

INVITED PAPER

HERBERT ROBBINS AND SEQUENTIAL ANALYSIS¹

BY DAVID SIEGMUND

Stanford University

This paper reviews Herbert Robbins' research in sequential analysis (excluding stochastic approximation) from 1952 until roughly 1980. Its relation to the research of his contemporaries and its impact on subsequent research are described.

1. Introduction. The period immediately following the Second World War saw a rapid growth in research across all scientific fields. In statistics, one of the most imaginative contributors to this growth was Herbert E. Robbins.

Herbert Robbins studied mathematics at Harvard in the 1930s. His initial research in probability and statistics began during the war, but it was his appointment to the faculty of the new Department of Mathematical Statistics at the University of North Carolina that laid the foundation for a career in statistics. His first paper on sequential analysis [Robbins (1952)] was published shortly after his two seminal papers on empirical Bayes [Robbins (1951)] and stochastic approximation [Robbins and Monro (1951)], which are discussed elsewhere in this issue. While the subject of sequential analysis can be traced back to the work of Dodge and Romig (1929) and Shewhart (1931), who were interested in problems of quality control, the modern history of the subject begins with the research of Wald conducted during World War II [cf. Wald (1947)], the independent, albeit less systematic contributions of Barnard (1946) and Anscombe (1946) and the research of Stein (1945) on fixed-width confidence intervals for a normal mean. Impetus to additional development immediately following the war came from the research of Wald and Wolfowitz (1948) and Arrow, Blackwell and Girshick (1949) on the optimality of the sequential probability ratio test.

The purpose of this article is to review Robbins' research in sequential analysis (excluding stochastic approximation), which during the 1960s and part of the 1970s was the main focus of his scientific activity. It is convenient to divide the review thematically into four sections: (i) sequential allocation, (ii) optimal stopping theory, (iii) sequential estimation and (iv) sequential hypothesis testing, although some research cuts across more than one category. A general feature of Robbins' research was its originality and elegance, which stimulated others to

Received December 2001; revised February 2002.

¹Supported by NSF Grant DMS-00-72523.

AMS 2000 subject classification. 62L10.

Key words and phrases. Sequential analysis, optimal stopping, dynamic programming.

think about related problems. I have tried to put his contributions to sequential analysis into historical perspective by discussing antecedent and subsequent related research, but I have not attempted a systematic review. To compensate for this selectivity, I have tried to indicate review articles that contain extensive bibliographies. See especially Lai (2001), which provides a broad review of research in sequential analysis (including stochastic approximation).

2. Sequential allocation. Robbins (1952) initiated the statistical discussion of sequential allocation with the following simple problem. An experimenter will receive a sequence of rewards X_1, X_2, \dots by choosing at each time $n = 1, 2, \dots$ to observe a random variable with distribution function F and mean value μ or a random variable with distribution function G and mean value $\nu \neq \mu$. The selection of a distribution to generate X_n can be based on the previous rewards, X_1, \dots, X_{n-1} , but once a distribution is selected X_n is generated independently of the previous X_i . Several different problems and procedures are discussed to achieve the goal of maximizing in an appropriate sense the expected average reward

$$(1) \quad E[(X_1 + \dots + X_n)/n].$$

One result is to show that if one chooses two sparse but infinite sequences of integers, makes forced choices of F at the times indicated by one sequence, forced choices of G at the times of the other sequence and at all other times chooses the distribution that in the past has given the larger average reward, then in the limit as $n \rightarrow \infty$ (1) converges to $\max(\mu, \nu)$. This result illustrates what has come to be recognized as a general issue in problems of sequential decision making under uncertainty: the conflict between maximizing one's immediate expected reward and gathering information that will be useful in the long run. This problem and a long line of successors are now referred to as "bandit problems," since they can be thought of as involving sequential strategies for playing slot machines ("one-armed bandits") having unknown statistical properties.

Robbins (1956) returned to this allocation problem by imposing the additional constraint that the experimenter was limited by a finite memory of the past. This led to a fascinating sequence of papers culminating in the research of Cover and Hellman [cf. Cover (1968) and Hellman and Cover (1970) and the references cited therein]. He returned to it again in Flehinger, Louis, Robbins and Singer (1972) and Robbins and Siegmund (1974a), but here the motivation and formulation were somewhat different. The sequential allocation was to be a choice between two treatments in a sequential clinical trial with immediate responses. The inferior treatment was to be allocated as infrequently as possible subject to the constraint that the clinical trial was to end at some finite time with a test of hypothesis concerning the better treatment. The same conflict between, in this case, allocating the treatment that appeared to be better and gathering information about the

relative merits of the two treatments manifested itself. The key to the analysis in these papers is a “separation” theorem, which, for the case that responses to the two treatments are normally distributed with possibly different means and a common variance, allows one to establish a class of valid tests, having essentially a fixed power function determined by the stopping rule, and then to study the allocation problem within this class of tests. For subsequent developments along the same lines, see Louis (1975) and Hayre (1979). A review and discussion of the gap between theory and practice is found in Siegmund (1985). For a somewhat different approach, references up to 1985 and another discussion of the gap between theory and practice, see Bather (1985), Armitage (1985) and the accompanying discussion. An application where the experimental treatment seemed so beneficial and the adaptive allocation so complete that very few patients were assigned to control, thus leading to some controversy about the validity of the conclusion, is described by Ware (1989).

