STIG I. ROSENLUND

SOCIAL DARWINIST ASPECTS OF UTILITY AND PROBABILITY

ABSTRACT. Utility and subjective probability are assessed from a normative social darwinist viewpoint. It is shown that utility is essentially equal to monetary reward and that probability must satisfy a reasonable frequency criterion.

1. INTRODUCTION

Bayesianism is unavoidable, as shown by the characterization of admissible decision rules as Bayes rules and limits of such rules. However, there has not, in my opinion, been put forward a satisfactory resolution of the problem of properly choosing prior distributions and utility functions, i.e., the problem of which criteria to apply to assess their usefulness. In this note I shall apply the social Darwinist criterion, i.e., decision rules are judged by the evolutionary advantage they give to the decision-maker. I shall try to deduce these two consequences:

(i) Utility is essentially equal to monetary reward.

(ii) Subjective probability must agree with the frequency interpretation in some reasonable sense, to be made more precise.

2. THE FORMAL SETUP

Consider a subject, such as a family or a tribe or a corporation or a government, confronting a sequence of decision situations. The ith situation is described by a probability space $(\Omega_i, \mathscr{F}_i, P_i)$, a decision space D_i, a reward space R_i, a reward function $r_i : (\Omega_i \times D_i) \to R_i$, and a utility function $U_i : R_i \to R$, the real line.

The measure P_i is necessarily subjective, nevertheless it is more or less suitable for the subject in the struggle for survival and natural selection. It depends in a more or less formal way on the previous experience

Theory and Decision **20** (1986) 123–132.

$(w_1, w_2, \ldots, w_{i-1}) \in \Omega_1 \times \Omega_2 \times \Omega_{i-1}$. One possible construction is to start out with the space $(\Omega_1 \times \Omega_2 \times \ldots, \mathscr{F}_1 * \mathscr{F}_2 * \ldots, Q)$, where Q is the overall prior probability, after which the subject obtains successively

$$P_i(A) = Q(w_i \in A \mid \mathscr{F}_1 * \mathscr{F}_2 * \ldots * \mathscr{F}_{i-1}), \quad A \in \mathscr{F}_i$$

(An activity only aiming at increasing knowledge for better assessment of some future P_i can be incorporated into experiment k for some $k < i$.) For example, the subject can be a Bayesian statistician observing a sequence of random variables X_i which are *i i d* with density $f(\cdot \mid \theta)$ conditionally on an unknown parameter θ with prior density $\pi(\cdot)$. Assume that the reward r_i depends on X_i and the decision d_i, which has to be taken before observation of X_i. Then $\Omega_i = R$ and

$$\begin{aligned} P_i(A) &= Q(X_i \in A \mid X_1, \ldots, X_{i-1}) \\ &= \int_A \int_{-\infty}^{\infty} f(x \mid \theta)\pi(\theta \mid X_1, \ldots, X_{i-1})\mathrm{d}\theta\, \mathrm{d}x. \end{aligned}$$

Regardless of the construction of P_i, it follows from my initial acceptance of Bayesianism that $d_i \in D_i$ should be chosen so as to maximize

$$E(U_i(r_i(w_i, d_i)) = \int_{\Omega_i} U_i(r_i(w, d_i))\mathrm{d}P_i(w).$$

Let the maximum be attained for $d_i = d_i^*$ and write for shortness

$$U_i^* = U_i(r_i(w_i, d_i^*)),$$

the utility realized from the optimal decision. I shall assume that the sequence $\{\mathrm{Var}(U_i^*)\}$ is reasonably bounded from below and from above. The bounding from below is obtained by grouping together decision situations that individually have small variances. If some decision situation is so important that it constitutes a singular survival situation, the outcome of which largely determines whether the subject lives or dies, then there is no reasonable bounding from above. This occurs for example when the U.S. president decides which nuclear buttons to press in response to a radar indication of an incoming enemy missile. In such a case the expected-utility hypothesis is vacuous; the theoretical possibility of constructing a utility function and a subjective probability measure such that it is best, in some sense, for the president to maximize expected utility does not aid in understanding the situation.

