ELICITATION OF SUBJECTIVE PROBABILITY DISTRIBUTIONS AND VON NEUMANN-MORGENSTERN UTILITY FUNCTIONS

by

Charles A. Holt, Jr.

Discussion Paper No. 79-128, December 1979

Center for Economic Research
Department of Economics
University of Minnesota
Minneapolis, Minnesota 55455

Abstract

In many situations, one individual (the "elicitor") would like to discover the precise form of a subjective probability distribution or of a von Neumann - Morgenstern utility function of another individual (the "subject"). This paper presents a general procedure or "scoring rule" which can be used by the elicitor to induce the subject to reveal a probability distribution representing the subject's beliefs about some uncertain event. This procedure contains several other known scoring rules as special cases. By reversing the roles of probability and utility, this procedure can also be used to induce a risk averse subject to reveal a utility function representing the subject's choices in risky situations.

Elicitation of Subjective Probability Distributions and von Neumann-Morgenstern Utility Functions*

by

Charles A. Holt, Jr.

It is well known that, under fairly general conditions, an individual's beliefs about the possible realizations of a random variable can be represented by a "subjective" probability distribution. In addition, there are sets of axioms (about rational behavior in risky situations) which imply that an individual's optimal decision is a decision which maximizes the expected value of a von Neumann - Morgenstern utility function.[1] In many situations, an economist or a decision-maker (the "elicitor") would like to discover the precise form of the subjective probability distribution or the utility function of another individual (the "subject"). In economic applications, the elicitor is often the owner of a firm or some other resource, and the subject is a hired manager.

This paper presents a procedure which can be used by the elicitor to induce the subject to reveal 1) a probability distribution representing the subject's beliefs about some uncertain event, and 2) a von Neumann - Morgenstern utility function representing a risk averse subject's attitudes toward risk.

I. Introduction

This section contains a brief introductory discussion of some approaches to the elicitation of information about probability distributions and utility functions. Probability elicitation schemes are commonly called "scoring rules," and the essential nature of a scoring rule can be illustrated with a simple example. Suppose the subject is a plant manager, and

*I am glad to acknowledge helpful discussions with James Jordan, George Mason, David Scheffman, Helen Tauchen, William Thomson, and Robert Townsend. I also received many thoughtful comments and criticisms on an earlier version of this paper from William R. Johnson, Roger Sherman, John Whitaker, and other members of the Microeconomics Workshop at the University of Virginia. Marie Childress assisted in the preparation of this manuscript.

there are two possible realizations of a random variable which affects the plant's output. Let the subject's subjective probability of a high output be denoted by p, so the corresponding probability of a low output is $1 - p$. The elicitor in this example is a central planner who asks the subject to report the probability that the plant's output will be high. This reported probability will be denoted by q. Suppose the elicitor can force the subject to accept a "reward" of $\ln(q)$ dollars if the output is subsequently observed to be high, and a reward of $\ln(1 - q)$ dollars if the output is observed to be low. The "ln" notation denotes a natural logarithm operator, so the subject is penalized severely with a large <u>negative</u> reward if the realized output is an event which was reported to have a relatively low probability. If the subject is risk neutral, then expected utility maximization is equivalent to choosing the reported probability q which maximizes the expected value of the reward payment:

$$p \ln(q) + (1 - p) \ln(1 - q)$$

A necessary and sufficient condition for an interior solution is:

$$\frac{p}{q} - \frac{1 - p}{1 - q} = 0$$

Thus the reported probability for each event is equal to the subject's subjective probability for that event. The reader can easily verify that this result is valid for all discrete probability distributions when the money reward payment in any realized event is the natural logarithm of the probability the agent reported for that event. The fact that a logarithmic reward function induces correct reporting of probabilities is well known; see Savage (1971), section 9.4.

