COWLES COMMISSION DISCUSSION PAPER: STATISTICS NO. 367

Rational Selection of a Decision Function in an Observational Situation

Roy Radner

October 23, 1951

Outline

1. Introduction.

Some statistical decision procedures which have been proposed in recent years, (mainly the minimax and allied rules), have been found unsatisfactory in certain cases because of their lack of responsiveness to relevant observations. My aim is to formalize our intuition behind this dissatisfaction in the form of one or more axioms about the rational selection of decision functions. Before this can be done it will be necessary to review several ideas which have already been set forth by various authors, concerning other axioms of rationality and their consequences.

The second section presents those definitions and axioms which apply to the general decision problem as represented by the normal form of the zero-sum, two person game. In order to discuss the response of a decision to observations, this very general representation will have to be particularized to include the concept of observations. I shall use the model of a truncated sequential decision problem (which is really more special than is necessary). The third section defines such a model and describes the structure of Bayes decisions in terms of "a posteriori" probability measures. The fourth section gives certain definitions and axioms about the relation between good decisions and observations, and some consequences.

These last mentioned axioms will not, in general, determine a unique optimal decision, in the sense that the minimax rule does, but, like the principle of admissibility, they effect a certain reduction in the class of decisions to be considered.

2. The General Decision Problem.

A decision problem Q is defined by:

1) A set Ω of states of nature ω

2) A set $\mathcal{D}$ of decisions δ

3) A (payoff) function K from $\Omega \times \mathcal{D}$ to the real numbers.

$K(\omega, \delta)$ may be interpreted as the expected utility to the decision maker associated with the decision δ when ω is the true state of nature.

A <u>solution of a problem</u> Q is a subset of $\mathcal{D}$. A <u>solution of a class</u> $\mathcal{Q}$ <u>of problems</u> Q is a function assigning a problem solution to each Q in $\mathcal{Q}$. A solution (of either type) is <u>rational</u> if it satisfies the several axioms which will be enumerated in the course of this paper. Not every decision problem will have a rational solution in this sense.

Two decisions δ_1 and δ_2 are <u>equivalent</u> if for every ω in Ω, $K(\omega, \delta_1) = K(\omega, \delta_2)$. δ_1 <u>is uniformly better than</u> δ_2 if for every ω in Ω, $K(\omega, \delta_1) \geq K(\omega, \delta_2)$ and for some ω in Ω, $K(\omega, \delta_1) > K(\omega, \delta_2)$.

<u>Axiom 1.</u> A solution shall be nonempty.

<u>Axiom 2.</u> If δ_1 is equivalent to δ_2, and δ_1 is in the solution, then so is δ_2.

<u>Axiom 3.</u> If δ is in the solution then there is no δ' in $\mathcal{D}$ which is uniformly better than δ.

Axiom 2 justifies our assuming that each equivalence class in $\mathcal{D}$ has only one representative.

Suppose that Ω has a σ-ring of subsets such that we can speak of probability measures (pr.m.) on it. On the other hand it will be convenient to assume that $\mathcal{D}$ is itself the set of all probability measures on some set of "pure" decisions. If G is a pr.m. on Ω, then a decision δ' is Bayes relative to G if it maximizes (over $\mathcal{D}$) the expectation $\int_{\Omega} K(\omega, \delta) dG$. G is also said to be the effective distribution for δ'.

A subset $\mathcal{C}$ of $\mathcal{D}$ is complete in $\mathcal{D}$ if for every δ not in $\mathcal{C}$ there exists a δ' in $\mathcal{C}$ such that δ' is uniformly better than δ. Wald has shown that under certain general conditions the class of all Bayes decisions is complete in $\mathcal{D}$. Under these conditions any decision in a rational solution has an effective distribution (by Axiom 3). This fact is useful because certain properties of a decision can be inferred from its effective distribution.

Two problems $Q = (\Omega, \mathcal{D}, K)$ and $Q' = (\Omega', \mathcal{D}', K')$ are δ-equivalent if both $\Omega = \Omega'$ and there is a one-one function ϕ from $\mathcal{D}$ onto $\mathcal{D}'$ such that for every δ in $\mathcal{D}$ and ω in Ω,

$$K(\omega, \delta) = K'[\omega, \phi(\delta)].$$

Axiom 4. If Q and Q' are in a class $\mathcal{Q}$, $\mathcal{S}$ is a solution of $\mathcal{Q}$ and Q is δ-equivalent to Q', then δ is in $\mathcal{S}(Q)$ if and only if the corresponding δ' is in $\mathcal{S}(Q')$.

3. The Truncated Sequential Decision Problem.

Ω is the set of all states of nature ω considered possible (although only one of them can be the true state).

$\xi = (x_1, \ldots, x_T)$ is a finite sequence of observations whose joint distribution is determined for every ω in Ω and is either discrete or absolutely continuous.

D is a set of terminal decisions d.

δ_n is a function of the first n observations which takes on either values in D, or the value e, which is defined to be the decision to take another observation, or a random mixture of these.

$\mathcal{D}$ is the (convex) set of all decision functions $\delta = (\delta_0, \delta_1, \ldots, \delta_T)$.

$k(\omega, d)$ is the "payoff," given ω and the terminal decision d.

c_n is the cost of making the n'th observation.

$K(\omega, \delta)$ is the expected value of $[K(\omega, d)$ minus the total cost of observation], given ω and the decision function δ.

