Eliciting Objective Probabilities via Lottery Insurance Games

Robin Hanson*
Department of Economics
George Mason University†

Abstract

Since utilities and probabilities jointly determine choices, event-dependent utilities complicate the elicitation of subjective event probabilities. However, for the usual purpose of obtaining the information embodied in agent beliefs, it is sufficient to elicit *objective* probabilities, i.e., probabilities obtained by updating a known common prior with that agent's further information. Bayesians who play a Nash equilibrium of a certain insurance game before they obtain relevant information will afterward act regarding lottery ticket payments as if they had event-independent risk-neutral utility and a known common prior. Proper scoring rules paid in lottery tickets can then elicit objective probabilities.

Introduction

Proper scoring rules and related methods are widely used to elicit event probabilities. Such probability elicitation is practiced in weather forecasting (Murphy & Winkler, 1984), economic forecasting (O'Carroll, 1977), risk analysis (DeWispelare, Herren, & Clemen, 1995), and the engineering of intelligent computer systems (Druzdzel & van der Gaag, 1995). The fact that choices are determined jointly by both utilities and subjective probabilities, however, makes such elicitation problematic. For example, if one knows only that an agent is a Bayesian, but cannot further constraint his utilities or probabilities, then it seems impossible to infer this agent's subjective probabilities from his actions (Kadane & Winkler, 1988).

When further constraints are available, this problem can be overcome. For example, simple scoring rules elicit subjective probabilities when utility is state-independent and risk-neutral (Savage, 1971). When utility and endowments are state-independent but utility is not risk-neutral, it is sufficient to use scoring rules that pay in lottery tickets (Smith, 1965;

*For their comments, I thank Chris Hibbert and Edi Karni.
†*rhanson@gmu.edu* http://hanson.gmu.edu 703-993-2326 FAX: 703-993-2323 MSN 1D3, Carow Hall, Fairfax VA 22030

Savage, 1971). When endowments but not utilities are event-dependent, endowments can be inferred by comparing local coefficients of absolute risk aversion, when this varies (Jaffray & Karni, 1999). And when utilities approach known upper and lower bounds, lotteries with extreme consequences can elicit subjective probabilities even for event-dependent utility (Karni, 1999; Jaffray & Karni, 1999). Unfortunately, the assumptions required for these solutions do not seem to apply to many contexts where probabilities are elicited.

In most applications of probability elicitation, however, including those listed above, the purpose is not to obtain someone's subjective probabilities, but rather to elicit the *information* embodied in his probabilities. We query a weather forecaster, for example, to find out what he knows about whether it will rain. We query a loan evaluator to find out what he knows about whether a loan will default. And we query an expert witness at a trial to find out what he knows about whether an accused person is guilty.

In such cases, we should want to minimize the influence of unknown variations in the prior beliefs and utilities of such an agent on the probabilities we elicit from him. We should thus want to elicit *objective* probabilities, i.e., the probabilities that an expert would hold had he started with a certain known prior, and then updated his beliefs based on his expert information. If we elicit sufficiently detailed objective probabilities, then we can directly infer the information this person has acquired since the known common prior.

For example, weather forecasters largely agree on their predictions about the weather a year in advance. This far in advance they basically revert to base rate frequencies; such as that the chance of rain on any given spring day in a certain city is about 25%. Then as the date in question approaches, a weather forecaster's estimate of the chance of rain typically changes as he acquires specific information and analyzes. Similarly, loan evaluators revert to base rates in evaluating the chance that an unknown 42 year old white female who makes $40,000 a year will repay her home loan, and polygraph experts also revert to base rates to evaluate the chance that any random statement in a polygraph test is a lie. These experts then typically change their estimates as they learn more about individual cases.

When we query weather forecasters, loan evaluators, or polygraph experts, what we mostly want is the information that induced them to change their estimates, rather than their disagreements about base rates. If one polygraph expert believes that women tend to lie more often, relative to what other polygraph experts believe, we usually prefer that his expert court testimony in a particular case not be influenced by this belief. We instead prefer to learn about how much his beliefs changed, relative to his base rate estimates, due to what he learned about a particular case. This change should be a reliable indication of his new information, and this change is what we learn when we elicit objective probabilities.

One approach to eliciting objective probabilities combines scoring rules and insurance markets. Insurance markets have long been known to induce equality of marginal rates of substitution across events (Debreu, 1959; Kadane & Winkler, 1988). Such an equality implies that agents will immediately afterward act as if they had a common prior and event-independent utility, at least regarding very small changes in event-dependent assets. So if an agent's event utilities do not change as he acquires new information, he should then respond to scoring rules with infinitesimal payments as if he had updated the common prior with his new information. Thus objective probabilities can be elicited by having agents first participate in an insurance market, and then later respond to an infinitesimal scoring rule.