This logarithmic scoring rule can also be used for continuous distributions: Suppose that the plant's output is the realization w^* of a continuous random variable w, and the realized output can be anywhere in the interval $[a, b]$. The subject's subjective beliefs are represented by a subjective probability distribution with a density function denoted by $f(w)$. The elicitation procedure is to have the subject report a density function $g(w)$ when it is understood that the subject's monetary reward will be $\ln(g(w^*))$ dollars when w^* is observed. A risk neutral subject would report a density function which maximizes the following functional:

$$\int_a^b \ln(g(w))\ f(w)\, dw \tag{1}$$

This is a special case of the variational problem to be analyzed in the next section. It is not surprising that the optimal reported density turns out to be the subject's subjective density when the reward function is logarithmic.

Suppose that the subject is not risk neutral and that the subject's attitudes toward risk can be represented by a von Neumann - Morgenstern utility function denoted by $U(\cdot)$. Also let the subject's monetary reward be a function $\pi(\cdot)$ of the reported density at w^*. Thus if $w = w^*$, the subject's composite reward function, denoted by $\psi(\cdot)$, can be expressed: $\psi(g(w^*)) = U(\pi(g(w^*)))$. If the elicitor knows the subject's utility function, then it is possible to specify $\pi(\cdot)$ so that $\psi(g) = \ln(g)$. As indicated above, the optimal reported density would then coincide with the subject's subjective density. But if practical considerations make it impossible to collect arbitrarily large negative rewards, then there may be no practical $\pi(\cdot)$ function which results in a logarithmic composite reward function. When the composite reward function $\psi(\cdot)$ is not logarithmic, then the optimal reported density may differ from the subject's actual subjective density. If the elicitor knows the subject's utility function up to an affine transformation, then the elicitor will know the $\psi(\cdot)$ function, and the reported density will still convey useful information to the elicitor--even though the reported and actual densities are different. For any $\psi(\cdot)$ function in a fairly general class of concave functions, theorem 1 in the next section can be used by the elicitor to infer the subject's actual density function from the optimal reported density function.[2]

I am not aware of any scoring rule procedure which works when the subject's von Neumann - Morgenstern utility function is not known by the elicitor. This is because it is difficult, if not impossible, to take a reported probability distribution and unscramble the separate effects of the subject's risk preferences and probabilistic beliefs. Thus it would be convenient to have scoring rule procedures for eliciting utility functions. This is the topic of section III.

Kenney and Raiffa (1976) have recently summarized current procedures for estimating a person's utility function. In their work, they begin by asking questions which will determine the general qualitative properties of a subject's utility function, properties such as monotonicity, decreasing absolute risk aversion, etc. This enables them to select a parametric family of utility functions that possesses the relevant qualitative characteristics. At this stage, they often work with utility functions with exponential terms. Next, they determine quantitative properties of the person's utility function by measuring specific certainty equivalents. They summarize this process:

> Then using the quantitative assessments, that is, the certainty equivalents, we try to find a specific member of that family that is appropriate for the decision maker. The information on certainty equivalents is used to specify values for the parameters of the original family of utility functions. If we are lucky, we will find a utility function satisfying all the qualitative and quantitative assessments simultaneously. Unfortunately, no general procedure exists either for determining whether a given set of qualitative and quantitative assessments are consistent or indicating an appropriate functional form of the utility function when the assessments are consistent.[3]

The utility elicitation procedures presented in sections III and IV do not have this consistency problem because the elicitor does not have to specify a parametric family of utility functions. These procedures are similar in structure to the probability elicitation procedures discussed previously. First, the subject is asked to announce a function, and the elicitor then uses a random device to compute a monetary reward for the subject according to a probability distribution which depends on the subject's announced function. Of course, the subject must know precisely how the announced function affects the distribution of rewards. Sections III and IV contain useful formulas for constructing the probability distribution of rewards. These formulas cause the optimal announced function to contain all essential information about the subject's utility function. This enables the elicitor, in a single step, to compute the subject's utility function up to an affine transformation. These elicitation procedures are, in this sense, scoring rules for utility functions.

The reader who is primarily interested in practical (as opposed to theoretical) applications may want to skip the continuous formulations of scoring rules and go directly to the simpler discrete formulations in section IV. Applications and limitations of scoring rule procedures are discussed in section V, and the problem of eliciting probability and utility information from the same subject is discussed. Section VI contains a summary.

