DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES
CALIFORNIA INSTITUTE OF TECHNOLOGY
PASADENA, CALIFORNIA 91125

A GROVES-LIKE MECHANISM IN RISK ASSESSMENT

Talbot Page

SOCIAL SCIENCE WORKING PAPER 507

December 1983

Abstract

This paper links two research areas that have developed independently -- incentives compatibility for public goods and elicitation of subjective probabilities. An analogy between incentives for reporting information in the two areas leads to the discovery of a new mechanism, based on the Groves mechanism, for eliciting subjective probabilities. In the public goods area, the analogy provides an extention of the basic theorem of truthful response to the more general case when one´s true valuation of the public good is state dependent. In the risk assessment area, the analogy provides a generalization of the traditional reporting mechanisms, proper scoring rules, and in doing so establishes a representation theorem for them.

The paper considers three goals which a principal might have while choosing a transfer mechanism. These goals are: information pooling, strong research incentives for the agents, and identifiability of the agent with the best information. For two structures of information and the specific cases considered, the new mechanism performs well, compared with four traditional mechanisms, in achieving these goals.

A Groves-like Mechanism in Risk Assessment

Talbot Page *

The purpose of this paper is to apply ideas from the theory of incentive compatibility for revealing willingnesses to pay for a public good to the theory of revealing probabilities in risk assessment. The link is achieved by an analogy between two models, one for public goods and the other for risk assessment. The analogy leads to the discovery of a new mechanism, based on the Groves mechanism, for eliciting subjective probabilities; a generalization of the basic theorem on the Groves mechanism; and generalization of proper scoring rules. Analysis of a few specific cases suggests that the new mechanism performs well, compared with existing mechanisms.

In risk assessment we are interested in developing estimates of the probability of some event, which might be rare (a reactor core meltdown) or unique (chemical X is a carcinogen). There is a principal who makes use of these estimates for some decision. We assume the principal has no direct information of his own on the probability of the event, but relies on the probability estimates provided by agents or assessors. The assessors have information on the state of nature and form inferences on the probability of the event. The assessors may have differing amounts and qualities of information, or they may have the same information but interpret if differently.

The principal may have several goals in mind: elicitation of a consensus estimate based on the pooled information; provision for strong research incentives, or identification of the assessor with the best information of inference skills (the principal may want to hire one of two or more assessors.) As a means toward some mix of these goals, the principal agrees to reward the assessors by a transfer mechanism, which is a function of the revealed probability estimates and whether or not the predicted event occurs in the appropriate time interval. It is assumed that the assessors have a single goal; to do "as best they can" in response to a given transfer mechanism. The emphasis of the paper will be on this latter question — how the assessors might respond to or manipulate various possible transfer mechanisms. A rationale for the focus is that no matter what mix of goals the principal might have, he will be unable to pick a transfer mechanism until he knows something of how the agents might respond to it. In a limited way we will also explore how well various mechanisms perform in meeting possible goals of the principal.

Concern with incentive compatibility in eliciting probability judgments has a venerable history. Thomas Bayes identified the method of revealing personal probabilities by means of choices of bets with differing monetary odds, and Ramsay, de Finetti, and Savage greatly developed this idea. In 1950 Brier [1950] proposed a method of eliciting probabilistic weather forecasts and a verification system which would be immune to manipulation, or as he put it, "playing the system." Savage (1971) devoted his last paper to mechanisms for eliciting truthful revelation of an individual's judgmental probability (such mechanisms came to be known as proper scoring rules). Recently Grether (1981) has developed a procedure, not depending on risk neutrality, for truthful revelation of probability estimates.

Interestingly, theorists in the two fields — demand revelation for public goods and elicitation of subjective probabilities — trace their work to a common source. Savage (1971) based his characterization of proper scoring rules on Marschak's seller's price auction. Similarly, Green and Laffont (1979, page 36), found Marschak's seller's price auction to contain the "essence" of the Groves mechanism for public goods. But Green and Laffont do not refer to the problem of public goods demand revelation, and even the references to Marschak are different. By and large it appears that the theories of the two types of revelation have grown up separately, with separate terminologies.

Section I sets out the risk assessment model and defines the new elicitation mechanism analogous with the Groves mechanism. Section II derives the Bayesian strategy for this mechanism. Section III carries the analogy back to Groves mechanism in the public goods model and derives the Bayesian strategy for agents who do not know their valuations with certainty. With the new elicitation mechanism in hand, Section IV turns to the traditional mechanisms for probability elicitation, proper scoring rules, and draws the link between the Groves-like mechanism and the traditional mechanisms. Section V relaxes the assumption of risk neutrality.

Sections II, IV and V emphasize the agents' problem, to react as best they can to a given reward rule. In Sections VI, VII and VIII the emphasis shifts toward the principal's problem and his concerns in choosing a reward rule. These sections briefly examine the performance of various mechanisms in meeting possible goals of the principal:

information pooling (Section VI); strong research incentives relative to the principal's expected budget (Section VII); and identification of the best assessor (Section VIII).

I. The Model

The model is basically in the form of a Bayesian game as defined by Myerson [forthcoming], and first developed by Harsanyi [1967-8]. The main difference is we focus on the assessors' beliefs about the others' actions rather than about the others' information sets.

We begin with an event X and a set of risk assessors $N=\{1,\dots,n\}$. If the event X occurs we write $X=1$, if not $X=0$. Each assessor has information upon which he forms his judgment of the probability of $X=1$. Each assessor i reports r_i, his revealed probability of $X=1$. The r_i are reported before X is observed, and unless otherwise stated assessor i reports r_i without observing the other assessors' reports. After X is observed each assessor is rewarded on the basis of his and the others' reports and whether $X=1$ or $X=0$. For example, each of n weather forecasters makes a probabilistic prediction of rain tomorrow. Various principals use this assessment to decide whether or not to harvest a crop today. The following day the forecasters are rewarded on the basis of the revealed predictions $(r_1,\dots,r_n)$ and on whether or not it rained. For a second example, each of n toxicologists assess the probability that a chemical will score positive in a bioassay. If the assessed probabilities are high, the chemical will be restricted from the market during the

testing program (testing may take several months, even up to four years). If the assessments are low, the chemical will not be restricted during testing. Once the test is completed, the assessors are rewarded on the basis of the reported $(r_1,\dots,r_n)$ and on the outcome of the test.

The reward rule is a function $t=(t_1,\dots,t_n)$ where $t_i=t_i(r_1,\dots,r_n,X)$ is the reward or transfer to assessor i. The strategy space for each i is [0,1].

Write $p=Pr(X=1)$. With his limited and costly information, i may not know p precisely. He expresses his uncertainty by forming $f_i(p)$, his subjective (generalized) p.d.f. on p. In doing so, he forms his expectation of X which is the same as his expectation of p $(E_i(X)=\int(1p+0(1-p))f_i(p)dp=\int pf_i(p)dp=E_i(p))$. We will write this expectation as $\overline{p}_i=E_i(p)$. An assessor's expectation of p, and more generally his strategic response to a reward rule, depend upon the general structure of information as well as his specific, private information.