3. UTILITY

The social Darwinist criterion implies that utility should be measured in resources for the survival and propagation of the subject in the natural selection. This is a normative view of utility. In any decision situation the outcomes can be ranked according to their survival value for the subject. Thus I dismiss the criticism of those who argue that utility is essentially a multidimensional concept which cannot be brought on a single scale. The normative aspect furthermore implies the irrelevance of much empirical and theoretical work on subjective utility. We should be more concerned with how people should choose than with how they actually choose between probability distributions over a set of rewards. An obvious reflection on for instance the example in DeGroot (1970, pp. 93–94) is that those who choose irrationally will lose in the natural selection in the long run. DeGroot cites factors such as feelings of remorse and ridicule, but such feelings are subject to natural selection to the effect that they become adapted to the subject's ability to survive and reproduce.

Several studies have recently explored the connection between game/decision-theory and evolutionary theory, see for instance Schuster, Sigmund *et al.* (1981), Smith (1982) and Leinfellner (1984). In these studies utility is taken as a synonym for reproductive value. Leinfellner (1984, p. 257), for example, states: "The equivalence of maximizing utility in social games and maximizing survival is based on the equation: utility = survival. The striving to win social games corresponds to the steady reproductive increase in the proportion of the population playing an optimal strategy."

How do utilities, i.e., resources for survival, depend on rewards of money? Traditional theory would have it that utility as a function of money is concave, and Jensen's inequality is called on to explain for example insurance, see DeGroot (1970, pp. 99–100). In economic welfare theory the allegedly decreasing marginal utility of income is used to justify progressive income taxes and transfer payments. I claim, however, that if a decision-maker is given sufficient time to plan for the use of a monetary reward, then utility is approximately proportional to it. The number of offspring afforded increases almost linearly with money within wide limits if the individual is given time to plan. One does, however, wish to reduce uncertainty at the expense of expectation if one cannot know how much income one will get until shortly before one must spend it, or how large outlays one will have until shortly before they occur. In such situations it is proper to apply concave utility functions. However, the financial and social network caters to the need for handling random fluctuations. One must look at the totality of similar individuals; if a decision rule gives a larger total money reward, though with larger random variation between individuals, then that rule is best, for the individuals will find ways to distribute rewards via for instance credit arrangements, so that they can be rationally used for the reproduction of the species. Insurance is important for this purpose. The credit function of insurance is more important than the function of leveling out of risks over individuals; if my factory burns down I have to get cash to rebuild it, whether or not I get it as an insurance reimbursement or as a loan. While there is a considerable need to reduce variability of consumption for a given individual over time, the need to reduce variability over the set of different individuals is less than commonly held.

Naturally people spend much money in ways that are not optimal from the viewpoint of reproduction, but social darwinism is concerned with the long run development of optimal behaviour through natural selection.

Lindley (1975, p. 113) states: "Fines should be in utiles. A wealthy man should pay more for a parking offense than an impecunious student." Such a position is clearly refuted by the arguments above.

Hence we can for practical purposes assume that $R_i = R$ and $U_i(r) = r$, where r stands for money. This, I hope, substantiates my assertion (i). It is implicitly understood above that a rule that maximizes expected monetary reward also in some sense maximizes the sum of a larger number of such rewards. This leads me to a discussion of subjective probability in the next section, which I hope shall substantiate my assertion (ii).

It should be noted that in addition to the immediately realized utility we have the utility resulting from the effect the decision has on the outcomes of future decision situations. For example, I might wish to decrease the expected immediate reward in a business negotiation by taking a tough bargaining attitude, thereby enhancing my prospects in future negotiations. In adding together successive utilities we have to subtract, in computing the ith utility, what we have already counted as utility in decision situations $1, \ldots i - 1$. This dynamic element will be present to some degree in any non-trivial real situation. I do not, however, wish to complicate my notation by introducing an explicit dynamic programming component. This component is to be implicitly understood.