A different line of development was the attempt to use the approach of dynamic programming to give an exact solution for a Bayesian formulation of Robbins’ original allocation problem [cf. Bradt, Johnson and Karlin (1956) and Bellman (1956)]. A breakthrough was provided by Gittins and Jones (1974), who showed that a discounted version of the allocation problem with product prior distributions could be reduced to a parameterized family of *optimal stopping problems*, which do not involve relations between the different “arms” and are comparatively easy to solve. See Gittins (1989) for a complete discussion.

Robbins (1952) also discussed what he called the problem of *optional stopping*, which is described below under sequential hypothesis testing.

3. Optimal stopping theory. Optimal stopping theory has its roots in the study of the optimality properties of the sequential probability ratio test of Wald and Wolfowitz (1948) and Arrow, Blackwell and Girshick (1949). The essential idea in both of these papers was to create a formal Bayes problem, the solution to which could be regarded either as an end in its own right or as a *deus ex machina* on the way to solving a non-Bayesian problem.

(The non-Bayesian formulation of the optimality property must certainly be one of the most surprising and profound results of mathematical statistics. It says that, for a sequence of independent identically distributed observations, given a sequential probability ratio test of a simple hypothesis against a simple alternative, any other test with error probabilities no larger than those of the sequential probability ratio test must have expected sample sizes at least as large as the sequential probability ratio test under *both* hypotheses.)

The formal Bayes problem is what we would now call an optimal stopping problem. A decision maker observes an adapted sequence $\{R_n, \mathcal{F}_n, n \geq 1\}$, with $E|R_n| < \infty$ for all n . At each time n a choice is to be made, to stop sampling and collect the currently available reward, R_n , or continue sampling in the *expectation* of collecting a larger reward in the future. An optimal stopping rule N is one that

maximizes the expected reward, $E(R_N)$. The key to finding an optimal or close to optimal stopping rule is the family of equations

$$(2) \quad Z_n = \max(R_n, E(Z_{n+1}|\mathcal{F}_n)), \quad n = 1, 2, \dots$$

The informal interpretation of Z_n is that it is the most one can expect to win if one has already reached stage n ; and equations (2) say that this quantity is the maximum of what one can win by stopping at the n th stage and what one can expect to win by taking at least one more observation and proceeding optimally thereafter. The plausible candidate for an optimal rule is to stop with $N = \min\{n : R_n \geq E(Z_{n+1}|\mathcal{F}_n)\}$, that is, stop as soon as the current reward is at least as large as the most that one can expect to win by continuing. Equations (2) show that $\{Z_n, \mathcal{F}_n\}$ is a supermartingale, while $\{Z_{\min(N,n)}, \mathcal{F}_n\}$ is a martingale. The equations do not have a unique solution, but in the case where the index n is bounded, say $1 \leq n \leq m$ for some given value of m , the solution of interest satisfies $Z_m = R_m$. Hence (2) can be solved and the optimal stopping rule can be found by “backward induction.” The general strategy of optimal stopping theory is to approximate the case where no bound m exists by first imposing such a bound, solving the bounded problem and then letting $m \rightarrow \infty$. It is easy to construct examples where this strategy fails, but it succeeds under broadly applicable conditions. (For examples of its failure, suppose Y_1, Y_2, \dots are independent Bernoulli variables equal to 0 or 1 with probability 1/2 and put $R_n = [n/(n+1)]2^n Y_1 \cdots Y_n$. Then $E(R_{n+1}|\mathcal{F}_n) > R_n$ unless $R_n = 0$, so an optimal rule when one is restricted to stop by stage m is to stop the first time $Y_n = 0$ or at stage m , whichever occurs first. The limit of such rules is to stop as soon as $Y_n = 0$, which leads with probability 1 to the smallest possible reward. For this problem no optimal stopping rule exists. Putting $R_n = -[n/(n+1)]2^n Y_1 \cdots Y_n$ yields an example where the optimal rule when the process is truncated after m observations is to stop after the first observation, no matter what the finite value of m , but the overall optimal rule is to stop the first time $Y_n = 0$, which is not the limit of rules for the truncated process.)