G_o denotes an "a priori" pr.m. on Ω, and

$G(\cdot \mid x_1, \ldots, x_n)$ denotes the corresponding "a posteriori" pr.m. on Ω, given $x_1, \ldots, x_n$, as in Bayes theorem; i.e., if $f_\omega(x_1, \ldots, x_n)$ is the joint elementary probability law for $x_1, \ldots, x_n$ given ω, then for any measurable set $W \subset \Omega$:

$$G(W \mid x_1, \ldots, x_n = \frac{\int_W f_\omega(x_1, \ldots, x_n) dG_o}{\int_\Omega f_\omega(x_1, \ldots, x_n) dG_o} .$$

Theorem. The following procedure is Bayes relative to the a priori pr.m. G_o:

At stage zero proceed according to any decision function which is Bayes relative to G_o. If observation x_n has been taken (n=1, ..., T) choose that decision function $(\delta_n, \ldots, \delta_T)$ of $(x_n, \ldots, x_T)$ which is Bayes relative to the a posteriori pr.m. $G(\cdot \mid x_1, \ldots, x_n)$.

4. Rational Decisions and Observations.

An observation x_n is defined to be relevant given $x_1, \ldots, x_{n-1}$ if there exist ω' and ω'' such that the distribution of x_n given ω' and $x_1, \ldots, x_{n-1}$ differs from the distribution of x_n given ω'' and $x_1, \ldots, x_{n-1}$. We should like a decision procedure to have the property that if Q is a given problem and x_n is a relevant observation, then provided observation costs are small enough and the penalty for making wrong decisions is high enough, the decision procedure will advise taking the observation x_n. This can be formalized as follows:

For any Q, let $\mathcal{Q}_1(Q)$ be the class of all problems which are identical with Q except possibly for the costs of observation c_i and the terminal pay-off function $k(\omega, d)$ (and, of course, all things defined in terms of these), provided that all the c_i are positive.

Axiom 5. If $\mathcal{Q}$ contains, along with every Q in it, $\mathcal{Q}_1(Q)$, $\mathcal{S}$ is a solution of $\mathcal{Q}$, and x_n is a relevant observation given $x_1, \ldots, x_{n-1}$, then there exists a Q' in $\mathcal{Q}$ such that for every δ in $\mathcal{S}(Q')$, $\delta_0(\xi) = \ldots = \delta_n(\xi) = e$ for all ξ with the given first n-1 coordinates.

For every ω in Ω define the point pr.m. $G^{(\omega)}$ on Ω by:

$$G^{(\omega)}(W) = \begin{cases} 1, & \text{if } \omega \text{ in } W \\ 0, & \text{if } \omega \text{ not in } W \end{cases} .$$

Let X be an infinite sequence of observations whose joint distribution is determined for every ω in Ω, and let G_n denote a function from the space of partial sequences $x_1, \ldots, x_n$ to the space of pr.m. on Ω. X is defined to be conclusive if there exists a infinite sequence G_n such that for any fixed ω in Ω, $W \subset \Omega$ and $\varepsilon > 0$:

$$\lim_{n \to \infty} P_r\left\{ |G_n(W) - G^{(\omega)}(W)| > \varepsilon \right\} = 0 .$$

In the presence of a conclusive sequence we should formulate a stronger requirement. For any Q let $\mathcal{Q}_2(Q)$ be the class of all problems identical with Q except for the number T of possible observations and the costs of observation c_i (provided these last are all positive) and such that the observations $x_1, \ldots, x_T$ are the first T elements of an infinite X.

Axiom 6. If $\mathcal{Q}$ contains, along with every Q in it, $\mathcal{Q}_2(Q)$, $\mathcal{S}$ is a solution of $\mathcal{Q}$ and X is conclusive, then for any integer N there exists a problem Q in $\mathcal{Q}$ such that for every δ in $\mathcal{S}(Q)$, $\delta_0(\xi) = \ldots = \delta_N(\xi) = e$ for all ξ.

Roughly speaking, Axiom 6 says that if Q is a problem with $x_1, \ldots, x_T$ part of a conclusive X, then a good decision rule will advise taking an arbitrarily large number of observations provided the costs of observation can be made small enough and it is possible to take enough observations.

We can relate the theorem of Section 3 to these axioms through the following concepts: If Ω is denumerable then a pr.m. G on Ω is <u>prejudiced</u> if there exists ω in Ω such that $G(\omega) = 0$. If Ω is a subset of Euclidean space with nonzero Lebesgue measure then a pr.m. G on Ω is <u>prejudiced</u> if there is some $W \subset \Omega$ with nonzero Lebesgue measure such that $G(W) = 0$.

If $\mathcal{S}$ is a solution of $\mathcal{Q}$ such that for every Q in $\mathcal{Q}$, $\mathcal{S}$ (Q) contains a decision which is Bayes relative to some prejudiced pr.m. on Ω, then $\mathcal{S}$ cannot satisfy Axiom 6, by the theorem of Section 2.

In examples in which minimax or related decision procedures lead to decisions which are unresponsive to further observations, no matter how cheap, it has been found that the effective distributions are prejudiced.

Axioms 5 and 6 do not by any means capture all desiderata with respect to the relationship between decisions and observations; e.g., consider the dissatisfaction with the minimax estimate of the binomial parameter, when the loss is the squared error. To me, an indication of the difficulty is the fact that the effective pr.m. on Ω, which is the closed interval (0, 1), has a density which is zero at 0 and 1. This is a weaker kind of "prejudice" than that defined above.

5. Further Questions.

It would be interesting to see what implications these axioms have for both the "generalized Bayes-Minimax" approach of Hurwicz and the "restricted Bayes" approach of Hodges and Lehmann. In the former case, the types of classes of "a priori" distributions one would be allowed to start with would be definitely restricted.

Another pertinent problem is that of finding the most general conditions under which the sequences of a posteriori probability measures converge in the sense of the sequences G_n in Section 4.