Assessor i's information may be focused on the probability of X or on X itself, as we can see in the following two examples.

<u>Information Structure p</u>. (Assessor i's private information on p). Each assessor i knows that event X is a trial from a Bernoulli process with unknown parameter p which is drawn from a uniform distribution over [0,1]. He knows that once p is drawn (but not observed by any of the assessors) he will observe M_i independent trials from the same process. He knows that j will observe another M_j independent trials from

the same process. This knowledge of the general structure is common knowledge for the n assessors. (i knows that j knows this structure, i knows that j knows that i knows, etc.) In addition, each i has some private information. Each i observes y_i, the number of successes in his M_i trials. Assessor i does not know y_j. Before acquiring his private information i forms a diffuse prior on p over [0,1]. After acquiring his private information i forms his subjective p.d.f. $f_i(p)$, which is a beta distribution with parameters $1+y_i$ and $1+M_i-y_i$.

Information Structure X (Assessor i's information is on the event itself). Each assessor knows that X is a Bernoulli trial with probability $\overset{*}{p}$. Each assessor knows $\overset{*}{p}$. If X=1 i observes M_i trials from a Bernoulli process with parameter a; if X=0 i observes M_i trials from a Bernoulli process with parameter 1-a (where a>.5). Thus much information is held in common knowledge with the other assessors. Assessor i's private knowledge is y_i, the number of successes in the M_i trials (he does not know which process he observes).

In the first structure $\overline{p}_i$ is the mean of a beta distribution with parameters $1+y_i$ and $1+M_i-y_i$; thus $\overline{p}_i=E_i(p)=(1+y_i)/(2+M_i)$. ($E_i(p)$ depends upon y_i but we suppress this argument.)

In the second structure $\overline{p}_i$ is obtained by Bayes theorem.

$$\overline{p}_i = Pr(X=1|y_i) = Pr(y_i|X=1)\overset{*}{p}/(Pr(y_i|X=1)\overset{*}{p} + Pr(y_i|X=0)(1-\overset{*}{p}))$$

$$= 1/[1+DA^{M-2y_i}] \quad \text{where} \quad D = (1-\overset{*}{p})/\overset{*}{p} \quad \text{and} \quad A = a/(1-a).$$

As an additional part of the assessors' common knowledge we will assume that each i knows the reward rule chosen by the principal. In contrast, the principal is in the dark — he must choose the reward rule and make use of the reported r_i's without knowing either the general structure of information nor the specific private information held by the individual assessors. The usual domain of transfer mechanisms considered by theorists and recommended for the principal to choose from is the set of proper scoring rules. We will define and discuss such rules later, but first we explore the properties of the new mechanism, based on an analogy with the Groves mechanism for public goods.

The idea is as follows. Each i reports a probability r_i of the event without knowing the others' reported assessments. For each assessor i, the consensus of the other n-1 assessors is defined and specified q_i. (The consensus of others can be defined in many ways. It might be the average of the others' reports, the median of their reports, or the geometric mean of the reported odds.) Then if the event occurs, i wins if $r_i \geq q_i$ (his reported probability is higher than the others' consensus) and i loses if $r_i < q_i$. And if the event does not occur, i wins if $r_i < q_i$ and i loses if $r_i \geq q_i$. How much he wins in each case is determined by the others' consensus. The amount of a win and the resolution of ties is provided by the definition:

<u>Definition</u>: The Groves-like risk assessment mechanism (mechanism G for short) is the function $t=(t_1,\ldots,t_n)$ where

$$t_i = \begin{cases} 1 - q_i & \text{if } X = 1 \text{ and } p_i \geq q_i \\ q_i & \text{if } X = 0 \text{ and } p_i < q_i \\ 0 & \text{otherwise} \end{cases}$$

The consensus of others q_i is the pivot of the mechanism for i, and as in the Groves mechanism the reward rule is split into two parts. the decision of who wins is determined by all the reports, and X; the decision of how much i gets if he wins is determined by all the reports but i's. Each i knows the aggregation rule defining the consensus, but since he does not know the others' reports in advance, he does not know the consensus at the time he makes his report.

Although i does not know the others' private information nor their reports, on the basis of his information he forms a belief as to their reports and hence as to q_i. We write i's p.d.f. on q_i as $g_i(q_i)$. Turning the matter around, if i learned the consensus of others he might infer something about their information and revise his expectation of p. Thus his beliefs about p and q_i may not be independent and i describes his joint beliefs on p and q_i by the joint p.d.f. $h_i(p,q_i)$. We will write i's expectation of p,

conditional on q_i as

$$\overline{p}_i(q_i) = E_i(p|q_i).$$

It turns out that even though i does not observe q_i, the function $\overline{p}_i(\cdot)$ plays a central role in the Groves-like mechanism.

II. Bayesian Strategy for the New Mechanism.

We define i's Bayesian strategy as his expected utility maximizing strategy. As background assumptions in this and the next few sections we will assume i is risk neutral and has a fixed amount of information (for now we don't consider research incentives). These assumptions mean that for now i's Bayesian strategy is his reporting strategy which maximizes his expected transfer. When i believes that there is at least some chance, however small, that q_i could be anywhere on the unit interval, and when i's conditional expectation $\overline{p}_i(q_i)$ is a continuous function of q_i, the main result is simply stated.

<u>Theorem 1</u>. If $\overline{p}_i(q_i)$ is continuous and $g_i(q_i)>0$ for $0\leq q_i\leq 1$, then i's Bayesian strategy under mechanism G is a fixed point of $\overline{p}_i(\cdot)$.

<u>Proof</u>. Assessor i's expected transfer, for a given true p and q and as a function of his report r (and omitting the subscript i) is

$$E(t|p,q,r) = \begin{cases} (1 - q)p & \text{if } r \geq q \\ p(1 - p) & \text{if } r < q \end{cases}$$

Taking the expectation over p and q, i's expected transfer as a function of r is

$$T(r) = \int_0^1\int_0^1 E(t|p,q,r)h(p,q)d\mu(p)d\mu(q)$$

$$= \int_0^1 d\mu(p) \int_0^r (1-q)ph(p,q)d\mu(q) + \int_0^1 d\mu(p) \int_{r^+}^1 q(1-p)h(p,q)d\mu(q)$$

$$= \int_0^r d\mu(q) \int_0^1 (1-q)ph_{p|q}(p|q)g(q)d\mu(p) + \int_{r^+}^1 d\mu(q) \int_0^1 q(1-p)h_{p|q}(p|q)g(q)d\mu(p)$$

$$= \int_0^r (1-q)\overline{p}(q)g(q)d\mu(q) + \int_{r^+}^1 q(1-\overline{p}(q))g(q)d\mu(q)$$