For Wald’s problem of testing a simple hypothesis against a simple alternative, the reward after n observations is the negative (since this problem is formulated in terms of *losses*) of the sum of two terms: (i) the posterior expected loss for an optimal Bayes test based on a sample of fixed size n and (ii) the cost, assumed to be proportional to n . Following the analysis of Arrow, Blackwell and Girshick (1949), the problem was recognized and discussed as an abstract optimal stopping problem by Snell (1952). Bellman (1957) created the field of “dynamic programming” by applying the heuristic principles underlying (2) to a wide class of sequential decision problems, both deterministic and stochastic.

Chow and Robbins (1961, 1963, 1967) [cf. also Siegmund (1967)] clarified and generalized the theoretical foundations of optimal stopping theory. In particular, they gave reasonably general sets of conditions under which an optimal rule exists

and can be computed by evaluating the limit of the truncated problems as $m \rightarrow \infty$. At the same time they identified several intriguing, concrete optimal stopping problems whose simplicity of description can conceal surprising subtleties.

Chow and Robbins (1961) define the “monotone” case, which is the only large class of optimal stopping problems that can be solved explicitly. Chow and Robbins (1965a) is concerned with a reward sequence R_n equal to the proportion of heads in n tosses of a fair coin. Subsequent contributions of importance are due to Dvoretzky (1967), who generalized the results of Chow and Robbins to averages of independent, identically distributed random variables with finite mean and variance, Shepp (1969), who obtained an exact description of the optimal stopping rule for an analogous problem involving Brownian motion and used the invariance principle to relate this to the Chow–Robbins–Dvoretzky version of the problem, and Klass (1973), who dealt with the case where the underlying random variables need not have finite second moment.

Chow, Moriguti, Robbins and Samuels (1964) is concerned with one version of the “secretary problem.” There are m candidates for a secretarial position, who are to be interviewed in a random order. After the n th interview, $n = 1, 2, \dots, m$, the interviewer can rank the candidates already interviewed and hence knows the *relative* ranks of the first n candidates. The interviewer is allowed to hire any of the candidates *at the times of their interviews, but cannot return to an earlier candidate after initiating the next interview*. The goal is to minimize the expected true rank of the candidate hired. (A well-known, much simpler, version of the problem has as the goal of the interviewer to maximize the probability of hiring the candidate who ranks 1 among all the candidates. That version has for large m the solution that the interviewer should wait until me^{-1} of the candidates have been interviewed and then hire the next candidate who has the relative rank of 1 at the time of the interview. The probability of hiring the candidate whose actual rank is 1 converges to e^{-1} .) For the problem of minimizing the expected rank of the selected candidate, Lindley (1961) attempted unsuccessfully to obtain an approximate solution by solving a single, heuristically derived differential equation. Chow, Moriguti, Robbins and Samuels (1964) report that a heuristic solution based on an infinite number of differential equations suggests that the limit as $m \rightarrow \infty$ of the *minimal* expected true rank of the secretary hired equals

$$(3) \quad \prod_{j=1}^{\infty} [1 + 2/j]^{1/(j+1)} \approx 3.8695.$$

In view of their inability to make this approach rigorous, they analyze the backward induction equations (2) characterizing the optimal rule to give a direct proof of (3). The use of differential equations was made precise by Mucci (1973) by direct calculation. An ingenious alternative approach was suggested by Rubin (1966), where in a brief abstract he described a version of the problem involving an infinite number of candidates. This approach and its natural relation to the

original problem were developed by Gianini and Samuels (1976), Gianini (1977) and Lorenzen (1979). For reviews of the many variations on this problem and the extensive related literature, see Freeman (1983), Petrucelli (1988) and Samuels (1991).

The general theory of optimal stopping and a number of examples are described in the books by Chow, Robbins and Siegmund (1971) and Shiryaev (1978). Using related arguments, Chow, Robbins and Teicher (1965) gave an elegant generalization of Wald's identity from the first moment to the second and higher moments.

Whereas optimal stopping theory in discrete time as described above is, in principle, completely understood, there are admittedly relatively few cases where one can compute an optimal rule without the use of numerical methods. For optimal stopping problems in continuous time, one finds analytic advantages in the computation of optimal rules and foundational difficulties. The pioneering efforts to exploit the advantages were due to Chernoff (1961), Shiryaev (1963) and Bather (1962). Moriguti and Robbins (1962) should also be mentioned, but this appears to have been a one-time diversion. The interplay between the analytic simplifications of continuous time, which often allow a substantial reduction in the dimensionality of the parameter space, and the numerical advantages of the equations (2) in discrete time has been beautifully exploited by Chernoff [e.g., Chernoff (1961, 1972) and Chernoff and Petkau (1985)]. In spite of considerable success in solving specific continuous-time problems, a completely general theoretical foundation seems still to be lacking. See Chernoff (1972) and Shiryaev (1978). For a thoroughly "modern" application to mathematical finance, see the AitSahlia–Lai (1999) solution of the American option.

