A NOTE ON THE APPLICABILITY OF SEQUENCE COMPOUND DECISION SCHEMES*

by

Stephen B. Vardeman
Purdue University

Department of Statistics
Division of Mathematical Sciences
Mimeograph Series #506

August 1977

*This work was supported through a Purdue Research Foundation Faculty XL Grant.

A NOTE ON THE APPLICABILITY OF SEQUENCE COMPOUND DECISION SCHEMES*

by

Stephen B. Vardeman
Purdue University

Abstract - We present an example which emphasises a limitation on the practical applicability of some well known asymptotically Bayes sequential compound decision rules.

I. INTRODUCTION

We consider a (component) decision problem with parameters $\theta \in \Theta$ indexing distributions P_θ on a sample space $\mathscr{X}$, possible actions $a \in \mathcal{A}$, loss function $L(\theta,a)$ and decision rules $d(\cdot)$ which are functions from $\mathscr{X}$ into $\mathcal{A}$. When faced with a sequence of such decision problems many authors, among them Hannan [4] and [5], Van Ryzin [7] and Ballard, Gilliland and Hannan [1] (see also the closely related work of Gilliland [3], and Cover and Shenhar [2]) have adopted a compound viewpoint, taken

$$N^{-1} \sum_{i=1}^{N} L(\theta_i, a_i)$$

as the N problem compound loss suffered and considered decision rules $\underline{\delta} = (\delta_1, \delta_2, \ldots)$ which are sequences of functions, $\delta_i(\cdot)$ mapping the first i

*This work was supported through a Purdue Research Foundation Faculty XL Grant.

observations $\underline{X}_i = (X_1,\ldots,X_i)$ to an action a_i to be taken in the ith problem. For a sequence of parameters $\underline{\theta} = (\theta_1,\theta_2,\ldots)$ (usually taken as fixed but unknown) $\underline{X}_i$ is supposed to have distribution $P_{\underline{\theta}_i} = P_{\theta_1} \times P_{\theta_2} \times\ldots\times P_{\theta_i}$ and the average risk of a sequence compound decision rule $\underline{\delta}$ through the first N problems becomes

$$R_N(\underline{\theta},\underline{\delta}) = N^{-1} \sum_{i=1}^{N} \int L(\theta_i,\delta_i(\underline{x}_i))dP_{\underline{\theta}_i}(\underline{x}_i).$$

In particular, the N problem risk of a "simple" rule $\underline{d}$ (a compound rule for which $\delta_i(\underline{X}_i) = d(X_i)$ for a fixed component rule d) is

$$R_N(\underline{\theta},\underline{d}) = N^{-1} \sum_{i=1}^{N} \int L(\theta_i,d(x_i))dP_{\theta_i}(x_i).$$

For many types of component problem it is possible to produce $\underline{\delta}$ with the optimality property that

$$\overline{\lim}\ (R_N(\underline{\theta},\underline{\delta}) - \min_d R_N(\underline{\theta},\underline{d})) \leq 0$$

for every $\underline{\theta}$. That is regardless of the value of the unknown $\underline{\theta}$, in the limit the risk of $\underline{\delta}$ is at least as good as could have been obtained if $\underline{\theta}$ were known and a best simple rule used. Indeed, Van Ryzin [7] for the mxn component was able to find $\underline{\delta}$ such that there exists a constant c not depending on $\underline{\theta}$ for which

(1) $$R_N(\underline{\theta},\underline{\delta}) - \min_d R_N(\underline{\theta},\underline{d}) \leq cN^{-\frac{1}{2}}.$$

The fact that bounds which like (1) are uniform in $\underline{\theta}$ can be obtained, makes it tempting to expect asymptotically optimal sequence compound procedures to perform well in any situation where a sequence of structurally identical decisions are to be made based on P_{θ_i} distributed X_i, regardless of whether $\underline{\theta}$

is thought of as fixed or randomly chosen. It is true that provided the conditional distributions of $\underline{X}_i$ given $\underline{\theta}_i$ are of the product form $P_{\underline{\theta}_i}$, an inequality like (1) guarantees good behavior of both the conditional on $\underline{\theta}$ compound risk $R_N(\underline{\theta},\underline{\delta})$ and its average with respect tc any distribution on $\underline{\theta}$. We present an example however, which emphasises the fact that assuming each conditional distribution of X_i given θ_i to be P_{θ_i} is not enough to guarantee good asymptotic performance of an optimal sequence compound decision rule.