Infinitesimal scoring rules have serious problems, however. Any small reason why reporting one probability distribution is easier than reporting another distribution can overwhelm the scoring rule incentives. For example, if there is a default distribution, and agents must make a fixed finite effort to change this default, then with a weak enough scoring rule agents should just leave the default unchanged, regardless of what they believe.

Larger scoring rule payments are possible if state-dependent marginal utilities are equalized regarding lottery tickets, since changes in lottery ticket holdings do not change one's marginal utilities across states. However, an insurance market equilibrium will not typically create such equalization, because the optimal number of lottery tickets is typically extreme; agents typically either want to be sure they win the lottery, or to be sure they do not win. To induce intermediate choices, we can follow the example of mixed strategy equilibria of simple two player games. In such equilibria, each player is indifferent between his possible mixtures, and yet still chooses the equilibrium mixture that makes the other player indifferent. Similarly, we can create an insurance game where each agent is paired with another agent, and where each agent chooses an allocation of lottery tickets that makes the other agent indifferent to his lottery allocation, via equalized marginal utilities.

This paper introduces such a procedure for eliciting objective probabilities, a procedure that combines scoring rules, lotteries, early insurance, and Nash equilibria of a certain game. This procedure goes as follows. First, a set of relevant events are chosen, and then well before some risk-averse agents acquire case specific information, they play a certain game. This game allocates two kinds of event-contingent assets: event-contingent cash, such as "Pays $1 if rain here Tuesday," and event-contingent lottery tickets, such as "Pays one lottery ticket if rain here Tuesday." Given a budget constraint and some equilibrium prices for event-contingent lottery tickets, each expert chooses his ideal mixture of event-contingent lottery tickets. Simultaneously, each agent's event-contingent cash is determined by the lottery ticket choices of some other random agent.

Similar to the way that some games only have mixed strategy Nash equilibria, in all Nash equilibria of this game, every agent chooses an intermediate amount of every kind of state-contingent lottery ticket. And this implies that, immediately after playing this game, all agents respond to payoffs denominated in state-dependent lottery tickets as if they had risk-neutral event-independent utility, and had the same known beliefs about event probabilities. Furthermore, if certain utility ratios do not change as agents then acquire new information about these events, agents will still respond to such payoffs as if they had risk-neutral event-independent utility. If so, then agents will later report their objective probabilities in response to a proper scoring rule paid in event-dependent lottery tickets.

For example, weather forecasters might anticipate regularly making forecasts about whether it will rain in a certain city, each day making a forecast for each of the next five days. A year in advance, these forecasters might participate in a special insurance market, where they acquire cash and lottery tickets that pay depending on whether it rains on particular combinations of days. After this insurance market reached equilibrium and closed, each forecaster should be indifferent to buying more or less of each kind of event-dependent lottery ticket. For example, a forecaster who beforehand prefered more to win the lottery if it rained on his birthday than if it did not should now no longer care. If during the next year this forecaster does not acquire new information which changes how much he cares, such as learning of special birthday visitors, then he should still not care in the few days before his birthday, when he makes forecasts in response to proper scoring rules that pay in lottery tickets. If so, then his forecasts should be objective probabilities, describing the information he had acquired in the previous year about rain on the day in question.

We will now review the logic of scoring rules, lottery ticket payments, early insurance, and an lottery insurance game which induces honest reports of objective probabilities. Finally, we consider some implementation issues.

Analysis

Consider a Bayesian principal concerned about her beliefs, represented by probabilities $\tilde{\pi}_i$, regarding a complete set I of disjoint events i (so that $\sum_{i\in I}\tilde{\pi}_i = 1$). In order to better inform her beliefs, this principal might consult one or more agents as experts, and ask them about their beliefs. For example, imagine that a Bayesian expected-utility-maximizing agent had previously honestly reported his beliefs to be $\pi_{i\,0}$, but now, after having acquired important relevant information, honestly reports his beliefs to be $\pi_{i\,t}$. In this case, the principal can infer that the likelihood ratio of this agent's new information was proportional to $\pi_{i\,t}/\pi_{i\,0}$. So if the principal was at first equally informed, and had initial beliefs $\tilde{\pi}_{i\,0}$ (and if all $\pi_{i\,0} > 0$), she could make her new beliefs embody this agent's new information by setting

$$\tilde{\pi}_{i\,t} = \frac{\tilde{\pi}_{i\,0}(\pi_{i\,t}/\pi_{i\,0})}{\sum_{j\in I}\tilde{\pi}_{j\,0}(\pi_{j\,t}/\pi_{j\,0})}.$$