(1) $$= \int_0^r \overline{p}(q)g(q)d\mu(q) + \int_{r^+}^1 qg(q)d\mu(q) + \int_0^1 q\overline{p}(q)g(q)d\mu(q)$$

So, except on mass points of $g(q)$,

(2) $$T'(r) = (\overline{p}(r) - r)g(r)$$

and since $g(r)>0$ for all $0\leq r\leq 1$, the maximum of $T(r)$ occurs at the fixed point $\overline{p}(r)=r$ or at 0 or 1. If the maximum occurs at 0, we must have $T'(r)\leq 0$ in a neighborhood (to the right) of 0. In this neighborhood $\overline{p}(0)-0\leq 0$. Since $\overline{p}(\cdot)$ is an expectation of a probability, $\overline{p}(\cdot)\geq 0$, so $\overline{p}(0)=0$, a fixed point. If the maximum occurs at 1, we must have $T'(r)\geq 0$, for a neighborhood of 1 (to the left), in which case $\overline{p}(1)-1\geq 0$. But $\overline{p}(\cdot)\leq 1$, so $\overline{p}(1)=1$, a fixed point.

In figure 1, i's conditional expectation of p is drawn as increasing in q (i would revise his expectation of p upward upon learning that the consensus of others was larger) but not as fast as q. While these appear to be plausible conditions neither are required by the theorem. (The second condition ($\overline{p}_i'(q_i)<1$) insures that there is no more than one fixed point, and thus the Bayesian strategy is unique, since clearly there is a fixed point, $\overline{p}_i(\cdot)$ being a continuous mapping from the unit interval into itself.) The difference between $T(r_2)$ and $T(r_1)$ is the little slice A weighted by $g(r_1)$ (and if there is a mass point in $[r_1,r_2]$ the weight is augmented). Clearly as long as $\overline{p}(\cdot)$ is above the diagonal for r in $[r_1,r_2]$, $T(r_2)>T(r_1)$.

Assessor i's Bayesian strategy can be characterized as his truthful, regret avoiding strategy (a similar regret avoiding strategy turns up in Milgrom [1981]). Consider what could happen if i could learn q_i after reporting r_i but before observing X. If he reported r_1 and then observed q_1 he would experience regret in the following sense. Upon learning q_1, i would revise his expectation of p to $\overline{p}_i(q_i)$. But in Figure 1 $r_1 < q_1$ so his expected transfer is $(q_1)(1-\overline{p}_i(q_1))=q_1-q_1\overline{p}_i(q_1)$. But in the figure $\overline{p}_i(q_1)>q_1$, so i would experience regret. He could have done better by reporting some $r>q_1$, in which case his expected transfer would be $(1-q_1)\overline{p}_i(q_1)=\overline{p}_i(q_1)-q_2\overline{p}_i(q_1)$, which is larger than his expected transfer obtained by reporting r_i. However, if i reported r^o he would not be subject to later regret. If he reported $r=r^o$ and q_i turned out to be less than r^o, then $\overline{p}_i(q_i)>q_i$, and i is

Figure 1

Bayesian Strategy for Mechanism G

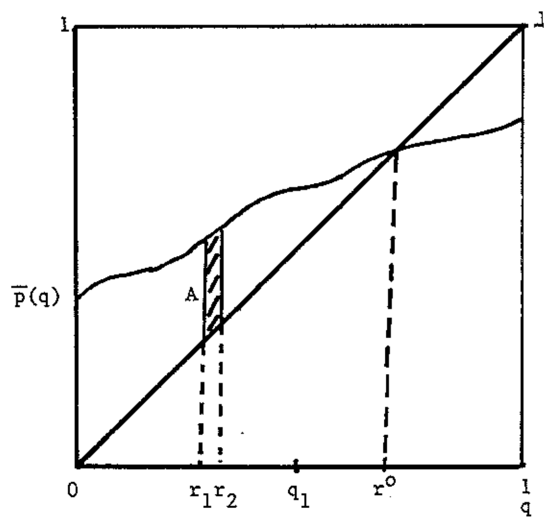

better off by a report greater than q_i, which he has done. If q_i turns out to be greater than r^o, then $\overline{p}_i(q_i)<q_i$, and i is better off by a report less than q_i, which he has done. Thus i's Bayesian strategy is to report that conditional expectation which would lead to no revision in strategy if information on q_i later became available. Theorem 1 says that this fixed point strategy is i's expected transfer maximizing strategy when he must choose r_i before observing q_i. Note that i's Bayesian strategy is a dominant strategy. He can do no better than report the fixed point of $\overline{p}_i(\cdot)$ no matter what probability weight he puts on others' actions.

To illustrate the theorem, and for later use, we derive Bayesian strategies for simple cases. For Information Structure p, consider the case where there are two assessors with symmetric information ($M_1=M_2=M$); the consensus of others is defined by $q_1=r_2$ and $q_2=r_1$; strategies are symmetric ($S_1(\cdot)=S_2(\cdot)=S(\cdot)$, where $S_i(y_i)=r_i$ is i's strategy function); and $S(\cdot)$ is invertible. Then assessor 1's conditional expectation is

$$\overline{p}_1(r_2) = (1 + y_1 + S^{-1}(r_2))/(2 + 2M)$$

Although the theorem is stated for $g(\cdot)>0$, we can use it to find a candidate equilibrium pair and then check to see if this is a Bayesian equilibrium. At the fixed point $\overline{p}_1(r_1)=r_1=(1+y_1+S^{-1}(r_1))/(2+2M) = (1+2y_1)/(2+2M)$. So we have the candidate strategy pair $(S_1,S_2)=((1+2y_1)/(2+2M),(1+2y_2)/(2+2M))$. But if i believes that $S_2(y_2)=(1+2y_2)/(2+2M)$, i maximizes his expected transfer by adopting $S_1(y_1)=(1+2y_1)/(2+2M)$ as his strategy function. Thus (S_1,S_2) is a Bayesian equilibrium. (At this equilibrium i believes that $g(\cdot)$ is concentrated on M mass points.)

For Information Structure X, consider the case where $M_1=M_2=M$; $q_1=r_2$ and $q_2=r_1$; and the strategies are symmetric and invertible. Then, recalling the previous notation of this information structure,

$$\bar{p}_1(r_2)=1/[1+DA^{2M-2y_1-2S^{-1}(r_2)}]$$

At the fixed point $\bar{p}_2(r_1)=r_1=1/[1+DA^{2M-2y_1-2S^{-1}(r_1)}]$

and $S(y_1)=r_1=1/[1+DA^{2M-4y_1}]$

III. Bayesian Strategy in the Groves Mechanism

Theorem 1 was derived after positing a mechanism for eliciting probabilities analogues with the Groves mechanism. In this section, with the idea the theorem in hand, we run the analogy the other way, and derive analogous theorem for the Groves mechanism in the public goods model. (We follow the discussion and as much as possible the notation of the treatment in Green and Laffont [1979].) To obtain the analogy we generalize the model by relaxing one of its usual assumptions. In the basic theorem which shows the Groves mechanism to have truthful dominant strategies (Green and Laffont Theorem 3.1) it is assumed that each i knows v_i, his true valuation of the public good, with certainty. However, suppose the public good in question is a dam which may or may not be built and i's valuation depends upon the unknown future state, which describes among other things whether there

will be much or little rain in the next ten years. For our generalization we will not assume that i knows v_i with certainty. Instead we will assume that i forms a subjective on the unknown future state. This induces a p.d.f. on v_i and the sum of others' reported willingnesses to pay. To facilitate the analogy we will write the sum of the others' willingness to pay as q_i, i's joint p.d.f. on (v_i,q_i) as $h_i(v_i,q_i)$, and his marginal p.d.f. on v_i as $f_i(v_i)$.