II. Elicitation of Subjective Probability Distributions

Let w^* denote the realization of a random variable $\tilde{w}$. The subject's subjective uncertainty about $\tilde{w}$ is represented by a cumulative probability distribution function $F(w)$ and a density function $f(w)$; $f(w) \geq 0$ on a closed interval $[a, b]$. The elicitation procedure to be analyzed in this section has four stages:

1) The elicitor announces reward functions denoted by ψ and ϕ.
2) The subject reports a probability distribution of a continuous random variable; the reported distribution and density functions are denoted by $G(\cdot)$ and $g(\cdot)$ respectively.
3) The realized value of $\tilde{w}$, denoted by w^*, is observed by all.
4) The subject receives a monetary reward from the principal which may be a function of both $g(w^*)$ and w^*. This reward is denoted by $\psi(g(w^*), w^*) + \int_a^b \phi(g(w), w)\, dw$.

Notice that the reward function has two parts, and only the term $\psi(g(w^*), w^*)$ depends on the realized value of w. The added generality of having ψ depend directly on w^* will be useful in section III.

A risk neutral subject will report a function $G(\cdot)$ which solves the following variational problem:

$$\underset{G(w) \in A}{\underline{\text{maximize}}} \quad \int_a^b [\psi(g(w), w)\ f(w) + \phi(g(w), w)]\ dw \tag{2}$$

$$\text{Subject to:} \quad G(a) = 0$$
$$G(b) = 1$$

where the set A of feasible $G(\cdot)$ functions must be specified.

There is a useful alternative interpretation of (2) when $\phi(g(w), w)$ is identically zero and the subject is risk averse. In this case, $\psi(g(w), w)$ denotes the subject's utility of the monetary reward determined by $g(w)$ and w. Thus the objective functional in (2) becomes an expected utility expression. The monetary reward function is still specified by the elicitor, but the elicitor cannot infer anything from the reported $G(\cdot)$ function unless the subject's utility function is known (or learned subsequently).

The following theorem can be used to solve the elicitor's problem of inferring the subject's probability distribution $F(\cdot)$ from the reported distribution $G(\cdot)$. In the statement of the theorem, partial derivatives are denoted by appropriate subscripts, and the derivative of $\hat{G}(w)$ is denoted by $\hat{g}(w)$.

Theorem 1: Let $\hat{G}(w)$ be an interior solution to the variational problem in (2) when the feasible set A is some subset of continuously differentiable functions on $[a, b]$. Suppose that $f(w)$ is continuous, $\psi(g, w)$ and $\phi(g, w)$ are continuously differentiable, and $\psi_g > 0$ on some open set which contains $\{(g, w): G \in A, w \in [a, b]\}$. Then

$$f(w) = \frac{\lambda - \phi_g(\hat{g}(w), w)}{\psi_g(\hat{g}(w), w)} \tag{3}$$

where

$$\lambda = \frac{1 + \int_a^b \frac{\phi_g(\hat{g}(w), w)}{\psi_g(\hat{g}(w), w)}\, dw}{\int_a^b \frac{1}{\psi_g(\hat{g}(w), w)}\, dw} \tag{4}$$

for $w \in [a, b]$.

proof: Given the assumed properties of $f(w)$ and the reward functions, the following Euler condition is necessarily satisfied in $[a, b]$:[4]

$$\frac{d}{dw}\left[\psi_g(\hat{g}(w), w)\, f(w) + \phi_g(\hat{g}(w), w)\right] = 0 \tag{5}$$

Thus the expression in brackets in (5) must equal a constant for all values of w. This constant is denoted by λ, so equation (3) follows directly from (5). It is straightforward to show that the appropriate constant (implied by the fact that $f(w)$ integrates to one) is given in equation (4). Q.E.D.