To induce her agents to honestly report their beliefs, a principal might commit to pay each agent a cash amount z according to a *proper scoring rule* (Savage, 1971) $z_i = s_i(\vec{r})$. Here z_i is the cash payment if i turns out to be the actual event, r_i is the probability reported by an agent for the event i, and $\vec{r} = \{r_i\}_{i\in I}$ is the full report. When the s_i constitute a proper scoring rule[1], an agent who sets his reports r_i to maximize his expected monetary payoff will honestly report $r_i = \pi_i$. That is, if Δ is the set of all probability distributions over the events I, then

$$\vec{\pi} = \text{argmax}_{\vec{r}\in\Delta}\sum_{i\in I}\pi_i s_i(\vec{r}).$$

For example, this condition holds for a logarithmic scoring rule $s_i = a_i + b\log(r_i)$ (Good, 1952), for a spherical scoring rule $s_i = a_i + b\,r_i/\sqrt{\sum_{j\in I} r_j^2}$, and for a quadratic scoring rule $s_i = a_i - b\sum_{j\in I}(1_{ij} - r_i)^2$, where $1_{ij} = 1$ when $i = j$ and 0 otherwise (Brier, 1950). Scoring

[1] For a scoring rule that is not proper, other reports as well may maximize an agent's expected payoff.

rules can also give agents incentives to acquire information they would not otherwise possess (Clemen, 2002).

An agent who was not risk neutral, however, would not maximize his expected payoff z, but rather an expected utility $u(z)$, where $u''(z) \neq 0$. In this situation the principal might use a *lottery scoring rule*, i.e. a proper scoring rule that pays in lottery tickets instead of cash (Smith, 1965; Savage, 1971). If e is an agent's initial cash endowment, L is the amount he could win in the lottery, and x denotes his chance of winning the lottery (i.e., his number of lottery tickets times the chance each has of winning), then if the principal paid $x_i = s_i(\vec{r})$ (and if[2] $s_i(\vec{r}) \in [0, 1]$ and he has no other way to win this lottery), then an agent who set his report r_i to maximize expected utility

$$\sum_{i \in I} \pi_i [(1 - x_i) u(e) + x_i u(e + L)] = u(e) + [u(L + e) - u(e)] \sum_{i \in I} \pi_i x_i$$

should again honestly report $r_i = \pi_i$.

A lottery scoring rule is not enough, however, if an agent has event-dependent cash endowments e_i or utilities $u_i(e)$. Such an agent will make reports to maximize

$$\sum_{i \in I} \pi_i [(1 - x_i) u_i(e_i) + x_i u_i(e_i + L)] = \sum_{i \in I} \pi_i u_i(e_i) + \sum_{i \in I} \pi_i \Delta u_i(e_i) x_i, \tag{1}$$

where $\Delta u_i(e) = u_i(e + L) - u_i(e)$. That is, such an agent will report as if he had beliefs proportional to $\pi_i \Delta u_i(e_i)$. A principal who was ignorant about an agent's utility steps Δu_i, or about his prior, would have difficulty inferring that agent's information from his reports (Kadane & Winkler, 1988).

Early insurance can mitigate problems caused by ignorance of agent priors and utilities. Assume that at some previous point in time, a set Λ of agents participate in a competitive insurance market, with prices that agents take as fixed. That is, at prices p_i (where $\sum_{i \in I} p_i = 1$) each agent adjusts his endowment e_i with hedges[3] h_i, so that his cash given event i becomes $z_i = e_i + h_i$. Such an agent would then maximize

$$\sum_{i \in I} \pi_i u_i(e_i + h_i) \quad \text{given} \quad \sum_i p_i h_i = 0.$$

Assuming an interior choice, each risk-averse ($u_i'' < 0$, $u_i' > 0$) agent would then hold z_i satisfying

$$\pi_i \, u_i'(z_i) = \lambda \, p_i,$$

where $\lambda = \sum_{i \in I} \pi_i u_i'(z_i)$. For very small asset changes Δz_i, such an agent would then maximize

[2]This cannot hold for all $\vec{r}$ under the logarithmic scoring rule, but can under the quadratic and spherical scoring rules.