At the same time that Chow and Robbins were developing their general theory of optimal stopping, there were two parallel, rather different approaches to conceptually similar problems. These were Blackwell's investigations of the foundations of dynamic programming [Blackwell (1962, 1965, 1967)] and Dubins and Savage's (1965) gambling theory. Unlike optimal stopping theory, neither of these problems starts from a fixed probability space with a given stochastic process defined on it, and one may encounter an uncountably infinite number of possible actions at each stage. As a consequence, there are substantial technical impediments to a rigorous interpretation of (an appropriate version of) (2), which, in turn, lead to some reflection on the proper role of measure theory at the foundation of probability theory. For a beautiful survey of a rich variety of models and applications that neatly avoids technical issues of foundations, see Whittle (1982, 1983).

4. Sequential estimation. Sequential hypothesis testing purports to be a more efficient way to accomplish something that can be accomplished with a fixed sample size. For problems of estimation, sequential methods allow the

experimenter to obtain estimators of specified precision, which is often impossible using a fixed sample.

For example, the usual confidence interval for a normal mean based on a sample of size n has a random length proportional to the sample standard deviation, s_n . Stein (1945) showed that a sequential (in fact, a two stage) procedure could produce a confidence interval having a length that is prescribed by the experimenter and does not depend on the variance of the distribution. After taking $n_0 \geq 2$ observations, one uses s_{n_0} to choose an appropriate random sample size $N - n_0$ for a second stage of sampling. Since Stein's procedure uses only s_{n_0} to estimate the population standard deviation, even though the (often much larger) final sample size contains more information, it presumably is inefficient. Stein (1949) and Anscombe (1953) suggested completely sequential methods, where one recalculates the sample variance after each observation and stops sampling as soon (after some minimum sample size) as the empirical length of the standard confidence interval is smaller than the desired length. Stein very briefly and Anscombe more systematically showed by clever heuristic analyses that this confidence interval has asymptotically the desired coverage probability and that the stopping rule is fully efficient asymptotically as the prescribed length converges to 0, in the sense that its expectation is asymptotic to what would be required if one *knew* the variance and used the easily calculated fixed sample size required to achieve the prescribed length.

Chow and Robbins (1965b) gave a rigorous version of (parts of) Anscombe's argument and generalized it by showing that for a first-order asymptotic analysis the assumption of a normal distribution could be replaced by application of an appropriate central limit theorem [due to Anscombe (1952)]. This pointed the way to a large number of generalizations and extensions [e.g., Gleser (1965), Simons (1968), Sen and Ghosh (1971), Srivastava (1971)]. Woodroffe (1977) [see also Woodroffe (1986)] used the nonlinear renewal theory that he had developed in the context of hypothesis testing (see below) to justify the second-order asymptotic approximations of Anscombe (1953). See also Hall (1981), who showed that similar asymptotic results could be obtained with three stages of sampling.

Motivated by Robbins and Siegmund (1976), which is concerned with fixed-width confidence intervals for the log odds, Siegmund (1982) and Lai and Siegmund (1983) studied examples (the log odds ratio in paired Bernoulli trials and the first-order nonexplosive autoregression, respectively) where the first-order asymptotic behavior of the fixed-width sequential confidence interval is found to hold uniformly in the unknown parameters over the entire parameter space. This contrasts with the corresponding fixed-sample-size procedures where asymptotic convergence near the ends of the parameter space is in one case very slow and in the other involves a completely different limit from what one finds in the interior of the parameter space.

5. Sequential hypothesis testing. Although the beginning of sequential analysis in Wald's wartime research centered on hypothesis testing, most of Robbins' research in sequential hypothesis testing occurred only in the late 1960s and 1970s. An important exception is found in Robbins (1952), where he commented on an earlier paper of Feller (1940) about statistical methods in extrasensory perception. Feller had observed that protocols in extrasensory perception did not specify an unambiguous sample size, but seemed to allow the experimenter the flexibility to choose the sample size based on the unfolding sequence of data. Thus, in testing a hypothesis that a coin is fair by observing a sequence of tosses of the coin, apparently one might choose the data-dependent sample size

$$(4) \quad N = \inf\{n : |S_n| \geq cn^{1/2}\},$$

where S_n equals the accumulated excess of heads over tails in the first n tosses and c is some suitable (large) constant. On the one hand, for a fair coin and any large fixed value of n , one knows from the central limit theorem that $P\{|S_n| \geq cn^{1/2}\} \approx 2[1 - \Phi(c)]$, where Φ is the standard normal distribution. For large c this probability is small, so the occurrence of the event that $|S_n| \geq cn^{1/2}$ would be interpreted in favor of the alternative hypothesis that the probability p of heads differs from $1/2$. On the other hand, by the law of the iterated logarithm, $P\{|S_n| \geq cn^{1/2} \text{ for infinitely many } n\} = 1$, so by choosing the value of n to be the random variable N defined in (4), one could be sure to reject a true null hypothesis. Feller concluded that consequently these tests were invalid. Robbins countered that although what Feller wrote was mathematically true, it might not render the tests completely invalid if one took into account what might be reasonable practical constraints in the choice of N . To illustrate his point, he discussed normally distributed random variables, which under the null hypothesis were to have mean 0, under the alternative were to have nonzero mean and in all cases were to have unit variance. He assumed that the experimenter was restricted to take a sample size that contained at least m_0 and at most m_1 observations. For the probability