II. A SEQUENCE OF DISCRIMINATION PROBLEMS

We specialize our discussion to a situation of discriminations between two Bernoulli distributions. That is, suppose $\Theta = \mathcal{A} = \{1,2\}$, $\mathcal{X} = \{0,1\}$, P_1 is the Bernoulli distribution with $p = 1/3$, P_2 is the Bernoulli distribution with $p = 2/3$ and $L(\theta,a) = I[\theta \neq a]$. This 2x2 component is covered by the work of Van Ryzin [7] and his Theorem 4.1 guarantees that the "play Bayes against the estimated past empiric distribution" procedure $\underline{\delta}^*$ defined for $i > 1$ by

$$\delta_i^*(\underline{X}_i) = \begin{cases} 1 & \text{if } X_i \leq 5-9\bar{X}_{i-1} \\ 2 & \text{if } X_i > 5-9\bar{X}_{i-1} \end{cases}$$

has the property (1) (i.e. provided the conditional distribution of $\underline{X}_i$ given $\underline{\theta}_i$ is $P_{\underline{\theta}_i}$, $\underline{\delta}^*$ is asymptotically at least as good as a best simple rule). Notice that for $\bar{X}_{i-1} \leq 4/9$ $\delta_i^* = 1$, for $\bar{X}_{i-1} > 5/9$ $\delta_i^* = 2$, and for $\bar{X}_{i-1} \in (4/9, 5/9]$ $\delta_i^* = 1$ when $X_i = 0$ and $\delta_i^* = 2$ when $X_i = 1$.

We can construct a distribution for $(\underline{\theta},\underline{X})$ with $X_i|\theta_i$ distributed as P_{θ_i} under which δ^* does very poorly in terms of expected compound loss. Take $\theta_1 \equiv 1$, the conditional distribution of X_i given $\underline{\theta}_i$ and $\underline{X}_{i-1}$ to be P_{θ_i}, and the conditional distribution of θ_i given $\underline{\theta}_{i-1}$ and $\underline{X}_{i-1}$ to be degenerate at 2 if $\bar{X}_{i-1} \leq 5/9$ and at 1 if $\bar{X}_{i-1} > 5/9$. For $i > 1$ the component risk suffered at the ith discrimination

(2) $$E\, L(\theta_i, \delta_i(\underline{X}_i))$$

can be written as

$$E\, L(\theta_i,\delta_i(\underline{X}_i))\{I[\bar{X}_{i-1} \leq 4/9] + I[\bar{X}_{i-1} \in (4/9,5/9]]+ I[\bar{X}_{i-1} > 5/9]\}.$$

But by the way we constructed the joint distribution of $\underline{\theta}$ and $\underline{X}$ and the definition of $\underline{\delta}^*$, for $i > 1$

$$\begin{aligned}(2) &= EI[\bar{X}_{i-1} \leq 4/9] + 1/3\, EI[\bar{X}_{i-1} \in (4/9,5/9]]+EI[\bar{X}_{i-1} > 5/9],\\ &\geq 1/3\, EI[\bar{X}_{i-1} \leq 5/9] + EI[\bar{X}_{i-1} > 5/9],\\ &= 1/3 + 2/3\, EI[\bar{X}_{i-1} > 5/9],\\ &= 1/3 + 2/3\, EI[\theta_i = 1].\end{aligned}$$

Hence the expected N problem compound loss of $\underline{\delta}^*$

$$E\, N^{-1} \sum_{i=1}^{N} L(\theta_i,\delta_i^*(\underline{X}_i))$$

is bounded below by

(3) $$\frac{N-1}{3N} + 2/3\, E\, N^{-1} \sum_{i=2}^{N} I[\theta_i=1].$$

In what follows we will show that the second term in (3) converges to 2/9 so that the compound risk of $\underline{\delta}^*$ is asymptotically at least 5/9. But a simple rule that in each problem uses the component minimax rule $d(x) = x+1$ has compound risk 1/3 for any distribution for $(\underline{\theta},\underline{X})$ for which $X_i|\theta_i$ has distribution P_{θ_i}. The asymptotic risk of $\underline{\delta}^*$ is thus nowhere near that of a best simple rule when the expectation is calculated according to the constructed distribution rather than according to one for which the conditional distribution of $\underline{X}_i$ given $\underline{\theta}_i$ is of a product form $P_{\underline{\theta}_i}$. (Indeed, in the present situation the risk of $\underline{\delta}$* is closer to 2/3, which is the asymptotic value of the risk of a worst simple rule.)