To use theorem 1, it is necessary to specify the feasible set A and the reward functions. If the integrand in the objective functional (2) is concave in g, then the functional is concave, and any local maximizer will be a global maximizer.[5] Besides choosing reward functions with the right concavity properties, it may be desirable to structure the rewards so that the necessary computations are not too messy. For example, suppose that $\psi(g(w), w) = \ln(g(w))$, $\phi(g(w), w) = 0$, and A is the set of continuously differentiable, strictly increasing functions on $[a, b]$. In this case, $\psi_g = 1/g(w)$, $\phi_g = 0$, $\lambda = 0$, and equation (3) becomes: $f(w) = \hat{g}(w)$. This logarithmic scoring rule is "proper" in the sense that the subject's optimal reported density is the subject's subjective density. Another well known proper scoring rule, the "quadratic rule," is constructed: $\psi(g(w), w) = 2\,g(w)$ and $\phi(g(w), w) = -[g(w)]^2$. In this case, it follows from theorem 1 that $\lambda = 0$ and $f(w) = \hat{g}(w)$.

Much of the previous work on probability elicitation is concerned with discovering and comparing scoring rules which induce correct reporting of probability distributions.[6] The implication of theorem 1 is that scoring rules which do not induce correct reporting may also be useful. For example, suppose that $\phi = 0$ and $\psi = [g(w)]^{1-r}$, where $0 < r < 1$. In this case, the implication of theorem 1 is:

$$f(w) = \frac{[\hat{g}(w)]^r}{\int_a^b [\hat{g}(x)]^r \, dx} \tag{6}$$

Thus the subject's density is easily calculated from the reported density, even though the optimal reported density differs from the subjective density. One advantage of this power rule over the logarithmic rule is that the monetary reward will never be negative.

Both the logarithmic and power rules discussed above have the convenient properties that ϕ is identically zero, $\psi_w = 0$, and ψ exhibits constant relative "risk aversion" with respect to g.[7] In other words, $r(g)$ defined in equation (7) is a constant:

$$r(g) \equiv -g[\psi_{gg}/\psi_g] \tag{7}$$

To see why this property is computationally convenient, the reader should differentiate both sides of (3) with respect to w. If $\phi = 0$ and $\psi_w = 0$, the result of this differentiation can be expressed:

$$-f'(w)/f(w) = r(\hat{g}(w))\ [-\hat{g}'(w)/\hat{g}(w)] \tag{8}$$

Note that $-f'(w)/f(w)$ and $-\hat{g}'(w)/\hat{g}(w)$ are the Arrow-Pratt indices of absolute "risk aversion" for $F(w)$ and $\hat{G}(w)$ respectively.[8] When $r(\hat{g}(w))$ is a constant, the subject should report a distribution function $\hat{G}(w)$ with an absolute risk aversion which is proportional to the absolute risk aversion of the subjective distribution function $F(w)$. For example, suppose that $F(w)$ is exponential: $F(w) = 1 - e^{-\alpha w}$ where $\alpha > 0$. This function exhibits constant absolute risk aversion in the sense that $-f'(w)/f(w) = \alpha$. It follows from (8) that the optimal reported distribution function is exponential: $\hat{G}(w) = 1 - e^{-\frac{\alpha w}{r}}$ where r is the relative risk aversion constant for the $\psi(\cdot)$ function.

III. Elicitation of Utility Functions

Suppose that a subject's preferences among money lotteries can be represented by a von Neumann - Morgenstern utility function defined on $[a, b]$. Let this function be denoted by $U(w)$, where w denotes money income. Without loss of generality, let $U(a) = 0$ and $U(b) = 1$. This section contains a mechanism for inducing the subject to report a function $G(w)$ which reveals all essential information about $U(w)$.

In the previous section, the subject's expected utility is determined by two factors: the monetary payoff functions, ψ and ϕ, and the subject's subjective probability distribution for the random variable $\tilde{w}$. The subject's reported function $G(w)$ affects the payoff functions but not the probability distribution. The roles of probability and utility can be reversed by letting the subject's reported $G(w)$ function affect the probability distribution of monetary payoffs, but not the monetary payoff functions. The specific mechanism to be considered has four stages:

1) The elicitor specifies a function $\theta(g(w), w)$. This function is revealed to the subject.
2) The subject reports a function $G(w)$ which is strictly increasing on $[a, b]$. In addition, it is required that $G(a) = 0$ and $G(b) = 1$.
3) The elicitor then uses a random device to generate the realization of a random variable with a distribution function $\theta(g(w), w)$, where $g(w)$ is the derivative of $G(w)$.
4) The subject's <u>money</u> payoff is the realized value of the random variable obtained in stage 3.