[3]Presumably in the aggregate, $\sum_{\alpha \in \Lambda} h_{\alpha \, i} = 0$.

$$\sum_{i\in I} \pi_i u_i(z_i+\Delta z_i) \approx \sum_{i\in I} \pi_i u_i(z_i) + \sum_{i\in I} \pi_i u_i'(z_i)\Delta z_i = \sum_{i\in I} \pi_i u_i(z_i) + \lambda \sum_{i\in I} p_i \Delta z_i.$$

Thus in response to an infinitesimal scoring rule payment $\Delta z_i = s_i(\vec{r})$ presented soon after early insurance, each agent should report $r_i \approx p_i$, with the approximation exact in the limit of zero asset changes Δz_i. Regarding such infinitesimal rules, agents thus act immediately after early insurance as if they had a common prior p_i. And an agent who later acquired new information about the events i, but had no change in his assets z_i, should report as if he had updated this common prior p_i based on his new information. (More on this below.)

Unfortunately, infinitesimal payments have serious practical problems. Not only do they eliminate the incentive for agents to acquire information, but they completely fail if there are any finite differential reporting costs. For example, if there is a default report $\vec{r}$, so that agents must exert finite effort to change this default report, then if the scoring rule payments are small enough, agents will always make the default report, regardless of what they believe. Fortunately, the insurance approach can be combined with a lottery approach to allow larger payments.

Consider a lottery scoring rule which pays an agent α in (state-contingent) lottery tickets $x_{\alpha i}$. It turns out to be sufficient to adjust this agent's assets $z_{\alpha i}$ until they satisfy

$$\pi_{\alpha i}\, \Delta u_{\alpha i}(z_{\alpha i}) = \lambda_\alpha\, p_i, \tag{2}$$

where $\lambda_\alpha = \sum_{i\in I} \pi_{\alpha i} \Delta u_{\alpha i}(z_{\alpha i})$.

Lemma 1 *Agents who satisfy equation 2 respond to a lottery scoring rule with $\vec{r} = \vec{p}$.*

Proof. According to equation 1, equation 2 induces agent α to make reports to maximize

$$\sum_{i\in I} \pi_{\alpha i}\, \Delta u_{\alpha i}(z_{\alpha i}) x_{\alpha i} = \lambda_\alpha \sum_{i\in I} p_i x_{\alpha i},$$

and hence to respond to a scoring rule $x_{\alpha i} = s_i(\vec{r}_\alpha)$ with $r_{\alpha i} = p_i$. QED.

Lemma 2 *If Bayesian α satisfied equation 2 at $t = 0$, and then at time t had new beliefs $\pi_{\alpha i t}$ but the same utility steps $\Delta u_{\alpha i}(z_{\alpha i})$, then this agent should then respond to a lottery scoring rule as if he had updated the prior p_i with his new information.*

Proof. According to equation 2, he at first reports

$$p_{\alpha i 0} = p_i = \frac{\pi_{\alpha i 0} \Delta u_{\alpha i}(z_{\alpha i})}{\sum_{j\in I} \pi_{\alpha j 0} \Delta u_{\alpha j}(z_{\alpha j})}$$

and by assumption he later substitutes $\pi_{\alpha i t}$ for $\pi_{\alpha i 0}$ and so reports according to

$$p_{\alpha i t} = \frac{\pi_{\alpha i t} \Delta u_{\alpha i}(z_{\alpha i})}{\sum_{j\in I} \pi_{\alpha j t} \Delta u_{\alpha j}(z_{\alpha j})} = \frac{p_i (\pi_{\alpha i t}/\pi_{\alpha i 0})}{\sum_{j\in I} p_j (\pi_{\alpha j t}/\pi_{\alpha j 0})}.$$

Since $\pi_{\alpha i t}/\pi_{\alpha i 0}$ is proportional to the likelihood ratio for agent α's new information, this is just as if agent α had updated the prior p_i based on his new information. QED.

Adjustments satisfying equation 2 can arise if the agents play a certain early insurance game. Let agents be paired as partners β, γ and let each agent's assets depend on the lottery tickets of his partner according to

$$z_{\beta i} = f(x_{\gamma i}, \underline{z}_{\beta i}, \bar{z}_{\beta i})$$

$$z_{\gamma i} = f(1 - x_{\beta i}, \underline{z}_{\gamma i}, \bar{z}_{\gamma i},),$$

where $\underline{z}_{\alpha i} < \bar{z}_{\alpha i}$,

$$f(x, \underline{z}, \bar{z}) = \underline{z} + (\bar{z} - \underline{z})g(x_{\beta i}),$$

and $g(x)$ is continuous, non-decreasing on $[0, 1]$, strictly increasing at $x = 1/2$, $g(x) = 0$ for $x \leq D \in (0, 1/2)$, and $g(x) = 1$ for $x \geq 1 - D$. In the *lottery insurance game*, the above equations apply, each agent faces a budget constraint

$$\sum_{i \in I} p_i \, x_{\alpha i} = 1/2,$$

and each agent chooses his favorite allocation of lottery tickets $x_{\alpha i} \in [0, 1]$.