Analogously with the risk assessment model, define

$\bar{v}_i + \int vf_i(v)d\mu(v)$ i's expected valuation of the public good (in the absence of knowledge of the others' willingness to pay) and

define $\bar{v}_i(q_i) = \int vh_{v|q_i}(v|q_i)dv$, i's valuation of K, conditional on the knowledge of the sum of others' willingness to pay. The transfer to i in Groves mechanism is

$$t_i = \begin{cases} q_i + H_i & \text{if } q_i + w_i \geq 0 \\ H_i & \text{if } q_i + w_i < 0 \end{cases}$$

where w_i is i's reported willingness to pay and H_i is an arbitrary function of the others' reported willingnesses to pay. The public good is produced if $q_i+w_i \geq 0$ in which case i enjoys utility v_i from it. The public good is not produced if $q_i+w_i<0$. With additive, separable utility, i's utility is

$$U^i = \begin{cases} v_i + q_i + H_i & \text{if} \quad q_i + w_i \geq 0 \\ H_i & \text{if} \quad q_i + w_i < 0 \end{cases}$$

<u>Theorem 2</u>. If $v_i(\cdot)$ is continuous and $g_i(q_i)>0$ $(-\infty<q_i<\infty)$, then agent i's Bayesian strategy is a reflecting point of $v(\cdot)$ or $\pm\infty$ (a reflecting point is a w satisfying $v(-w)=w$).

<u>Proof</u>. Write $V_i = U^i - H_i$. Since H_i is not a function of w_i i's maximizing of his expected utility is the same as maximizing V_i. Omitting the subscript i,

$$E(V) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} Vh(v,q)d\mu(v)d\mu(q)$$

$$= \int_{-\infty}^{\infty} d\mu(v) \int_{-\infty}^{-w} (0)h(v,q)d\mu(q) + \int_{-\infty}^{\infty} d\mu(v) \int_{-w}^{\infty} (v+q)h(v,q)d\mu(q)$$

$$= \int_{-w}^{\infty} d\mu(q) \int_{-\infty}^{\infty} (v+q)h(v,q)d\mu(v)$$

$$= \int_{-w}^{\infty} g(q)d\mu(q) \int_{-\infty}^{\infty} (v+q)h_{v|q}(v|q)d\mu(v)$$

$$= \int_{-w}^{\infty} (\overline{v}(q)+q)g(q)d\mu(q)$$

So, except on mass points

$$\frac{dE[V(w)]}{dw} = (\overline{v}(-w) - w)g(-w).$$

With $g(-w) > 0$, the maximum either occurs where $\overline{v}(-w) = w$ or at $\pm\infty$.

The Bayesian strategy of Theorem 2 has the same interpretation of regret avoidance. In Figure 2, $v_i(q_i)$, i's conditional expected valuation of the public good is depicted as increasing in q_i (again this condition is not required for the theorem) and i's Bayesian strategy is at w^o. If i reports w^o and it turns out that $q_i > -w^o$ the public good is produced. After reporting w^o and learning q_i, i would experience no regret and no wish to change his strategy, because $v_i(q_i)>q_i$. And similarly i would not wish to change his strategy if q_i turned out to be less than $-w^o$ (and the public good was not produced). In contrast if i reported other than w^o, situations could arise where i would wish it possible to have a different strategy.

Note that when i believes that v_i and q_i are independent, $\overline{v}_i(q_i)=\overline{v}_i$ and the reflecting point of $\overline{v}_i(\cdot)$ is $\overline{v}_i$. A special case of independence is when i believes he knows v_i with certainty (so $v_i=\overline{v}_i$). In this case Theorem 2 says that i's Bayesian strategy is to report v_i, no matter what are his beliefs as to the others' actions and information. In this case i's Bayesian strategy is his dominant strategy and Theorem 2 specializes to its old form.

Figure 2

Bayesian Strategy for the Groves Mechanism

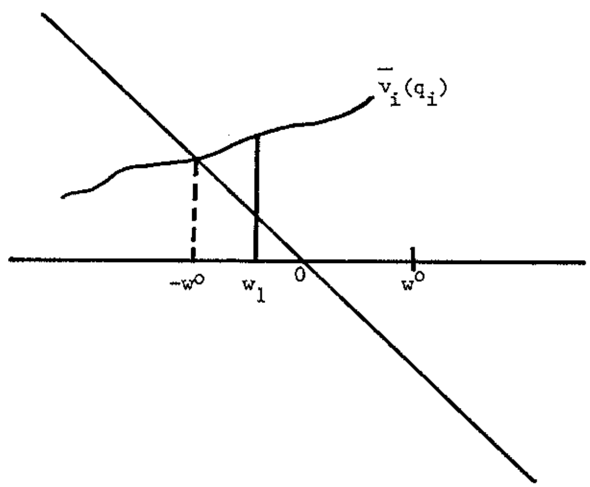

IV. Proper Scoring Rules

We turn now from mechanism G to the traditional class of incentive compatible mechanisms. This class is composed of proper scoring rules. Our purposes in doing so are to draw a link between theories of incentive compatibility in the two areas — demand revelation for public goods and probability elicitation — and to provide a basis for comparing the new mechanism with traditional ones.

A proper scoring rule is a special type of transfer mechanism, distinguished by two features. First, i's transfer is only a function of r_i and X. Second, a proper scoring rule is defined as a rule for which i maximizes his expected transfer by reporting $r_i = \bar{p}_i$. Since i maximizes his expected transfer by $r_i = \bar{p}_i$, no matter what the others report, $\bar{p}_i$ is not only i's Bayesian strategy but also his dominant strategy.