4. PROBABILITY

In the axiomatic Bayesian development of subjective probability and the expected-utility hypothesis, as in DeGroot (1970), no long run frequency considerations are invoked. In social Darwinist theory, however, such considerations are paramount. The larger the total utility from the decisions made by a large number of members of a group, the better the group will fare in the natural selection. (In the present setup the subject is the whole group.) If probability has some reasonable interpretation in terms of frequencies, then the maximization of expected utility will, under the assumption of bounded variances, tend to maximize the sum of utilities from many decisions and hence it will entail evolutionary advantage. If not, the expected-utility hypothesis is vacuous.

I shall try to give a frequency interpretation for a sequence of non-repeatable experiments. This interpretation is tied to utilities and decisions.

DEFINITION: Let

$$\phi_n(A_1, A_2, \ldots, A_n) = n^{-1} \sum_{i=1}^{n} (1_{A_i} - P_i(A_i)),\ A_i \in \mathscr{F}_i.$$

Assume that $\text{limsup}_n |\phi_n|$ is small for those sequencies $(A_1, A_2, \ldots)$ such that $\{w_i \in A_i \text{ for all } i\}$ implies a large value for

$$\text{limsup}_n |n^{-1} \sum_{i=1}^{n} (U_i(r_i(w_i, d_i)) - E(U_i(r_i(w_i, d_i))))|,$$

where $d_i = d_i(U_i, P_i)$ is an *a priori* reasonable decision rule. In particular assume this for $d_i = d_i^*$. Then $(P_1, P_2, \ldots)$ has a reasonable frequency interpretation.

Note that this interpretation has a meaning only for the whole sequence of probability measures. It is an attempt at an answer, albeit vague, to the question posed in Lindley (1978, p. 12):

For example, is it necessary, or even desirable, that a subject who gives the same value x say to the probabilities of each of a number of events should ultimately see a proportion x of them to be true? . . . I experience the same doubts that I have in considering frequency statements in statistics; namely, which sequence of events are we considering when we demand the proportion of true ones be x?

My interpretation does not, however, resolve the problem of determining any single probability P_i.

Reichenbach (1949, §§ 71–73), in discussing the single-case interpretation of probability, claims that for each single statement there is a unique appropriate class of statements of which the statement is a member. The relative frequency of true statements in this class is then the unique probability of the single statement being true. The statement can be "this die will show a six next roll" or "Julius Caesar was in Britain". De Finetti (1974, Ch. 3) takes the opposite stand in asserting that probability has nothing to do with frequencies but is instead always an expression of personal and subjective uncertainty. Both Reichenbach and de Finetti assert that a probability can be given to any statement. I am here trying to develop a synthesis between these two extremes, the frequentist and subjectivist interpretations.

Note that in the strictly frequentist model, where successive experiments are independent, by a strong law of large numbers

$$\lim_n \phi_n(A_1^{(k)}, A_2^{(k)}, \ldots, A_n^{(k)}) = 0 \text{ for all } k \in \{1, 2, \ldots\}$$

almost surely, for any sequence of sequences $(A_1, A_2, \ldots)$. Also

$$\lim_n n^{-1} \sum_{i=1}^{n} (U_i^* - E(U_i^*)) = 0 \quad \text{a.s.}$$

since I assumed bounded variances (it suffices that $\Sigma_1^\infty \operatorname{Var}(U_i^*)/i^2 < \infty$, see Breiman (1968, p. 51)). One cannot require such strict frequency agreement from decision-makers operating in the real world, but some degree of agreement must be present if subjective probability is to be useful at all.