Each partner β who plays this lottery insurance game should seek to choose the $x_{\beta i} \in [0, 1]$ to maximize

$$\sum_{i \in I} \pi_{\beta i} \Delta u_{\beta i}(f(x_{\gamma i}, \underline{z}_{\beta i}, \bar{z}_{\beta i}))x_{\beta i} \qquad (3)$$

given his budget constraint (and similarly for partners γ). Let μ_α be the Lagrange multiplier of the budget constraint in such an optimization, and let us define marginal value of an agent in a state as

$$m_{\alpha i}(z_{\alpha i}) = \pi_{\alpha i} \, \Delta u_{\alpha i}(z_{\alpha i})/p_i.$$

Risk-aversion implies $m'_{\alpha i} < 0$, and the Kuhn-Tucker conditions of such optimizations say that unless $m_{\alpha i}(z_{\alpha i}) = \mu_\alpha$, we must have either $m_{\alpha i}(z_{\alpha i}) > \mu_\alpha$ and $x_{\alpha i} = 1$, or $m_{\alpha i}(z_{\alpha i}) < \mu_\alpha$ and $x_{\alpha i} = 0$. That is, an agent buys as many tickets as possible for states with above the critical marginal value, and buys no tickets at all for states with below the critical marginal value. Equation 2, the equation we want satisfied, can be written as $m_{\alpha i}(z_{\alpha i}) = \lambda_\alpha$ for all i, i.e., all states have the same marginal value. Let us say that the critical marginal value is *interior* for state i whenever

$$m_{\alpha i}(\underline{z}_{\alpha i}) > \mu_\alpha > m_{\alpha i}(\bar{z}_{\alpha i}),$$

that it is *high* when $\mu_\alpha \geq m_{\alpha i}(\underline{z}_{\alpha i})$, and *low* when $m_{\alpha i}(\bar{z}_{\alpha i}) \geq \mu_\alpha$. It turns out to be sufficient to have interior critical marginal values.[4]

[4]Note that if during the lottery insurance game an agent anticipates later making one or more lottery scoring rule reports, he should realize that his final $x_{\alpha i}$ will not equal the $x_{\alpha i}$ he chooses in the game. This should not change his equilibrium game choices, however, if such choices are interior $x_{\alpha i} \in [0, 1]$, and after the later change in lottery tickets he still satisfies $x_{\alpha i} \in [0, 1]$.

Lemma 3 *For risk-averse partners in a Nash equilibrium of the lottery insurance game, if these agents have interior critical marginal values for state i, then equation 2 is satisfied for i, with $x_{\alpha i} \in [D, 1-D]$.*

Proof. Let $z^*_{\alpha i}$ be defined by $\mu_\alpha = m_{\alpha i}(z^*_{\alpha i})$, which is unique by risk-aversion. The assumption $m_{\alpha i}(\underline{z}_{\alpha i}) > \mu_\alpha > m_{\alpha i}(\bar{z}_{\alpha i})$ can then be written as $\underline{z}_{\alpha i} < z^*_{\alpha i} < \bar{z}_{\alpha i}$. Because of this, neither $x_{\beta i} = 0$ nor $x_{\beta i} = 1$ can be equilibria. If $x_{\beta i} = 0$, then $z_{\gamma i} = \bar{z}_{\gamma i} > z^*_{\gamma i}$, which induces γ to choose $x_{\gamma i} = 0$, which implies $z_{\beta i} = \underline{z}_{\beta i} < z^*_{\beta i}$, which induces β to choose $x_{\beta i} = 1$, which is a contradiction. Similar contradictions follow from $x_{\beta i} = 1$, and from $x_{\gamma i} = 0$ or 1. But if $x_{\alpha i} \in (0,1)$ then by the Kuhn-Tucker conditions $m_{\alpha i}(z_{\alpha i}) = \mu_\alpha$. Thus equation 2 must be satisfied for $\alpha = \beta, \gamma$. Since $x_{\beta i}, x_{\gamma i} \neq 0, 1$, the constraint that they be in $[D, 1-D]$ follows directly from $g(x) = 0$ for $x \leq D > 0$, and $g(x) = 1$ for $x \geq 1-D$. QED.

