A MODIFIED BAYES STOPPING RULE[1]

By Sigmund J. Amster[2]

University of North Carolina

1. Introduction of the modified Bayes rule. This paper describes a stopping rule for sequential sampling that weighs the cost of additional sampling against the expected gain to be derived from additional observations. The modified Bayes rule (MBR) requires one more observation to be taken as long as the posterior risk is larger than the expected posterior risk for some additional fixed-size sample. In this investigation, risk is defined within the Wald framework of statistical decision theory [6], using losses and costs.

A subjective probability justification for the MBR may be found in interpreting the value of ξ_0, the prior distribution on the parameter θ, at any set of θ values, as representing the original relative conviction that the true value of θ lies in this set. Once the first n observations have been taken, the belief has been changed as reflected in the values of the posterior distribution ξ_n. In either case, the distribution ξ determines *which* risk one would like to minimize. The MBR will accept R_n, the present average risk, (i.e., stop sampling) only if the cost of increasing one's convictions, through knowledge of a sample of any *fixed* size, is more than the expected amount to be gained. If not, one more sample will be taken and the same problem posed with the (hoped-for) better knowledge of the true state of nature.

In the context treated above, the defining property of the Bayes sequential rule (BSR), that of minimizing the *original* average risk, does not seem particularly relevant. However, the method which determines this rule, by comparing R_n with the average risk "resulting from a continuation if at each future stage we did the *best* we could with the resulting observations" ([1], page 243), is really the optimal property. The MBR tries to approximate this by considering only the average risk if any fixed-size future sample were taken.

The calculation of the MBR is feasible whenever the Bayes fixed sample procedure can be explicitly obtained. (Part III of [5] can offer assistance in the evaluation of the integrals involved.) However, except for certain cases, as testing two simple hypotheses, when Wald's sequential probability ratio test (SPRT) is such a rule, or if the BSR is truncated or fixed sample, it is not usually possible to carry out the Bayes procedure. Even in these cases, the determination of the appropriate SPRT is not simple and in the truncated cases the necessary computations are exceedingly tedious.

Received July 12, 1962; revised June 17, 1963.

[1] Submitted in partial fulfillment of the requirements for the Ph.D. degree, University of North Carolina. This research was partially supported by the Office of Naval Research under contract No. Nonr-855(09).

[2] Now at Bell Telephone Laboratories, Whippany, New Jersey.