It is well known that there is an infinite number of proper scoring rules; nonetheless, three have received the most attention. These are the

Brier, $$t_i(r_i, X) = \begin{cases} 2r_i - r_i^2 & \text{if } X = 1 \\ 1 - r_i^2 & \text{if } X = 0 \end{cases}$$

Spherical, $$t_i(r_i,X) = \begin{cases} r_i/\sqrt{r_i^2 + (1-r_i)^2} & \text{if } X = 1 \\ (1-r_i)/\sqrt{r_i^2 + (1-r_i)^2} & \text{if } X = 0 \end{cases}$$

and Logarithmic, $$t_i(r_i,X) = \begin{cases} \log(r_i) & \text{if } X = 1 \\ \log(1-r_i) & \text{if } X = 0 \end{cases}$$

The Brier rule has been used extensively in the evaluation of weather forecasters. A fourth example is particularly interesting because Savage (1971) based his construction of proper scoring rules on the idea underlying it. This rule is Marschak's seller's price mechanism, where

$$t_i(r_i,X) = \begin{cases} 1 & \text{if } X = 1 \text{ and } r_i \geq Z \\ q_i & \text{if } r_i < Z \\ 0 & \text{otherwise} \end{cases}$$

and where Z is a random variable uniform over the unit interval.

The idea for this definition, which appears (in different form) in Becker, DeGroot, and Marschak (1964), is as follows. Assessor i is

given a lottery, where he wins $1 if X=1 and 0 if X=0, and asked to reveal a price r_i for which he will sell it. The principal draws a price $Z randomly from the unit interval. Assessor i reveals r_i before knowing Z. If $r_i<Z$ the principal buys the lottery from i, at price $Z. If $r_i \geq Z$ the principal does not buy the lottery, and in that case i receives $1 if X=1 and nothing if X=0. It is easy to check that for all four rules, $E(t_i)$ is maximized at $r_i=\overline{p}_i$.

At first glance it would appear that these proper scoring rules are quite different from mechanism G. Proper scoring rules have different Bayesian strategies, and for each i's transfer is decoupled from the others' actions. Nonetheless, there is a close link, as we shall see in this section.

We begin by noting that the Brier, Spherical and Marsèhak rules, along with mechanism G, are normalized in the sense that transfers are bound between 0 and 1 for each. We will use the concept of a normalized transfer rule later, but for now we need a slightly different concept.

<u>Definition</u>. A proper scoring rule is standardized if $t(0,1)=t(1,0)=0$ and $t(1,1)+t(0,0)=1$.

(A well known property of proper scoring rules is that $t(r,1)$ is an increasing function of r and $t(r,0)$ is a decreasing function, so a standardized rule is also normalized.)

Neither the Brier not the Spherical rules are standardized, but they can be easily standardized by dividing all transfers by 2 (linear transformations do not affect the properness of a rule). Since in the Logarithmic rule $t_i \longrightarrow -\infty$ as $r_i \longrightarrow 0$ when X=1, this rule is not standardizable by a linear transformation. Marschak's rule, being probabilistic is standardized, but as we will see it is much like the Brier rule in expectation.

The following theorem draws a link between mechanism G and a large class of proper scoring rules including the standardized versions of the Brier and Spherical rules. Again we omit the subscript i.

Theorem 3. If $t=t(r,X)$ is a standardized, differentiable proper scoring rule, then it has the same expected transfer as mechanism G for an assessor i who believes $\overline{p}(q)=\overline{p}$ (all $0\leq q\leq 1$) and $g(q)=t_1(q,1)-t_1(q,0)$ where $t_1(q,X)=\frac{\partial}{\partial r}t(r,X)|_q$.

Proof. First we show that $g(\cdot)$ is well defined as a probability density function. For a given p and r the expected transfer of t is

$$E(t|p,r) = pt(r,1) + (1-p)t(r,0)$$

and its expectation over p is

$$(3) \qquad T(r) = \int_0^1 (pt(r,1) + (1-p)t(r,0))f(p)d\mu(p) = \overline{p}[t(r,1) - t(r,0)] + t(r,0)$$

$$T'(r) = \overline{p}[(t_1(r,1)-t_1(r,0)] - t_1(r,0)=\overline{p}g(r) + t_1(r,0)$$

Since T is maximized by $t=\overline{p}$ for any $0 \leq \overline{p} \leq 1$ (t is proper), we have $t_1(\overline{p},0) = -\overline{p}g(\overline{p})$, to satisfy the first order condition, and since this is true for all $0 \leq \overline{p} \leq 1$, we have $t_1(r,0) = -rg(r)$. So

$$T'(r) = \overline{p}g(r) - rg(r)$$

$$T''(r) = (\overline{p} - r)g'(r) - g(r)$$

For $r = \overline{p}$ to be a maximum, $T'' \leq 0$ at $r = \overline{p}$, which implies $g(\overline{p}) \geq 0$. This last condition must hold for $\overline{p}$, so $g(q) \geq 0$ for $0 \leq q \leq 1$

(and $g(q) > 0$ for a unique maximum).

From definition

$$(4) \qquad t(r,1) - t(r,0) = \int_0^r g(q)dq + k_1$$

$$(5) \qquad t(0,1) - t(0,0) = k_1$$

and since $t(0,1) = 0$, $k_1 = -t(0,0)$.

Thus,

$$t(1,1) - t(1,0) = \int_0^1 g(q)dq - t(0,0)$$

or

$$t(1,1) + t(0,0) = \int_0^1 g(q)dq \quad \text{as} \quad t(1,0) = 0$$

But $t(1,1) + t(0,0) = 1$. So we know $g(\cdot)$ is well defined as a probability density function.

Because $t_1(r,0) = -rg(r)$

$$(6) \qquad t(r,0) = \int_r^1 qg(q)dq + k_2$$

But $t(1,0) = 0$, so $k_2 = 0$. From (3), (4), (5), and (6)

$$T(r)= \overline{p}\left(\int_0^r g(q)dq - t(0,0)\right) + \int_r^1 qg(q)dq$$

$$(7) \qquad = \overline{p}\int_0^r g(q)dq + \int_r^1 qg(q)dq - \overline{p}\int_0^1 qg(q)dq$$

which is the same as (1) in Theorem 1, in the case where $\overline{p}(q)=\overline{p}$ for all $0 \leq q \leq 1$ and no mass points in $g(\cdot)$.

We can apply this theorem to the standardized version of the Brier rule. From its definition $t_1(r,1) = 1-p$ and $t_1(r,0) = -p$, so $g(r) = 1 - p - (-p) = 1$. In other words the standardized Brier rule is equivalent in expectation to a special case of mechanism G where i believes that q_i is distributed uniformly on the unit interval and that q_i and p are independent (so $\overline{p}(r) = \overline{p}$).

Theorem 3 is a representation theorem. Any differentiable standardized proper scoring rule can be represented by a p.d.f. over the unit interval. Alternatively the theorem says that mechanism G can be used to generate proper scoring rules. The principal chooses some $g(\cdot)$, a p.d.f. over the unit interval, and announces that instead of defining q_i to be the consensus of others, q_i is to be drawn from a distribution with density $g(\cdot)$. In that case i believes (rationally) that p and q_i are independent, and Theorem 3 applies.

V. Bayesian Strategies Without Risk Neutrality

In this section we use the notion of a normalized transfer rule to relax the assumption of risk neutrality. For normalized rules, such as the Brier, Spherical, Marschak, and mechanism G we can define corresponding lottery versions.