Let me illustrate with a simple example, which will clarify my interpretation. The subject operates a roulette-wheel, on which the only outcomes are green (the zero), red and black, with (subjective) probabilities $P_i(\text{green}) = 1/37$, $P_i(\text{black}) = P_i(\text{red}) = 18/37$. He plays against a single opponent who bets one dollar on either red or black each time. The subject gives the opponent one dollar in addition to the betted dollar if the opponent wins, otherwise the betted dollar is kept. This is the optimal decision d_i^* for the opponent would not, by assumption, bet more if the win was greater and would not bet at all if the win was smaller. We can partition Ω_1 into the sets

$$E_1 = \{\text{the opponent bets red and green comes up}\}$$

$$E_2 = \{\text{the opponent bets red and red comes up}\}$$

$$E_3 = \{\text{the opponent bets red and black comes up}\}$$

$$E_4 = \{\text{the opponent bets black and green comes up}\}$$

$$E_5 = \{\text{the opponent bets black and red comes up}\}$$

$$E_6 = \{\text{the opponent bets black and black comes up}\}.$$

Now $P_i(E_k)$ may vary with i if the opponent is judged by the subject to change his betting inclination as the games roll on. From the judgment that the roulette-wheel is symmetric and that the opponent does not possess paranormal capacities, however, the subject obtains $P_i(E_2 \cup E_6) = 18/37$. Since $U_i^* = -1$ on $E_2 \cup E_6$ and $U_i^* = 1$ on the complement, it is sufficient that $\limsup_n |\phi_n(A_1, A_2, \ldots, A_n)|$ is small when $A_i = E_2 \cup E_6$ for all i, because then

$$n^{-1} \sum_{i=1}^{n} U_i^* = 1/37 - 2\phi_n(A_1, A_2, \ldots, A_n).$$

It is also trivially seen that we cannot require that $\limsup_n |\phi_n|$ be small for all sequences $(A_1, A_2, \ldots)$. With $A_i = \{\text{red comes up}\}$ when red in fact comes up, and $A_i = \{\text{black comes up}\}$ when black or green in fact comes up, we have $\limsup_n |\phi_n| = 18/37$. But we should require that $\limsup_n |\phi_n|$ be small for sufficiently many sequences where A_i is determined before knowledge of the ith outcome.

If it is observed that $|\phi_n|$, for $A_i = E_2 \cup E_6$, does not get sufficiently small as the games roll on, then naturally the subject will change his probabilities, concluding that either the roulette-wheel is unsymmetric or that the opponent has precognitive powers. Thus d_i^* will have to be changed so that less is payed for a win. Some degree of frequency agreement is thus forced upon the subject, lest he wants to go bankrupt.

These observations can be generalized to the general setting. I require that $\limsup_n |\phi_n|$ be sufficiently small for sufficiently many sequences of $A_i = \{U_i(r_i(w_i, d_i)) \in I_i\}$, for different decision rules d_i depending on the utility function and the subjective probability measure and which can reasonably be expected to be satisfactory in some sense. In particular I require it for $d_i = d_i^*$. If the requirement is satisfied the average of a large number of utilities will be approximately maximized by taking $d_i = d_i^*$ and will be approximately $n^{-1} \Sigma_{i=1}^n E(U_i^*)$. It is of course nice if the average is much larger than its subjective expectation, but then there are reasons to believe that the optimal decision rules in general would have been different from d_i^*, leading to a still larger gain.

A well-known example from insurance, where the subject was forced to adjust probabilities to get a better frequency agreement, concerns mortality for on one hand life insurance policies with positive risk sum (meaning the company loses from the death of the insured) and on the other hand policies with negative risk sum. It has been found that mortality is higher for the first kind of policies. This real example exhibits similar, although not paranormal, precognitive powers of the opponent as in the roulette example above.

5. IMPLICATIONS FOR BAYESIAN INFERENCE

In the example of Section 2, where the subject is a Bayesian statistician observing a sequence of variables which are *i i d* conditionally on an unknown parameter θ, my frequency interpretation holds under weak regularity conditions on $\pi(\theta)$ and $f(x \mid \theta)$ guaranteeing the almost sure convergence of posterior distributions to a one-point distribution at the true value of θ. That Bayesians emphasize this convergence constitutes an admission that some frequency interpretation is required.