Note that the logic of this equilibrium is similar to that of simple two-player two-action games, such as the matching-penny game. In such simple games, each player chooses a probability distribution over actions that makes the other player indifferent between his actions, and thus willing himself to choose a non-extreme probability distribution. In this lottery insurance game, choosing lottery ticket holdings $x_{\alpha i}$ is similar to choosing a probability distribution, in that agents will want to choose extremal holdings unless they are indifferent between all values. The functions relating $z_{\beta i}$ to $x_{\gamma i}$ and $z_{\gamma i}$ to $x_{\beta i}$ have been chosen so that each player will induce his partner to become indifferent regarding lottery ticket holdings.

To ensure interior critical marginal values, we need to restrict the values $\underline{z}_{\alpha i}, \bar{z}_{\alpha i}$. Let us say that an insurance game is *sufficiently broad* for paired agents β, γ, it is known that for both partners $\alpha = \beta, \gamma$ and all states i, j, k,

$$m_{\alpha i}(\underline{z}_{\alpha i}) > m_{\alpha j}(f(1/2, \underline{z}_{\alpha j}, \bar{z}_{\alpha j})) > m_{\alpha k}(\bar{z}_{\alpha k}).$$

Lemma 4 *If a lottery insurance game is sufficiently broad, then for risk-averse agents its Nash equilibria have interior critical marginal values for all states i.*

Proof. For each agent α and state i, the critical marginal value can be high, interior, or low. If we define $\underline{M}_\alpha = \min_j m_{\alpha j}(\underline{z}_{\alpha j})$ and $\overline{M}_\alpha = \max_j m_{\alpha j}(\bar{z}_{\alpha j})$, then sufficiently broad games satisfy $\underline{M}_\alpha > \overline{M}_\alpha$, and there are three cases overall to consider for each agent. The critical marginal value is either high overall $\mu_\alpha \geq \underline{M}_\alpha$, interior overall $\underline{M}_\alpha > \mu_\alpha > \overline{M}_\alpha$, or low overall $\overline{M}_\alpha \geq \mu_\alpha$. If an agent's marginal critical value is high overall, then for each state it is high or interior. If it is low overall, for each state it is low or interior, and if it is interior overall, then it is interior for all states. To prove that only this last case occurs, we show that in all other cases some agent violates his budget constraint.

If an agent's marginal critical value is high overall, then by the definition of sufficiently broad, $\mu_\alpha > m_{\alpha i}(f(1/2, \underline{z}_{\alpha i}, \bar{z}_{\alpha i}))$ for every i. If partner β in a pair is high overall, then $x_{\gamma i} = 1$ or $x_{\gamma i} < 1/2$, because $x_{\gamma i} \in (0,1)$ implies $\mu_\alpha = m_{\alpha i}(f(x_{\gamma i}, \underline{z}_{\alpha i}, \bar{z}_{\alpha i}))$, m is decreasing, and g is non-decreasing. Similarly, if β is low overall, then $x_{\gamma i} = 0$ or $x_{\gamma i} > 1/2$, if γ is high overall $x_{\beta i} = 0$ or $x_{\beta i} > 1/2$, and if γ is low overall, $x_{\beta i} = 1$ or $x_{\beta i} < 1/2$.

With two partners β, γ, there are nine possible combinations of the two critical marginal values being high overall, interior overall or low overall. There are two classes of combinations to consider, when one agent is interior overall, and when neither agent is interior overall. Situations within each class are very similar.

As an example of the first class, consider the situation where β is high overall and γ is interior overall. In this situation, for some state i either both β and γ are interior, or β is high while γ is interior. In the former case, $x_{\gamma i} \in (0,1)$ implying $x_{\gamma i} < 1/2$, and in the latter case $\mu_\beta > m_{\beta j}(z_{\beta i})$ implying $x_{\beta i} = 0$ implying $\mu_\gamma > m_{\beta j}(z_{\beta i})$ implying $x_{\gamma i} = 0$. Thus $x_{\gamma i} < 1/2$ for all states i. But then $\sum_{i \in I} p_i x_{\gamma i} < 1/2$, and so γ's budget constraint is violated. A similar construction shows that in other situations where one agent is interior overall, this agent's budget constraint is violated if his partner is not interior overall.

As an example of the class of situations where neither agent is interior overall, consider the situation where β is high overall and γ is high overall. In this situation, for some state i both β and γ are interior, or β is high while γ is interior, or γ is high while β is interior, or both β and γ are high. As before, when both are interior, $x_{\gamma i} < 1/2$, and when β is high $x_{\gamma i} = 0$. In addition, when both are high or when only γ is high, $\mu_\gamma > m_{\gamma j}(z_{\gamma i})$ implying $x_{\gamma i} = 0$. Thus again $x_{\gamma i} < 1/2$ for all states i, violating γ's budget constraint. A similar construction shows that in other situations where neither partner is interior overall, one partner's budget constraint is violated. QED.