For this mechanism to work, the $\theta(g(w), w)$ function must be a valid probability distribution function for any $G(\cdot)$ function reported by the subject in stage 2.

An expected utility maximizing subject will report a function which solves the following variational problem:

$$\underset{G(w) \in A}{\underline{\text{maximize}}} \qquad \int_a^b U(w)\ d\theta(g(w),\ w) \tag{9}$$

$$\text{Subject to:}\quad G(a) = 0$$
$$G(b) = 1$$

where the set A of feasible functions must be specified.

In order to apply theorem 1 to the utility elicitation problem, it is necessary to introduce some assumptions about the subject's utility function. Let $U(\cdot)$ be strictly increasing, continuously differentiable, and concave on $[a, b]$. Hildreth (1977) presents an equilibrium argument that risk aversion (concavity) is not a restrictive assumption in an economy with well developed futures markets. However, the differentiability assumption is restrictive.

If $\theta(g(w), w)$ is non-decreasing in w on $[a, b]$, then the assumption that $U(\cdot)$ is continuously differentiable makes it possible to integrate the objective functional in (9) by parts to obtain:

$$1 + \int_a^b -\theta(g(w),\ w)\ U'(w)\ dw \tag{10}$$

where $U'(w)$ denotes the derivative of $U(w)$. This functional has the same essential structure as that in (2): $-\theta(g(w), w)$ corresponds to $\psi(g(w), w)$, and $U'(w)$ corresponds to $f(w)$.

To apply theorem 1 to the objective functional in (10), it is necessary to specify the feasible set A and an appropriate $\theta(g(w), w)$ function. Let A be the set of strictly increasing, concave functions which have a continuous, bounded derivative on $[a, b]$. Thus if $G \in A$, the right-hand derivative of $G(w)$ at $w = a$, denoted by $g(a)$, will be less than or equal to a finite constant γ: $g(a) \leq \gamma$. Then equation (11) defines a $\theta(g, w)$ function which is twice continuously differentiable and strictly increasing in g for all $G \in A$.

$$\theta(g(w), w) = 1 - \left[\frac{g(w)}{\gamma}\right]^{1-r} \left[\frac{b - w}{b - a}\right] \tag{11}$$

Recall that $g(w)$ is required to be less than or equal to γ, so $\theta(g(w), w)$ is a valid probability distribution function for all $G \in A$. If the subject reports a $G(w)$ function with $g(a) < \gamma$, then $\theta(g(a), a) > 0$ and there is a "spike" of probability at the point $w = a$.

If $\hat{g}(w)$ denotes the derivative of the optimal reported function, then it is straightforward to show that the implication of theorem 1 is:

$$U'(w) = \frac{\dfrac{\hat{g}(w)^r}{b - w}}{\displaystyle\int_a^b \frac{\hat{g}(x)^r}{b - x}\, dx} \tag{12}$$

for $w \in (a, b)$. The subject's utility function can be recovered by integrating (12).

If the right-hand limit of $U'(w)$ at $w = a$ is finite, then it follows from (12) that the right-hand limit of $\hat{g}(w)$ at $w = a$ is finite. Thus it is possible to select a finite value of the constraint constant γ which is not binding. Of course, there is no reason to expect $U'(a)$ to be finite. This difficulty at the end point does not arise in the discrete scoring rule for utility functions presented in the next section.