Definition. Let $t = (t_1,\ldots,t_N)$ be a normalized transfer rule. Then the lottery version of t is defined by $t' = (t'_1,\ldots,t'_N)$ where

$$t'_i = \begin{cases} 1 \text{ with probability } t_i \\ 0 \text{ with probability } 1 - t_i \end{cases}$$

Lottery versions are well defined because $0 \leq t_i \leq 1$ for normalized mechanisms and of course $t_i + (1 - t_i) = 1$. In place of risk neutrality we assume von Neuman-Morgenstern utilities.

Theorem 4. If i has a von Neuman-Morgenstern utility function his Bayesian (expected utility maximizing) strategy is the fixed point of $\overline{p}_i(\cdot)$ under the lottery version of mechanism G and $\overline{p}_i$ under the lottery version of a proper scoring rule.

Proof. The proof of Theorem 1 goes through as before except what was previously an expected transfer is now a probability of winning a zero-one lottery. By the axiom of monotonicity, when faced with dichotomous lottery i prefers higher probabilities of winning to lower probabilities. Similarly for proper scoring rules, the $r_i = \overline{p}_i$ which maximizes the expected transfer of a proper scoring rule is the same $r_i = \overline{p}_i$ which maximizes the probability of winning the lottery of the lottery version of the proper scoring rule.

The idea is to avoid the assumption of risk neutrality by "paying off in probability." In one form or another the idea can be traced back to Savage [1954], Smith [1961], and Marschak [1975]. Grether's mechanism is a lottery version of Marschak's rule, and is thus an application of the idea. McKelvey and Ordeshook [1984] also "paid in probabilities" in an experimental setting.

VI. Information Pooling

In this and the next two sections we consider three possible goals the principal might have: information pooling, sharp research incentives relative to the principal's budget, and identification of the best assessor.

After the principal receives the reported probabilities he must somehow aggregate them to a consensus estimate of the probability of X. An ideal situation would be if the principal's aggregation of the $(r_1,\ldots,r_n)$ produced the pooled information expectation $E(p|y_1,\ldots,y_n)$. However, since the principal doesn't know the underlying information structure, he doesn't know what aggregation rule, if any, would produce the pooled information expectation. Being in the dark, the principal may choose a simple aggregation rule; the rule must often be discussed or recommended by theorists appears to be an average of the reported probabilities. (Savage [1971], DeGroot [1974]).

A way out of this quandry is for the principal to have the assessors do the information aggregation themselves. Geanakoplos and Polemarchakis [1982] have shown that for two assessors who start with common knowledge of the general (but finite) structure of information and individual private knowledge as well, and who then successively report their posterior probabilities to each other, their posterior probabilities eventually become the same. The final consensus is not necessarily the same as the pooled information expectation of p, and the iterative process can take many rounds to terminate, but under favorable conditions the process terminates in the pooled information expectation in the second round. The simplest of the favorable conditions is that each i's posterior probability be a 1-1 function from his informaion set. McKelvey and Page

[1983] have generalized this result to n assessors and for successive revelations of functions of the reports (instead of revelations of all n reports individually for each round).

These results mean that the principal could use a proper scoring rule to elicit $(\overline{p}_1,\ldots,\overline{p}_n)$, make these reports known, elicit new $(\overline{p}_1,\ldots,\overline{p}_n)$ (conditional on the new information), and so on until all the $\overline{p}_1$ are equal. At this point the principal would not know for sure that he had obtained the pooled information expectation, but if the posteriors were 1-1 functions he would obtain the pooled information expectation in the second round. Alternatively the principal could use mechanism G to elicit the Bayesian strategies $(r_1,\ldots,r_n)$, make them known and elicit a new round of reported probabilities, and so on until all the reports are the same. And similarly if the $S_i(\cdot)$ are invertible, the principal would obtain the pooled information expectation at the second round. A difference between the two types of mechanisms is that under a proper scoring rule, the principal could elicit $(\overline{p}_1,\ldots,\overline{p}_n)$ in the first round and then could elicit just one probability in the second round, saving on his budget, whereas under mechanism G at least two reports have to be elicited in the second round to make the mechanism work.

It is interesting to note how close the two types of mechanisms come to the pooled information expectation in the first round. For Information Structure p and the case $n = 2$, $M_1 = M_2 = M$, and symmetric and invertible strategy functions, recall that

$$\overline{p}_i = (1 + y_i)/(2 + M)$$

$$S_i(y_i) = (1 + 2y_i)/(2 + 2M) \qquad \text{(Bayesian strategy for mechanism G)}$$

and note that the pooled information expectation is

$$(1 + y_1 + y_2)/(2 + 2M).$$

Thus if the principal used mechanism G and took the average of the reports he would obtain the pooled information expectation in the first round. If he used a proper scoring rule he would not.

For Information Structure X and the case n=2, $M_1=M_2=M$, and symmetric and invertible strategy functions, recall that

$$\overline{p}_i = 1/[1 + DA^{M - 2y_i}]$$

$$S(y_i) = 1/[1 + DA^{2M-4y_i}]$$

and note that the pooled information expectation is

$$1/[1 + DA^{2M - 2y_1 - 2y_2}]$$

Thus a simple average of the reports under either proper scoring rules or mechanism G will not yield the pooled information expectation. (In mechanism G, when $y_1 = y_2$ the pooled information expectation is obtained, and more generally for mechanism G, with n=2, if the principal took the geometric mean of the reported probabilities he would obtain the pooled information expectation.)

For these two cases, mechanism G does a little better than proper scoring rules in the first round. But in more complicated cases it is difficult to compute Bayesian strategies for mechanism G and there can be multiple equilibria, in which case information might be lost in an averaging process.

VII. Research Incentives

A possible goal for the principal is to find a transfer rule which maximizes research incentives for the assessors. The search must be over a constrained set; otherwise the principal could increase research incentives without limit by simply scaling up a transfer rule. We begin by looking for the standardized proper scoring rule which maximizes the value of information to i, when his current expectation of p is $\overline{p}_i$. Any transfer t from one of these rules is constrained by $0 \leq t \leq 1$.

By (7) in Theorem 3 i's expected transfer for his Bayesian strategy $r_i = \overline{p}_i$ is (omitting i)

$$T(\overline{p}) = \overline{p}\int_0^{\overline{p}} g(q)dq + \int_{\overline{p}}^1 qg(q)dq - \overline{p}\int_0^1 qg(q)dq$$

The assessor can simply reveal i's current $\overline{p}$, or undertake research to update $\overline{p}$. Suppose, if he undertakes research, the research is positive, and the expectation of p is $p_b > \overline{p}$. And suppose, if the new research is negative, his new expectation of p is $p_a < \overline{p}$. Write the probability of a positive research finding as α. (Then of course $\alpha p_b + (1 - \alpha)p_a = \overline{p}$). Then the expected value of information is

(8) $$VOI = \alpha T(p_b) + (1 - \alpha)T(p_a) - T(\overline{p}) .$$

By Theorem 3 we can associate each standardized proper scoring rule with a p.d.f., and vice versa. Thus the principal's problem is to find the p.d.f. which maximizes the VOI in (8). Theorem 4 provides the condition for this.