$$(5) \quad P\{S_n \geq cn^{1/2} \text{ for some } m_0 \leq n \leq m_1\},$$

he then derived an upper bound in the form of a function of c and $(m_1/m_0 - 1)$, which is small if c is sufficiently large and m_1/m_0 is not too large. Robbins concludes that by imposing a suitable restriction on the range of possible values of the random sample size, one could devise a test that allowed for some flexibility in the choice of sample size without falling prey to the difficulty identified by Feller, and he poses the problem of finding a good approximation for the probability (5).

The same statistical issue of sampling to a foregone conclusion was addressed again in Darling and Robbins (1967) with an altogether different approach. It was assumed that a sequence of independent and normally distributed random variables with mean μ and variance 1 is observed sequentially. If $\mu \leq 0$, one wants to

observe the process indefinitely, but if μ is positive, one is required to indicate this condition as soon as possible. This leads to the notion of an α -level *test of power one* of the hypothesis $\mu \leq 0$ against the alternative $\mu > 0$, which is by definition a stopping rule N with the property that $P_\mu(N < \infty) \leq \alpha$ (< 1) for all $\mu \leq 0$, while $P_\mu(N < \infty) = 1$ for all $\mu > 0$. Invoking the stopping rule is interpreted as a declaration that $\mu > 0$. In addition to satisfying the constraints on $P_\mu(N < \infty)$, one wants to minimize in some suitable sense the expected amount of sampling, $E_\mu(N)$, for $\mu > 0$. Darling and Robbins (1967) showed that a test of power one can be obtained as a stopping rule of the form $N = \inf\{n : S_n \geq c_n\}$ ($\inf \phi = +\infty$), where $c_n \geq cn^{1/2}$ for large n , to avoid the problem of the law of the iterated logarithm identified by Feller, and $c_n/n \rightarrow 0$ as $n \rightarrow \infty$. They also discussed a two-sided version of the problem along with the related notion of a *confidence sequence*: a sequence of random intervals I_n with the property that $P_\mu(\mu \in I_n \text{ for all } n \geq 1) \geq 1 - \alpha$.

The problem was discussed in a number of subsequent papers, for example, Robbins (1970) and Robbins and Siegmund (1970, 1974b). Cornfield (1966) is also motivated by the issue of optional stopping and arrives at somewhat similar ideas, expressed from a Bayesian perspective. Related results of considerable technical virtuosity motivated by quite different statistical issues are found in Farrell (1964).

To describe these and related results, let Y_1, Y_2, \dots be a sequence of random variables and let \mathcal{F}_n denote the σ -algebra generated by Y_1, \dots, Y_n . Let P and Q denote two probability measures on the sequence of Y 's and let P_n and Q_n denote the restriction of these probabilities to \mathcal{F}_n . (In applications P will denote probability under a specific hypothesis, e.g., that the Y 's are independent and normally distributed with mean 0 and variance 1, while Q is related in a natural way to an alternative hypothesis.) Suppose that Q_n and P_n are mutually absolutely continuous for every value of n and let $Z_n = \log dQ_n/dP_n$. For any stopping time, possibly infinite valued, following Wald, we have

$$(6) \quad P\{N < \infty\} = \sum_1^\infty P\{N = n\} = \sum_1^\infty \int_{\{N=n\}} e^{-Z_n} dQ = \int_{\{N < \infty\}} e^{-Z_N} dQ.$$

For the special case $N = \inf\{n : Z_n \geq b\}$ ($\inf \phi = +\infty$), we obtain

$$(7) \quad P\{N < \infty\} = \exp(-b) \int_{\{N < \infty\}} \exp[-(Z_N - b)] dQ \leq \exp(-b).$$

Similar arguments in the special case that Y_1, Y_2, \dots are independent, identically distributed random variables under both P and Q are the basis of Wald's analysis of the sequential probability ratio test. For example, let P_μ be the probability that makes the Y 's independently and normally distributed with mean μ and variance 1. Let $P = P_0$ and $Q = P_{\mu_1}$ for some fixed $\mu_1 > 0$. Then

$$(8) \quad Z_n = \mu_1 S_n - n\mu_1^2/2,$$

so the inequality $Z_n \geq b$ is equivalent to $S_n \geq n\mu_1/2 + b/\mu_1$, and (7) yields a well-known inequality. However, since $P_\mu\{N = \infty\} > 0$ whenever $0 < \mu < \mu_1/2$, this special case cannot be used for a test of power one.

