THE EMPIRICAL BAYES APPROACH TO STATISTICAL DECISION PROBLEMS[1]

By Herbert Robbins

Columbia University

1. Introduction. The empirical Bayes approach to statistical decision problems is applicable when the same decision problem presents itself repeatedly and independently with a fixed but unknown *a priori* distribution of the parameter. Not all decision problems in practice come to us imbedded in such a sequence, but when they do the empirical Bayes approach offers certain advantages over any approach which ignores the fact that the parameter is itself a random variable, as well as over any approach which assumes a personal or a conventional distribution of the parameter not subject to change with experience. My own interest in the empirical Bayes approach was renewed by recent work of E. Samuel [10], [11] and J. Neyman [6], to both of whom I am very much indebted. In keeping with the purpose of the Rietz Lecture I shall not confine myself to presenting new results and shall try to make the argument explicit at the risk of being tedious. In the current controversy between the Bayesian school and their opponents it is obvious that any theory of statistical inference will find itself in and out of fashion as the winds of doctrine blow. Here, then, are some remarks and references for further reading which I hope will interest my audience in thinking the matter through for themselves. Considerations of space have confined mention of the non-parametric case, and of the closely related "compound" approach in which no *a priori* distribution of the parameter is assumed, to the references at the end of the article.

2. The empirical Bayes decision problem. We begin by stating the kind of statistical decision problem with which we shall be concerned. This comprises

(a) A *parameter space* Λ with generic element λ. λ is the "state of nature" which is unknown to us.

(b) An *action space* A with generic element a.

(c) A *loss function* $L(a, \lambda) \geqq 0$ representing the loss we incur in taking action a when the parameter is λ.

(d) An *a priori distribution* G of λ on Λ. G may or may not be known to us.

(e) An *observable random variable* x belonging to a space X on which a σ-finite measure μ is defined. When the parameter is λ, x has a specified probability density f_λ with respect to μ.

The problem is to choose a *decision function* t, defined on X and with values in

Received 22 October 1963.

[1] Presented as the Rietz Lecture at the 26th Annual Meeting of the Institute of Mathematical Statistics, August 29, 1963, at Ottawa, Canada. Research sponsored by National Science Foundation Grant NSF-GP-316 at Columbia University.