To verify the claim that $EN^{-1}\sum_{i=2}^{N} I[\theta_i = 1]$ converges to 1/3, notice first that

$$N^{-1}\sum_{i=2}^{N} I[\theta_i = 1] = f_N - N^{-1}$$

where f_N is the fraction of θ_i's for i = 1,...,N such that $\theta_i = 1$. Then observe that

$$E\bar{X}_N = N^{-1}E(1/3 + \sum_{i=2}^{N} E[X_i|\underline{X}_{i-1}]),$$

$$= N^{-1}E(1/3 + \sum_{i=2}^{N} (1/3\ I[\bar{X}_{i-1} > 5/9] + 2/3\ I[\bar{X}_{i-1} \leq 5/9])),$$

$$= E[1/3\ f_N + 2/3(1-f_N)],$$

$$= 2/3 - 1/3\ Ef_N.$$

Thus it will be sufficient to show that $E\bar{X}_N$ converges to 5/9, and for this it in turn suffices to show that $E(\bar{X}_N - 5/9)^2$ converges to 0. Now with $S_N = \sum_{i=1}^{N} X_i$

$$E(S_{N+1} - 5/9(N+1))^2 = E(S_N - 5/9N)^2 + E(X_{N+1} - 5/9)^2$$
$$+ 2E(S_N - 5/9N)(X_{N+1} - 5/9).$$

The last term above is nonpositive since

$$E(S_N - 5/9N)(X_{N+1} - 5/9) = E(S_N - 5/9N)E[X_{N+1} - 5/9|S_N]$$

and by construction $(S_N - 5/9N)$ and $E[X_{N+1} - 5/9|S_N]$ have opposite signs. Thus dropping the last term

$$E(S_{N+1} - 5/9(N+1))^2 \leq E(S_N - 5/9N)^2 + E(X_{N+1} - 5/9)^2,$$
$$\leq E(S_N - 5/9N)^2 + 1.$$

Since $E(X_1 - 5/9)^2 < 1$, $E(S_{N+1} - 5/9(N+1))^2 \leq N+1$ so that $E(\bar{X}_{N+1} - 5/9)^2 \leq \frac{1}{N+1}$ and the claim has been established.

III. REMARKS

In a sense, the example is unfair to the sequence compound methodology in that it points out possible bad performance of an optimal rule only when used in a situation other than one for which it was designed. But the example is important, as the temptation to lose sight of the distributional assumptions implied in a statement of compound optimality is a strong one. Correct statements like "nature can not set up the predictor for future disasterous performance" (see Cover and Shenhar [2]) are frequent in the compound decision literature and are capable of misinterpretation. When a real situation is adequately modeled by a statement that $\underline{\theta}$ is fixed but unknown, an asymptotically optimal sequence compound rule is appropriate. On the other hand, such a rule may not be appropriate when the real situation is more complicated. For instance the example shows that if the parameters are being generated by a well informed antagonist (i.e. one that knows $\underline{X}_{i-1}$ when generating θ_i) use of an "optimal" rule can be disasterous.

It should be remarked that while we have spoken in terms of sequentially faced decisions, similar cautions apply to the application of the asymptotically optimal nonsequential compound rules of Robbins [6] and others.

ACKNOWLEDGEMENT

The author is grateful to Professor Burgess Davis for suggesting the proof of the convergence of $EN^{-1} \sum_{i=1}^{N} I[\theta_i = 1]$ to 1/3.

REFERENCES

[1] Robert Ballard, Dennis Gilliland and James Hannan, "$O(N^{-\frac{1}{2}})$ convergence to k-extended Bayes risk in the sequence compound decision problem with mxn component," RM-333, Statistics and Probability, Michigan State University.

[2] Thomas Cover and Aaron Shenhar, "Compound Bayes predictors for sequences with apparent Markov structure," IEEE Trans. Syst., Man, Cybern., Vol. SMC-7, pp. 421-424, June 1977.

[3] Dennis, Gilliland, "Asymptotic risk stability resulting from play against the past in a sequence of decision problems," IEEE Trans. Information Theory, Vol. IT-18, pp. 614-617, Sept. 1972.

[4] James, Hannan, "The dynamic statistical decision problem when the component problem involves a finite number, m, of distributions," (Abstract), Ann. Math. Statist., Vol. 27, pp. 212, March 1956.

[5] James Hannan, "Approximation to Bayes risk in repeated play" Contributions to the Theory of Games, Vol. III, Annals of Mathematics Studies No. 39, pp. 97-139, Princeton 1957.

[6] Herbert Robbins, "Asymptotically subminimax solutions of compound statistical decision problems," Proc. Second Berkeley Symposium on Mathematical Statistics and Probability, University of California Press, pp. 131-148, 1951.

[7] John Van Ryzin, "The sequential compound decision problem with mxn finite loss matrix," Ann. Math. Statist., Vol. 37, pp. 954-975, August 1966.