Now let $P = P_0$ be as above and

$$(9) \quad Q = \int_{-\infty}^{\infty} P_\mu \phi(\mu) d\mu,$$

where ϕ is the standard normal density. Then

$$(10) \quad Z_n = [S_n^2/(n+1) - \log(n+1)]/2,$$

so $Z_n \geq b$ is equivalent to

$$(11) \quad |S_n| \geq \{(n+1)[\log(n+1) + 2b]\}^{1/2}.$$

This can be used to define a two-sided test of power one of $\mu = 0$ against $\mu \neq 0$, and by inversion a confidence sequence for the normal mean μ . The case where the integral in (9) is restricted to the positive half-line yields a one-sided test of power one.

As suggested above, there are natural theoretical questions connecting tests of power one and the law of the iterated logarithm. See Farrell (1964) and Robbins and Siegmund (1970, 1974b). Indeed, the investigation of these issues in Robbins and Siegmund (1970) led to the discovery of an error in the classical paper of Feller (1946), with complete resolution in Bai (1989), who used a result in Feller (1970).

Although tests of power one and confidence sequences are theoretically and intuitively appealing, because of their unbounded and indeed possibly infinite sample size they appear to have no direct application. However, by focusing attention on simple conceptual and technical problems, these and related ideas have proved very useful in the design and analysis of practical procedures. Some examples follow.

(i) Armitage, McPherson and Rowe (1969), Armitage (1975) and Pocock (1977) suggested the use of "repeated significance tests" for clinical trials. In this context, the random variable Y_n represents the difference in response to two treatments within the n th pair in a sequential clinical trial, or the sum of such differences within the n th group in a group sequential trial. For simplicity and because a number of more complicated models reduce to this one asymptotically by the central limit theorem, the Y 's are assumed to be independent and normally distributed with mean μ and variance 1. The hypothesis of no treatment effect is $\mu = 0$. The sample size is $\min(N, m)$, N is defined by (4) and m is a fixed upper bound for the sample size. The hypothesis is rejected if and only if $N \leq m$, so the significance level is $P_0\{N \leq m\}$. This is closely related to the probability in (5) and was evaluated by Armitage, McPherson and Rowe (1969) by repeated numerical integration.

Theoretical approximations to the significance level of a repeated significance test and other related sequential tests were given by Lai and Siegmund (1977), who used the representation (6) and the log likelihood ratio arising from (9), and by Woodroffe (1976), who developed a completely different method having roots in Anscombe (1953). Generalizations to multiparameter exponential families were obtained by Woodroffe (1978), Lalley (1983) and Hu (1988). See Woodroffe (1982) and Siegmund (1985) for more complete discussions and related references.

(ii) Similarly, the expected sample size approximations given by Siegmund (1985) for repeated significance tests build on the earlier results for tests of power one obtained by Robbins and Siegmund (1974b), Pollak and Siegmund (1975) and Lai and Siegmund (1979).

(iii) Sequential change-point detection utilizes methods discussed in the preceding paragraphs in conjunction with optimal stopping theory. Lorden (1971) observed that the sequential CUSUM test of Page (1954) for detection of a change in a normal mean from an initial value of $\mu = 0$ to the value $\mu = \mu_1 \neq 0$ can be described as follows: for each $k = 0, 1, 2, \dots$, let N_k denote N applied to the "shifted" data $Z_{k+n} - Z_k$, $n = 1, 2, \dots$, where Z_n is defined in (8). Detection of a change occurs at $\min(N_k + k)$. Lorden used inequality (7) in proving the asymptotic optimality of this sequential CUSUM test. Siegmund and Venkatraman (1995) and Pollak (1985, 1987) give an in-depth study of a related stopping rule suggested initially by Shirayev (1963), which was derived from optimal stopping theory applied to a particular Bayesian version of the problem.

(iv) For the fixed-sample problem of testing whether a sequence of normally distributed observations has a constant mean against the alternative of at most one change in the mean when (for simplicity) the variance is known, the log likelihood ratio statistic is

$$\max_{1 \leq j \leq n} |S_j - jS_n/n|/[j(1 - j/n)]^{1/2},$$

which is closely related to the $\max_{1 \leq j \leq n} |S_j|/j^{1/2}$ of a repeated significance test. Siegmund (1985, 1986) uses a mixture of likelihood ratios similar to (9) and an argument along the lines of Lai and Siegmund (1977) to give an approximation to the significance level. A related problem of interest is to test the hypothesis of a constant mean value against the alternative that over some interval, say from j to k , where $j < k$ are both unknown, the observations have a mean value that differs from the baseline value that they have outside the interval $[j, k]$. Since this test involves a two-dimensional maximization with respect to the two putative change points, an argument using (6), which depends strongly on the linear structure of the indexing set and the stopping time N , does not appear to generalize. Siegmund (1988) has adapted Woodroffe's (1976) alternative method to deal with this and other problems involving maxima of multiply indexed random fields. Yakir and Pollak (1998) have recently introduced a change of measure different from (6),

which does not use stopping times or intrinsically require a linearly ordered indexing set. See Siegmund and Yakir (2000a) for applications to a number of change-point-like problems and Siegmund and Yakir (2000b) for an application to pairwise local sequence alignments in searching protein data bases.