IV. Scoring Rules for Discrete Cases

Suppose that the elicitor is only interested in obtaining utility numbers representing the subject's preferences at a finite number of money values. Specifically, let the subject's von Neumann - Morgenstern utility numbers for the integers $\{0, 1, 2, \ldots, n\}$ be denoted $\{U_0, U_1, U_2, \ldots, U_n\}$, where $U_w \geq U_{w-1}$ for $w = 1, 2, \ldots, n$. There is no loss of generality if it is also assumed that $U_0 = 0$ and $U_n = 1$. The "marginal" utilities can be defined: $f_w \equiv U_w - U_{w-1}$ for $w = 1, \ldots n$. The subject is assumed to be risk averse in the sense that $f_w \geq f_{w+1}$ for $w = 1, \ldots n - 1$.

The elicitation mechanism has the same structure as those considered previously. First, the elicitor announces a function denoted by $\theta(\cdot)$. The subject is then asked to report a vector of non-negative numbers denoted by $\underline{g} \equiv [g_1, g_2, \ldots, g_n]$, subject to a constraint that these numbers sum to one. Finally, the subject receives a monetary reward which is the realization of a random variable with a distribution determined:

$$\theta_{w-1} = \theta(g_w), \qquad w = 1, \ldots n \tag{13}$$

where θ_w denotes the probability that the reward will be less than or equal to w. Consequently, $\theta_n = 1$.

For example, consider a formula which is similar to some of the scoring rules developed in sections II and III:

$$\theta_{w-1} = 1 - (g_w)^{1-r}, \qquad w = 1, \dots n \tag{14}$$

where $0 < r < 1$. This formula will determine a discrete probability distribution if $0 \leq g_{w+1} \leq g_w$ for $w = 1, \dots n-1$. These inequalities determine the set of feasible $\underline{g}$ vectors which can be reported by the subject in this example.

With this notation, the subject's expected utility will be:

$$\sum_{w=1}^{n} U_w \left[\theta_w - \theta_{w-1}\right]$$

It follows from (13) and the utility normalization that an expected utility maximizing subject will solve:

$$\underset{[g_1,\, g_2,\, \dots\, g_n] \in A}{\underline{\text{maximize}}} \quad 1 - \theta(g_n) + \sum_{w=1}^{n-1} U_w[\theta(g_{w+1}) - \theta(g_w)] \tag{15}$$

$$\text{Subject to:} \quad \sum_{w=1}^{n} g_w = 1$$

where A denotes the set of feasible $\underline{g}$ vectors. The discrete analogue of integration by parts is simply to express the objective function in (15) in terms of the marginal utilities:

$$1 + \sum_{w=1}^{n} \left(-\theta(g_w)\, f_w \right) \tag{16}$$

If the $\theta(\cdot)$ function is strictly convex, then the objective function in (16) will be strictly concave in the choice variables. Let λ denote a Lagrange multiplier associated with the constraint that the choice variables sum to one. Then the following conditions are necessary and sufficient for an interior solution to (15):

$$-\theta'(\hat{g}_w)\, f_w - \lambda = 0, \qquad w = 1, 2, \ldots n. \tag{17}$$

where $\hat{g}_w$ denotes the optimal value of g_w. This is the discrete analogue of the Euler condition in equation (5).[9] It follows from (17) and the fact that $\sum_{w=1}^{n} f_w = 1$ that

$$\lambda = -1/[\sum_{w=1}^{n} 1/\theta'(\hat{g}_w)],$$

and therefore

$$f_w = \frac{1/\theta'(\hat{g}_w)}{\sum_{i=1}^{n} [1/\theta'(\hat{g}_i)]} \tag{18}$$

for $w = 1, \ldots n$. Thus the elicitor can use (18) and the optimal reported values of g_w to infer the subject's utility function:

$$U_w = \sum_{i=1}^{w} f_i, \qquad w = 1, \ldots n.$$

For the specific $\theta(\cdot)$ function in equation (14), equation (18) becomes:

$$f_w = \frac{(\hat{g}_w)^r}{\sum_{i=1}^{n} (\hat{g}_i)^r} \tag{19}$$

Equation (19) and the fact that $\sum_{w=1}^{n} \hat{g}_w = 1$ can be used to show that

$$\hat{g}_w = \frac{(f_w)^{\frac{1}{r}}}{\sum_{i=1}^{n} (f_i)^{\frac{1}{r}}} \tag{20}$$

Note that if $f_1 > 0$ and $f_w < f_{w-1}$ for $w = 2, 3, \ldots n$, then the analagous strict inequalities will hold for the optimal values of g_w. In this case, the $\hat{g}_w$ numbers determined in (20) constitute a unique interior solution to the subject's maximization problem.