Lemma's 1-4 together imply the main result of this paper. Let $x_{\alpha i t}$ be agent α's lottery tickets at time t, after responding to a lottery scoring rule, and let $x_{\alpha i 0}$ be his lottery ticket holdings before responding to this lottery scoring rule. Thus $\Delta x_{\alpha i} = x_{\alpha i t} - x_{\alpha i 0}$ is the transfer of lottery tickets due to the lottery scoring rule.

Theorem 1 *After a Nash equilibrium of a sufficiently-broad lottery insurance game, and assuming risk-averse Bayesians whose utility steps do not change before responding to lottery scoring rules $\Delta x_{\alpha i} = s_i(\vec{r}_\alpha) \in [-D, D]$, such agents respond as if they had event-independent risk-neutral utility, and as if they had updated a known common prior with their further private information. (I.e., they maximize $\sum_{i \in I} p_{\alpha i t} \Delta x_{\alpha i}$ where $p_{\alpha i t}$ updates the insurance price prior p_i with an agent α's new information up to time t.)*

Proof. Assuming risk-averse agents, Lemma 4 shows that if a lottery insurance game is sufficiently broad, its Nash equilibria have interior critical marginal values for all states i, and Lemma 3 shows that this implies equation 2 is satisfied afterward, with $x_{\alpha i} \in [D, 1-D]$. A lottery scoring rule with $\Delta x_{\alpha i} = s_i(\vec{r}_\alpha) \in [-D, D]$ then produces final $x_{\alpha i t} \in [0,1]$. Lemma 1 shows that, given equation 2, agents respond to a lottery scoring rule with the insurance prices p_i, and Lemma 2 shows that if a Bayesian later acquired new information, but had not changes in his event utility steps, he would then respond to a lottery scoring rule as if he had updated the insurance price prior p_i with his new information. QED.

Implementation Issues

Let us now consider some more practical implementation issues with lottery insurance games.

First consider that a Nash equilibrium is not a plausible outcome if agents interact only once, and are ignorant of the utility functions of the agents that they are paired with. In this case we might more plausibly expect a Bayes-Nash equilibrium, where agents average over their beliefs about the various possible utility functions they might encounter. A Nash equilibrium might more plausibly result, however, if agent β was continuously told the current tentative values of his $z_{\beta i}$, computed from the current tentative value of his partner's $x_{\gamma i}$, and could adjust his tentative $x_{\beta i}$ choices in response to changes in $z_{\beta i}$ resulting from changes in $x_{\gamma i}$, and similarly for partner γ. If a situation were reached where no agent wanted to change his $x_{\alpha i}$ values, and if those were the values implemented, that would constitute a Nash equilibrium. (Such repeated interaction is a common technique for seeking a Nash Equilibrium in the face of agent uncertainty (Kalai & Ledyard, 1998; Kalai & Lehrer, 1993).)

A Nash equilibrium would of course not obtain if agents colluded with one another, jointly agreeing to choices of $x_{\alpha i}$ and side payments. Agents do not need to be told which other agents they are paired with, however, or even which half of the pair they are. The more agents there are, the harder it should be for each agent to discover and collude with his counterpart. (A large number of agents would also make it easier to minimize the per-agent cost of covering the excess demand in lottery tickets due to errors in choosing appropriate insurance prices p_i.)

Voluntary participation by all agents should be possible via appropriate choices of assets $\underline{z}_{\alpha i}, \bar{z}_{\alpha i}$. For example, if by not participating agent α would receive $\tilde{z}_{\alpha i}$, then his participation seems ensured if $\underline{z}_{\alpha i} \geq \tilde{z}_{\alpha i}$ for all i.

Our assumption that the utility steps $\Delta u_{\alpha i}(z_{\alpha i})$ do not change clearly holds if agents only get information about the events I. If, however, agents can get information about a finer partition of events Ω, then the question is whether the utility ratios

$$\frac{\Delta u_{\alpha i}}{\Delta u_{\alpha j}} = \frac{\sum_{\omega \in \Omega} \pi_\alpha(\omega|\, i) \Delta u_{\alpha \omega}(z_{\alpha \omega})}{\sum_{\omega \in \Omega} \pi_\alpha(\omega|\, j) \Delta u_{\alpha \omega}(z_{\alpha \omega})},$$

remain constant in the face of such information. For an agent who is insured regarding a sufficient number of events, these ratios should not change. And information that agents may receive that is uncorrelated with the events I is irrelevant. But if agents could plausibly receive information that would change these ratios, then it would be desirable to refine the events I into more detailed events until this constant ratio condition holds for the new events. On the other hand, the insurance market events can be a coarsening of the partition I if it is known that this does not prevent exchanges. That is, kinds of insurance for which there is no demand need not be offered.