(v) Statistical analysis of genetic mapping leads naturally to problems very similar to that of testing for a constant normal mean against the alternative of at most one change, but now the basic underlying observations are the discrete skeleton of an Ornstein–Uhlenbeck-like process, which can be studied by a change of measure or by Woodroffe’s (1976) method. See Feingold, Brown and Siegmund (1993) for a detailed discussion.

Another example of the intersection of sequential hypothesis testing and optimal stopping theory, which again involves Robbins’ reaction to the ideas of Anscombe, arises out of Anscombe’s (1963) review of the first edition of Armitage’s *Sequential Medical Trials* [Armitage (1960)]. Anscombe criticized Armitage’s formulation on the grounds that it was not explicitly decision theoretic or Bayesian in its approach. For an idealized model, which involved pairwise allocation of two treatments until a (random) time when a decision is made to treat all future patients up to a previously specified horizon with what was judged to be the better of the two treatments, Anscombe formulated a set of costs that balanced the risk of making a wrong decision against the delay in choosing the better treatment for application to all patients in the future. After assigning a prior distribution to the difference in mean values of responses to the two treatments, the problem becomes one of optimal stopping. Anscombe conjectured an approximation to an optimal rule, which leads to something quite different from the tests suggested in Armitage (1960) and the repeated significance tests he favored in the second edition of his book in 1975. Lai, Levin, Robbins and Siegmund (1980) showed that Anscombe’s suggested procedure is not particularly close to the optimal stopping rule, but it nevertheless performs about as well as the optimal rule and another heuristically motivated rule they suggested. Chernoff and Petkau (1981, 1985) have made related contributions from a similar point of view. Armitage (1985) discusses reasons why these models have not found their way into the practice of clinical trials.

6. Summary. I have discussed Robbins’ research in sequential analysis—its relation to previous and to contemporaneous research and its influence on subsequent research. We see that unlike stochastic approximation and empirical Bayes, where his first articles were arguably unprecedented, sequential analysis was a rapidly developing subject when Robbins made his first contributions. Consequently, his recurrent contributions can be seen as a dialogue with his contemporaries that has influenced succeeding generations. Although necessarily selective, I hope this review has shown that by virtue of his ability to identify simple

conceptual problems containing the germ of important general ideas, Herbert Robbins' research in sequential analysis was a major intellectual achievement that has changed forever the way we think about a large class of challenging problems.