However, in practical inference situations it is often not possible to take as many observations from the same distribution as one would like. The prior distribution is important when there are few observations, and care must be taken to make it conform to the frequency interpretation of this note. This interpretation, referring to a sequence of disparate experiments, is unfortunately not constructive. One can say, however, that the routine use of vague priors should certainly be avoided.

Essentially, the criteria of this note are already largely met in the economic life where subjects act with full responsibility for the results of their actions. Businessmen would hardly alter their behaviour and insurance actuaries would hardly alter their formulas after reading this note, although they might benefit from a deeper understanding of what they are doing and why. It is different with social science research work, where university statisticians are asked to supply methods. The decisions and actions, if any, resulting from a typical sociological investigation are not at all clear, and anyhow the sociologists who made the investigation will certainly not suffer from any mistaken actions. This clearly explains why so many social and medical science investigations are so bad in statistical methodology, in particular why they are not Bayesian. Naturally even Bayesian statistics can be bad statistics, and in fact often will be when the inferrer is separate from the decision-maker. But at least the unavoidable subjective element is clearly exposed rather than hidden in presumed objective clothing, and others are free to criticize the priors as unreasonable and at variance with a reasonable frequency interpretation.

I would say the following is a rather typical situation. Let decision situation i consist of inference concerning a parameter $w_i \in R$ and action on the basis of the inference. It could be the energy-saving effect of thermostat-regulated radiators or the life-saving effect of a specific type of heart operation. Assume that $w_i > 0$ means a beneficial effect, $w_i = 0$ means no effect and $w_i < 0$ means a harmful effect. The decision d_i is either to introduce or not to introduce the method in (construction respectively medical) practice. The reason that the method is statistically investigated is that a pilot study has indicated that it could have a good effect. However, the statistician wants to appear objective and chooses therefore a non-informative prior distribution for w_i which is symmetric around the origin. It is only possible to take rather few observations on some random variable with density $f_i(\cdot \mid w_i)$. The non-informative prior means that the investigation will give about the same result as a standard sampling theory investigation. Typically the posterior distribution P_i for w_i will have $E(w_i)$ somewhat greater than zero and a considerable dispersion. The pilot study would have justified a prior with expectation greater than zero and a not too large dispersion. The posterior distribution would then have had larger mean and smaller dispersion than obtained under the non-informative prior. With A_i for example a sequence of intervals around the origin, $\limsup_n |\phi_n|$ would have been smaller under the informative than under the non-informative priors. A better regard to the pilot studies would have given a better frequency agreement and a larger sum of utilities over many investigations. In my view the statistical profession has here a responsibility for enlightening the users of statistics, so that the statistician no longer would have the need to appear objective by using standard sampling theory or non-informative priors when in fact there exists prior information pointing in some direction.

REFERENCES

Breiman, L.: 1968, *Probability*, Addison-Wesley, Reading.
De Finetti, B.: 1974, *Theory of Probability*, Vol. 1, Wiley, New York.
DeGroot, M.: 1970, *Optimal Statistical Decisions*, McGraw-Hill, New York.
Leinfellner, W.: 1984, 'Evolutionary Causality, Theory of Games and Evolution of Intelligence', in F. M. Wuketits (ed.), *Concepts and Approaches in Evolutionary Epistemology*, D. Reidel, Dordrecht, 233–277.
Lindley, D. V.: 1975, 'The Future of Statistics – A Bayesian 21st Century', *Adv. Appl. Prob.* (Suppl), 105–115.
Lindley, D. V.: 1978, 'The Bayesian Approach', *Scand. J. Statistics* **5**, 1–26.
Reichenbach, H.: 1949, *The Theory of Probability*, University of California Press, Berkeley.
Schuster, P., Sigmund, K. *et al.*: 1981, 'Self Regulation of Behaviour', *Biological Cybernetics* **40**, 1–8.
Smith, M.: 1982, *Evolution and the Theory of Games*, Cambridge University Press, Cambridge.

Länsförsäkringsobolagen,
Box 26057,
S-100 41 Stockholm,
Sweden