Note that having too coarse a set of events I can also make it difficult to infer an agent's information from the likelihood ratios $\pi_{\alpha i t}/\pi_{\alpha i 0}$, which are revealed by a lottery scoring rule after an lottery insurance game. And of course the fact that people are not Bayesians can also make it difficult to infer their information via a lottery insurance game.

There are clearly many practical difficulties which might prevent lottery insurance games from exactly or always eliciting objective probabilities. Probability elicitation is a common practice, however, even using mechanisms which in theory require even stronger assumptions to work than lottery insurance games. So even if lottery insurance games functioned with a substantial error rate, they might still be attractive if they has a substantially smaller error rate than other probability elicitation mechanisms.

Conclusion

The usual reason to elicit probabilities is to obtain information that experts have acquired. This purpose can be served by eliciting each agent's objective probabilities, i.e., the beliefs he would have if he had updated a known common prior with their further information. Simple insurance can achieve this result for infinitesimal scoring rules payments, and this paper has shown how a more complex approach, lottery insurance games, can achieve this result for substantial lottery ticket payments.

That is, the following approach can induce Bayesian expected-utility-maximizing agents to make honest scoring rule reports as if they had updated a common and known prior with their further information. First a principal chooses a set of events of interest, and supplements these events with other events on correlated risks. The principal also chooses a wide enough and high enough range of state-dependent agent cash levels for each agent. Before the agents acquire their differing relevant information on these events, the agents play a lottery insurance game, using common known prices, and come to an equilibrium where no agent wants to change his lottery ticket holdings, given the choices of the other agents. After this, and after acquiring new relevant information, unless an agent acquires information that changes his value of winning the lottery given different events, he should respond to a scoring rule that pays in lottery tickets as if he had updated the insurance price prior with his new information.

References

Brier, G. W. (1950). Verification of Forecasts Expressed in Terms of Probability. *Monthly Weather Review*, *78*, 1–3.

Clemen, R. T. (2002). Incentive Contracts and Strictly Proper Scoring Rules. *Test*, *11*, 195–217.

Debreu, G. (1959). *Theory of Value*. Wiley, New York.

DeWispelare, A. R., Herren, L. T., & Clemen, R. T. (1995). The use of probability elicitation in the high-level nuclear waste regulation program. *International Journal of Forecasting*, *11*(1), 5–24.

Druzdzel, M. J., & van der Gaag, L. C. (1995). Elicitation of Probabilities for Belief Networks: Combining Qualitative and Quantitative Information. In *Uncertainty in Artificial Intelligence*, pp. 141–148.

Good, I. J. (1952). Rational Decisions. *Journal of the Royal Statistical Society. Series B (Methodological)*, *14*(1), 107–114.

Jaffray, J.-Y., & Karni, E. (1999). Elicitation of Subjective Probabilities When the Initial Endowment is Unobservable. *Journal of Risk and Uncertainty*, *18*(1), 5–20.

Kadane, J. B., & Winkler, R. L. (1988). Separating Probability Elicitation From Utilities. *Journal of the American Statistical Association*, *83*(402), 357–363.

Kalai, E., & Ledyard, J. O. (1998). Repeated Implementation. *Journal of Economic Theory*, *83*(2), 308–17.

Kalai, E., & Lehrer, E. (1993). Rational learning leads to Nash equilibrium. *Econometrica*, *61*(5), 1019–1045.

Karni, E. (1999). Elicitation of Subjective Probabilities when Preferences are State-Dependent. *International Economic Review*, *40*(2), 479–486.

Karni, E., & Schmeidler, D. (1993). On the Uniqueness of Subjective Probabilities. *Economic Theory*, *3*(2), 267–77.

Murphy, A. H., & Winkler, R. L. (1984). Probability Forecasting in Meterology. *Journal of the American Statistical Association*, *79*(387), 489–500.

Nau, R. (1995). Coherent Decision Analysis with Inseparable Probabilities and Utilities. *Journal of Risk and Uncertainty*, *10*(1), 71–91.

O'Carroll, F. M. (1977). Subjective Probabilities and Short-Term Economic Forecasts: An Empirical Investigation. *Applied Statistics*, *26*(3), 269–278.

Savage, L. J. (1971). Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association*, *66*(336), 783–801.

Smith, C. A. B. (1965). Personal Probability and Statistical Analysis. *Journal of the Royal Statistical Society. Series A (General)*, *128*(4), 469–499.