Theorem 4. If i's expectation of p is currently $\overline{p}$, the standardized proper scoring rule which maximizes the agent's expected value of information from further research is characterized by an associated p.d.f. which is fully concentrated at $\overline{p}$.

Proof. Applying (7) to (8), for a standardized proper scoring rule,

$$\begin{aligned} \text{VOI} \;=\; & \alpha\left(p_b\int_0^{p_b} g(q)dq + \int_{p_b}^{1} qg(q)dq - p_b\int_0^1 qg(q)dq\right) \\ & (1-\alpha)\left(p_a\int_0^{p_a} g(q)dq + \int_{p_a}^{1} qg(q)dq - p_a\int_0^1 qg(q)dq\right) \\ & - \overline{p}\int_0^{\overline{p}} g(q)dq + \int_{\overline{p}}^{1} qg(q)dq - \overline{p}\int_0^1 qg(q)d(q) \end{aligned}$$

$$\begin{aligned} = \; & \alpha p_b\int_0^{p_b} g(q)dq + (1-\alpha)p_a\int_0^{p_a} g(q)dq - \overline{p}\int_0^{\overline{p}} g(q)dq \\ & + \alpha\int_{p_b}^{1} qg(q)dq + (1-\alpha)\int_{p_a}^{1} qg(q)dq - \int_{\overline{p}}^{1} qg(q)dq \end{aligned}$$

$$\begin{aligned} = \; & \alpha p_b\int_0^{\overline{p}} g(q)dq + \alpha p_b\int_{\overline{p}}^{p_b} g(q)dq + (1-\alpha)p_a\int_0^{\overline{p}} g(q)dq \\ & - (1-\alpha)p_a\int_{p_a}^{\overline{p}} g(q)dq - \overline{p}\int_0^{\overline{p}} g(q)dq \\ & + \alpha\int_{\overline{p}}^{1} qg(q)dq - \alpha\int_{\overline{p}}^{p_b} qg(q)dq + (1-\alpha)\int_{\overline{p}}^{1} qg(q)dq \\ & + (1-\alpha)\int_{p_a}^{p} qg(q)dq - \int_{\overline{p}}^{1} qg(q)dq \end{aligned}$$

$$(9) \qquad = \alpha\int_{\overline{p}}^{p_b} (p_b - q)g(q)dq - (1-\alpha)\int_{p_a}^{\overline{p}} (q - p_a)g(q)dq$$

We are looking for a function $g(q)$ which maximizes (9) subject to $\int_0^1 g(q)dq = 1$ and $g(q) \geq 0$. This is a simple isoperimetric control problem and its solution is a $g(\cdot)$ fully concentrated at $\overline{p}$.

We can interpret Theorem 4 by returning to (7) from Theorem 3. Note that T is a function of r and $\overline{p}$. Write $T = T(r,\overline{p})$ and consider $J(\overline{p}) = T(\overline{p},\overline{p})$, i's expected transfer when he reports his Bayesian strategy $r = \overline{p}$ for a proper scoring rule. From (7) $J''(\overline{p}) = g(\overline{p})$.

Since g is a p.d.f. with $g(\cdot) \geq 0$, i's expected transfer under Bayesian reporting $J(\cdot)$ is a convex function, and more convex where $g(\cdot)$ is more concentrated. Where J is more convex, there is a greater VOI from research which revises $\overline{p}$.

Theorem 4 is not of direct help to the principal, because he is in the dark as to the current $\overline{p}_i$ and thus cannot define individual scoring rules with p.d.f.'s concentrated in the $\overline{p}_i$.

An attractive feature of mechanism G is that it may allow the principal to achieve some of the benefits of Theorem 4, in stronger research incentives, without the principal himself knowing the current $\overline{p}$'s of the agents. If i believes that the consensus of others is likely to be close ot his own current $\overline{p}_i$ -- a plausible assumption -- he will feel some of the incentives described in Theorem 4.

From (2)

$$T''(r) = (\overline{p}'(r)-1)g(r) + (\overline{p}(r) - r)g'(r)$$

Thus i's expected transfer function at the point $r = \overline{p}(r)$ (his Bayesian report) has convexity $(\overline{p}'(r) - 1)g(r)$. The more concentrated $g(r)$ at the point of i's current Bayesian report, the greater the convexity and greater i's incentive to undertake research which might revise his expectation of p. Conversely the more sensitive i's conditional expectation $\overline{p}_i(q_i)$ to q_i (the closer $\overline{p}_i(\cdot)$ to the diagonal in Figure 1, the smaller the convexity and the smaller i's research incentives.

We can see how these two factors trade off in specific cases. In each case in Table 1 assessor 1's information is specified by $M_1 = 5$. Assessor 2's information is better, with M_2=10 or M_2=15. For information Structure X, the case is taken for D=.5 and A=1.5. The difference in expected transfers, as a percent of assessor 1's expected transfer is shown in Table 1. (Comparing the expected transfers in relative terms normalizes differences in expected transfers among the various rules.)

Table 1

Relative Value of Information:
The Difference between Assessor 2's Expected Transfer and Assessor 1's Expected Transfer as a Percent of Assessor 1's Expected Transfer.

	M_2	G	Sphere	Brier	Log	Marschak
Information Structure p	10	12.4	5.2	1.2	.8	.8
	15	18.7	7.3	1.7	1.1	1.1
Information Structure X	10	28.9	5.1	4.1	2.3	2.5
	15	57.1	9.1	7.3	4.3	4.5

$M_1 = 5$

As can be seen, Mechanism G performs better than the others in providing a higher relative value of information to the assessor with the better information.[1]

The Marschak and Brier rules can be compared as follows. For a given p i's expected transfer for the Marschak rule is

$$\int_0^r (p)(1)dq + \int_r^1 (q)(1)dq = pr + 1/2 - r^2/2$$

and, taking the expectation over p, his expected transfer is

$$T_M = \int_0^1 prf(p)dp + 1/2 - r^2/2 = r\overline{p} + 1/2 - r^2/2$$

and for the Bayesian strategy $r = \overline{p}$, his expected transfer is

$$J_M(\overline{p}) = \overline{p}^2/2 - 1/2.$$

$$J''_M = 1$$

In comparison the i's expected transfer for the standardized Brier rule is

$$T_B = \overline{p}(2r - r^2)/2 + (1 - \overline{p})(1 - p^2)/2 = \overline{p}r + 1/2 - r^2/2 - \overline{p}/2$$

and for this rule J''_B also is 1. Note $T_M - \overline{p}/2 = T_B$. These two conditions mean that the VOI, relative to i's expected transfer is everywhere higher for the Brier rule, compared with Marschak's. (If i truly believes $\overline{p} = 0$, he has no need of further research.) Grether's rule is the lottery version of Marschak's rule, so the same comparison can be made between Grether's rule and the Brier rule.