To convert the utility elicitation scheme into a probability elicitation scheme, the roles of probability and utility must be reversed. Suppose that the realization w^* of a random variable $\tilde{w}$ can be any integer from 1 to n. The subject's subjective probability that $w^* = w$ is denoted by f_w, and $\sum_{w=1}^{n} f_w = 1$. The first step in the elicitation scheme is for the elicitor to announce a strictly concave function denoted by $\psi(\cdot)$. Then the subject is asked to report n numbers denoted by $\{g_1, g_2, \ldots, g_n\}$, and those numbers must sum to 1. The subject's money payoff, when the realized value w^* is observed, is $\psi(g_{w^*})$. A risk neutral subject's expected payoff is

$$\sum_{w=1}^{n} \psi(g_w)\, f_w \tag{21}$$

This objective is equivalent to the objective function in (16), where $-\theta(g_w)$ corresponds to $\psi(g_w)$. Then equation (18) can be used by the elicitor to infer the subject's subjective probabilities from the optimal reported probabilities. For example, if $\psi(g_w) = [g_w]^{1-r}$ for $0 < r < 1$, then the subject's subjective probabilities can be computed from (19). Thus equation (19) is the discrete analogue of equation (6).

V. Applications

Actual and potential applications of probability scoring rules are discussed in Savage (1971). I am not aware of any previous work on utility scoring rules, and I have not yet attempted to apply the utility scoring rules presented here. The discrete utility scoring rule in section IV is surely easier to administer than the continuous version in section III. This discrete scoring rule would probably require less time than the Keeney and Raiffa iterative procedures described briefly in section II. Speed and simplicity would be important in field situations such as the interviewing of farmers. When the elicitor has plenty of time to spend with the subject, a utility scoring rule procedure may be useful in conjunction with standard iterative procedures.

If scoring rule procedures are to provide good information, the potential monetary payoffs must be large enough to induce the subject to think carefully about the reported function. The scoring rule procedures presented in this paper need not be designed to induce correct reporting of utility functions or probability distribution functions because the elicitor can use the optimal reported function to infer the subject's actual function. Therefore the subject can concentrate on the potential money payoffs without directly considering any preconceived ideas about the form of the function being elicited. For example, a person who believes himself or herself to be risk neutral may allow this preconception to affect choices between risky prospects. In the discrete utility scoring rule procedure, the subject is asked to report a non-increasing sequence of numbers which sum to one, but the subject does not have to know that these numbers will be used to infer a utility function, so preconceptions may not interfere. Indeed, this elicitation procedure can be used when the subject does not know anything at all about utility theory. For the same reason, the probability scoring rules can be used to elicit probability information from subjects who have never been exposed to probability theory.

It is important to notice that the elicitation procedures in this paper have only been analyzed in the context of a single, once-and-for-all interaction between the elicitor and the subject. In each case considered, the subject is solving a single stage optimization problem. Such myopic behavior may be appropriate in an experimental situation. For example, suppose that the subjects are participating in an experiment in which they must make decisions under uncertainty. After the experiment is over, the experimentor could use a utility scoring rule to elicit subjects' utility functions. Then the experimentor could use this utility information (corrected for changes in wealth) to analyze the effects of risk aversion on the subjects' behavior in the experiment. Similarly, a utility elicitation procedure, used after a probability elicitation procedure, may allow the experimentor to adjust the reported probability distribution function for the effects of risk aversion.

