifying extremists. *Theory and Decision* 1998;44(3):293–301.
- Hanson R. Decision markets. *IEEE Intelligent Systems* 1999;14(3):16–19.
- Hanson R. Eliciting objective probabilities via lottery insurance games. Tech. rep., George Mason University Economics, 2002a. (<http://hanson.gmu.edu/elicite.pdf>).
- Hanson R. Logarithmic market scoring rules for modular combinatorial information aggregation. Tech. rep., George Mason University Economics, 2002b. (<http://hanson.gmu.edu/mktscore.pdf>).
- Hanson R. Disagreement is unpredictable. *Economics Letters* 2003. (<http://hanson.gmu.edu/unpredict.pdf>).
- Jensen FV. *Bayesian Networks and Decision Graphs*. New York: Springer, 2001.
- Kadane J, Winkler R. Separating probability elicitation from utilities. *Journal of the American Statistical Association* 1988;83(402):357–363.
- Krahnen JP, Weber M. Does information aggregation depend on market structure? Market makers vs. double auction. *Zeitschrift für Wirtschaftsund Sozialwissenschaften* 1999;119:1–22.
- Lo AW. *Market Efficiency: Stock Market Behaviour in Theory and Practice*. Elgar, Lyme, 1997.
- Milgrom P, Stokey N. Information, trade and common knowledge. *Journal of Economic Theory* 1982;26:17–27.
- Murphy AH, Winkler RL. Probability forecasting in meteorology. *Journal of the American Statistical Association* 1984;79(387):489–500.
- O'Carroll FM. Subjective probabilities and short-term economic forecasts: An empirical investigation. *Applied Statistics* 1977;26(3):269–278.
- Pearl J. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, California: Morgan Kaufmann, 1988.
- Pennock DM, Giles C, Nielsen F. The real power of artificial markets. *Science* 2001;291:987–988.
- Pennock DM, Wellman MP. Compact securities markets for Pareto optimal reallocation of risk. In: *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, San Francisco: Morgan Kaufmann, 2000:481–488 (<http://rome.exp.sis.pitt.edu/UAI/Abstract.asp?articleID=56>).
- Roll R. Orange juice and weather. *The American Economic Review* 1984;74(5):861–880.
- Savage LJ. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association* 1971;66(336):783–801.
- Smith CAB. Personal probability and statistical analysis. *Journal of the Royal Statistical Society. Series A (General)* 1965;128(4):469–499.
- Winkler RL. Scoring rules and the evaluation of probability assessors. *Journal of the American Statistical Association* 1969;64(327):1073–1078.

Robin Hanson is an assistant professor of economics at George Mason University. In 1998 Robin received his Ph.D. in social science from the California Institute of Technology, and was afterward a Robert Wood Johnson Foundation health policy scholar at the University of California at Berkeley. Earlier, in 1984, Robin received a masters in physics and a masters in the philosophy of science from the University of Chicago, and then spent nine years researching artificial intelligence, Bayesian statistics, and hypertext publishing at Lockheed, NASA, and independently.

Robin has published in *CATO Journal*, *Communications of the ACM*, *Economics Letters*, *Econometrica*, *Entropy*, *IEEE Intelligent Systems*, *Information Systems Frontiers*, *International Joint Conference on Artificial Intelligence*, *Journal of Evolution and Technology*, *Journal of Public Economics*, *Social Epistemology*, *Social Philosophy and Policy*, and *Theory and Decision*.

Robin has unusually diverse research interests, with papers on spatial product competition, health incentive contracts, group insurance, explaining product bans, evolutionary psychology of health care, bioethics, voter incentives to become informed, Bayesian statistics and classification, agreeing to disagree, self-deception in disagreement, information aggregation via speculative markets, governance based on such markets, probability elicitation, combinatorial information market makers, wiretaps, image reconstruction, the history of science prizes, reversible computation, the origin of life, very long term economic growth, growth given machine intelligence, and interstellar colonization.