Table 2

Probability of Identifying Assessor 2 After Forty Trials

	M_2	G	Sphere	Brier	Log	Marschak
Information Structure p	5	.446	.456	.443	.436	.453
	10	.554	.548	.484	.466	.472
	15	.602	.584	.507	.482	.479
Information Structure X	5	.449	.448	.445	.438	.453
	10	.700	.614	.593	.538	.513
	15	.841	.737	.704	.621	.563

$M_1 = 5$

VIII. Identification

Suppose that a principal has two assessors as consultants and wishes, after a limited number of assessments, to hire permanently the one who has consistently better information (Roberts [1965] has an early discussion of this identification problem). The principal chooses a transfer rule and after R rounds hires the assessor with the higher sum of transfers over the R rounds. The principal wants to choose a transfer rule under which the assessor who has the best information has a high probability of having the higher sum of transfers. These probabilities are easily computed if the principal uses lottery versions of the transfer rules. In the lottery version i's probability of winning a zero-one lottery is his expected transfer $E(t_i)$ under the original definition of the transfer rule. With the same information structure for R rounds the sum of wins for i is a binomial random variable with parameters R and $E(t_i)$. For the two information structures, the probability that assessor 2 will have the higher sum of transfers after 40 rounds[2] is shown in Table 2. It can be seen that Mechanism G performs better than the proper scoring rules in identifying the best assessor after a limited number of rounds. (Note that in several cases the probability of assessor 2 being ahead after 40 rounds is less than .5; this is because of ties. The frequency of ties can be seen by the first and fourth rows in Table 2.)

While mechanism G performs well compared with the four proper scoring rules in terms of information pooling, research incentives and identifiability, it should be noted that the Bayesian equilibrium

is hard to caluclate when there are more than two assessors and there can be multiple Bayesian equilibria, which can lead to difficulties. In contrast, the Bayesian equilibria are easy to calculate for proper scoring rules and they are unique.

IX. Conclusion

This paper is a bridge between two areas of research: theories of incentive compatibility in the reporting of willingness to pay for a public good and theories of incentive compatibility in the reporting of subjective probabilities. A new mechanism for reporting subjective probabilities is discovered by analogy with the Groves mechanism for public goods, and the Bayesian (also dominant) strategy for it characterized in Theorem 1. The idea of the theorem is then carried back to the public goods model and the basic theorem of the Groves mechanism is generalized in Theorem 2 for the situation where true valuations are uncertain because they depend upon future states of nature. Theorem 3 shows how the new mechanism is a generalization of the traditional elicitation mechanisms for probability reporting. Theorem 4 obtains the condition for maximizing research incentives for standardized proper scoring rules and relates this condition to the new mechanism.

For the cases considered the new mechanism appears to perform about as well as proper scoring rules in pooling information in two or more reporting rounds. In the first round the new mechanism comes closer to pooled information expectation than do proper scoring rules (for the few cases considered). And, for the numerical cases considered the new mechanism performs better than proper scoring rules in allocating a relatively larger expected transfer to the assessor with the best information and in identifying the assessor with the best information over a limited number of rounds.

Footnotes

*This research was supported by the National Science Foundation and by the Mellon Foundation. I would like to thank Jennifer Reinganum, Richard McKelvey, Ed Green, John Ledyard, John Ferejohn, Roger Noll, David Grether and Joshua Foreman for many helpful comments. This research was supported by the National Science Foundation RPA 8114463 and by the Mellon Foundation.

[1] To bound the range of the Logarithmic rule, the rule was truncated by setting $r_i=.02$ for $r_i<.02$. The truncated rule is not proper for $r_i<.02$ but there is little effect of the truncation in Table 1 since $r_i<.02$ arises infrequently for Information Structure X and not at all for Information Structure p. The truncation makes it possible to define a lottery version of the Logarithmic rule and the construction of Table 2.

[2] When $M_1 \neq M_2$ we modified the assumption about symmetric strategy functions. We assumed that the strategies are "symmetrically asymmetric," in the sense that $M_1 S_2^{-1}(\cdot) = M_2 S_1^{-1}(\cdot)$. This reduces to the symmetric case when $M_1=M_2$ and provides for $S_1^{-1}(\cdot)$ and $S_2^{-1}(\cdot)$ having full ranges compared with the expectations on the full information $y_1 + y_2$. This assumption yields Bayesian strategies of $r_1 = (1+(1+M_1/M_2)y_1)/(2\underline{+}M_1+M_2)$ and $r_2 = (1+(1+M_2/M_1)y_2)/(2+M_1+M_2)$ under Information Structure p and $r_1 = 1/[1+A^{M_1+M_2-2y_1(1+M_2/M_1)}]$ and $r_2 = 1/[1+A^{M_1+M_2-2y_2(1+M_1/M_2)}]$ under Information Structure X.

As in Table 1, the Logarithmic rule is truncated by setting $r_i=.02$ for $r_i<.02$.

References

Becker, G.M., M.H. DeGroot, and J. Marschak (1964), "Measuring Utility by a Single-Response Sequential Method," Behavioral Science, Vol. 9, pp. 226-32.

Brier, G. (1950), "Verification of Forecasts Expressed in Terms of Probability," Monthly Weather Review, Vol. 78, No. 1, Jan. 1950.

DeGroot, M.H. (1974), "Reaching a Consensus," Journal of the American Statistical Association, (Theory and Methods), Vol. 69, No. 345, March 1974.

Green, J. and J.J. Laffont (1979), Incentives in Public Decision-Making, North-Holland Pub. Co., New York.

Grether, D.M. (1981), "Financial Incentives Effects and Individual Decisionmaking," Caltech Social Science Working Paper No. 401.

Geanakoplos, J.D. and H.M. Polemarchakis (1982), "We Can't Disagree Forever," Journal of Economic Theory, Vol. 28, pp. 192-200.

Marschak, J. (1975), "Do Personal Probabilities of Probabilities Have an Operational Meaning?" In Theory and Decision, Vol. 6, No. 2, May 1965.

McKelvey, R.D. and P.C. Ordeshook (1984), " An Experimental Study of the Effects of Parliamentary Rules on Committee Behavior," Journal of Politics, (forthcoming).

Milgrom, P.R. (1981), "Rational Expectations, Information Acquisition, and Competitive Bidding," Econometrica, Vol. 49, No. 4, pp. 921-43. July 1981.

Roberts, H. (1965), "Probabilistic Production," American Statistical Association Journal, March 1965.

Savage, L.J. (1954), The Foundations of Statistics, New York: Dover Publications.

Savage, L.J. (1971), "Elicitation of Personal Probabilities and Expectations," J. of the American Statistical Association, Vol. 66, No. 336.

Smith, C.A.B. (1961), "Consistency in Statistical Inference and Decision," Journal of the Royal Statistical Society, Ser. B, Vol. 23, pp. 1-25.