Finally, the elicitation results in this paper may turn out to have useful theoretical applications. Economic agents are often assumed to face the problem of choosing a function from some set of feasible functions, e.g., insurance contracts, investment strategies, or portfolios. It may be possible, in some contexts, to use theorem 1 to associate the function chosen by an economic agent (i.e., the optimal function) with the agent's utility or probability distribution function. In fact, the problem of inferring a utility function from an investor's asset demand correspondence has recently been solved by Green, Lau, and Polemarchakis (1979). In theoretical applications of this nature, the continuous versions of the elicitation results may be just as convenient as the discrete results presented in section IV.

VI. Summary

This paper considers a situation in which the elicitor wants to discover the precise form of a <u>function</u> which characterizes a subject's preferences or probabilistic beliefs (i.e., a von Neumann - Morgenstern utility function or a probability distribution function). All of the elicitation procedures developed have a common structure: the subject's expected utility depends on a function reported by the subject, and the elicitor uses the subject's reported function to calculate the subject's actual (utility or distribution) function. By reversing the roles of probability and utility, it is shown that the method of eliciting probabilities is similar to the method of eliciting utilities. The analysis is done for both continuous and discrete formulations, and analogous results are obtained.

References

M. H. DeGroot, Optimal Statistical Decisions, McGraw-Hill, New York, 1970.

J. R. Green, L. J. Lau, and H. M. Polemarchakis, "A Theorem on the Identifiability of the von Neumann - Morgenstern Utility Function from Asset Demands," Economic Letters, 59(3), pp. 217-220.

C. Hildreth, "What Do We Know about Agricultural Producers' Behavior under Price and Yield Instability?," American Journal of Agricultural Economics, 59, December 1977, pp. 898-902.

R. L. Keeney and H. Raiffa, Decisions with Multiple Objectives: Preferences and Value Tradeoffs, John Wiley and Sons, New York, 1976.

D. G. Luenberger, Optimization by Vector Space Methods, John Wiley and Sons, New York, 1969.

J. E. Matheson and R. L. Winkler, "The Elicitation of Continuous Probability Distributions," Technical Report No. 86, Dept. of Statistics, Stanford University, August 1974.

J. W. Pratt, "Risk Aversion in the Small and in the Large," Econometrica, 32, Jan./Apr. 1964, pp. 122-136.

L. J. Savage, "Elicitation of Personal Probabilities and Expectations," Journal of the American Statistical Association, 66, Dec. 1971, pp. 783-801.

W. Thomson, "Eliciting Production Possibilities from a Well-Informed Manager," Journal of Economic Theory, 20, June 1979, pp. 360-380.

Footnotes

[1]For example, see chapters 6 and 7 in DeGroot (1970).

[2]This paper does not deal with the equilibrium determination of reward contracts offered by the elicitor. In a market setting, such contracts would presumably be Pareto optimal in some appropriate set of feasible contracts.

[3]Keeney and Raiffa (1976), p. 197.

[4]The necessity of (5) follows the discussion in Luenberger (1969), section 7.5, especially Lemma 3. Note that the necessary condition requires that the expression in square brackets in (5) be differentiable, even though f(w) may not be differentiable at some points.

[5]See Luenberger (1969), Proposition 1, p. 191.

[6]Matheson and Winkler (1974) contains an interesting discussion of strictly proper scoring rules for continuous probability distributions. In addition to the literature on eliciting probability distributions, there is a separate literature on the elicitation of a fractile of an unknown distribution. See Thomson (1979) and the references therein.

[7]David Scheffman first suggested to me that the constant relative risk aversion property would be convenient. The "Arrow-Pratt" indices of absolute risk aversion and relative risk aversion are defined and analyzed in Pratt (1964). It is important to remember that the ψ function in this paper is not the subject's utility function, so the relative risk aversion of ψ is not a direct measure of the subject's attitudes toward risk.

[8]Of course, $F(\cdot)$ and $\hat{G}(\cdot)$ are not utility functions. But because of the symmetric roles of utility and probability, concepts which are useful in analyzing utility functions may provide insights about probability distribution functions and vice versa.

[9]Note that $-\theta'(\hat{g}_w)f_w$ in (17) corresponds to $\psi_g(\hat{g}(w), w)f(w)$ in (5).