REFERENCES

- AITSAHLIA, F. and LAI, T. L. (1999). A canonical optimal stopping problem for American options and its numerical solution. *J. Comput. Finance* **3**(2) 33–52.
- ANSCOMBE, F. J. (1946). Linear sequential rectifying inspection for controlling fraction defective. *J. Roy. Statist. Soc. Suppl.* **8** 216–222.
- ANSCOMBE, F. J. (1952). Large-sample theory of sequential estimation. *Proc. Cambridge Philos. Soc.* **48** 600–607.
- ANSCOMBE, F. J. (1953). Sequential estimation. *J. Roy. Statist. Soc. Ser. B* **15** 1–29.
- ANSCOMBE, F. J. (1963). Sequential medical trials. *J. Amer. Statist. Assoc.* **58** 365–383.
- ARMITAGE, P. (1960). *Sequential Medical Trials*. Blackwell, Oxford.
- ARMITAGE, P. (1975). *Sequential Medical Trials*, 2nd ed. Blackwell, Oxford.
- ARMITAGE, P. (1985). The search for optimality in clinical trials (with discussion). *Internat. Statist. Rev.* **53** 15–36.
- ARMITAGE, P., MCPHERSON, C. K. and ROWE, B. C. (1969). Repeated significance tests on accumulating data. *J. Roy. Statist. Soc. Ser. A* **132** 235–244.
- ARROW, K. J., BLACKWELL, D. and GIRSHICK, M. A. (1949). Bayes and minimax solutions of sequential decision problems. *Econometrica* **17** 213–244.
- BAI, Z. D. (1989). A theorem of Feller revisited. *Ann. Probab.* **17** 385–395.
- BARNARD, G. A. (1946). Sequential tests in industrial statistics. *J. Roy. Statist. Soc. Suppl.* **8** 1–26.
- BATHER, J. A. (1962). Bayes procedures for deciding the sign of a normal mean. *Proc. Cambridge Philos. Soc.* **58** 599–620.
- BATHER, J. A. (1985). On the allocation of treatments in sequential medical trials (with discussion). *Internat. Statist. Rev.* **53** 1–13, 25–36.
- BELLMAN, R. (1956). A problem in the sequential design of experiments. *Sankhyā* **16** 221–229.
- BELLMAN, R. (1957). *Dynamic Programming*. Princeton Univ. Press.
- BLACKWELL, D. (1962). Discrete dynamic programming. *Ann. Math. Statist.* **33** 719–726.
- BLACKWELL, D. (1965). Discounted dynamic programming. *Ann. Math. Statist.* **36** 226–235.
- BLACKWELL, D. (1967). Positive dynamic programming. *Proc. Fifth Berkeley Symp. Math. Statist. Probab.* **1** 415–418. Univ. California Press, Berkeley.
- BRADT, R. N., JOHNSON, S. M. and KARLIN, S. (1956). On sequential designs for maximizing the sum of n observations. *Ann. Math. Statist.* **27** 1060–1074.
- CHERNOFF, H. (1961). Sequential tests for the mean of a normal distribution. *Proc. Fourth Berkeley Symp. Math. Statist. Probab.* **1** 79–91. Univ. California Press, Berkeley.
- CHERNOFF, H. (1972). *Sequential Analysis and Optimal Design*. SIAM, Philadelphia.
- CHERNOFF, H. and PETKAU, A. J. (1981). Sequential medical trials involving paired data. *Biometrika* **68** 119–132.
- CHERNOFF, H. and PETKAU, A. J. (1985). Sequential medical trials with ethical cost. In *Proc. of the Berkeley Conference in Honor of J. Neyman and J. Kiefer* (L. Le Cam and R. Olshen, eds.) **2** 521–537. Wadsworth, Belmont, CA.
- CHOW, Y. S., MORIGUTI, S., ROBBINS, H. and SAMUELS, S. M. (1964). Optimal selection based on relative rank—the “secretary” problem. *Israel J. Math.* **2** 81–90.
- CHOW, Y. S. and ROBBINS, H. E. (1961). A martingale system theorem and applications. *Proc. Fourth Berkeley Symp. Math. Statist. Probab.* **1** 93–104. Univ. California Press, Berkeley.
- CHOW, Y. S. and ROBBINS, H. E. (1963). On optimal stopping rules. *Z. Wahrsch. Verw. Gebiete* **2** 33–49.

- CHOW, Y. S. and ROBBINS, H. E. (1965a). On optimal stopping rules for S_n/n . *Illinois J. Math.* **9** 444–454.
- CHOW, Y. S. and ROBBINS, H. E. (1965b). On the asymptotic theory of fixed-width sequential confidence intervals for the mean. *Ann. Math. Statist.* **36** 457–462.
- CHOW, Y. S. and ROBBINS, H. E. (1967). On values associated with a stochastic sequence. *Proc. Fifth Berkeley Symp. Math. Statist. Probab.* **1** 427–440. Univ. California Press, Berkeley.
- CHOW, Y. S., ROBBINS, H. E. and SIEGMUND, D. (1971). *Great Expectations: The Theory of Optimal Stopping*. Houghton Mifflin, Boston.
- CHOW, Y. S., ROBBINS, H. E. and TEICHER, H. (1965). Moments of randomly stopped sums. *Ann. Math. Statist.* **36** 789–799.
- CORNFIELD, J. (1966). A Bayesian test of some classical hypotheses—with applications to sequential clinical trials. *J. Amer. Statist. Assoc.* **61** 577–594.
- COVER, T. M. (1968). A note on the two-armed bandit problem with finite memory. *Information and Control* **12** 371–377.
- DARLING, D. A. and ROBBINS, H. E. (1967). Iterated logarithm inequalities. *Proc. Nat. Acad. Sci. U.S.A.* **57** 1188–1192.
- DODGE, H. F. and ROMIG, H. G. (1929). A method of sampling inspection. *Bell System Tech. J.* **8** 613–631.
- DUBINS, L. and SAVAGE, L. J. (1965). *How to Gamble if You Must: Inequalities for Stochastic Processes*. McGraw-Hill, New York.
- DVORETZKY, A. (1967). Existence and properties of certain optimal stopping rules. *Proc. Fifth Berkeley Symp. Math. Statist. Probab.* **1** 441–452. Univ. California Press, Berkeley.
- FARRELL, R. (1964). Asymptotic behavior of expected sample size in certain one-sided tests. *Ann. Math. Statist.* **35** 36–72.
- FEINGOLD, E., BROWN, P. O. and SIEGMUND, D. (1993). Gaussian models for genetic linkage analysis using complete high resolution maps of identity-by-descent. *Amer. J. Human Genetics* **53** 234–251.
- FELLER, W. (1940). Statistical aspects of ESP. *J. Parapsychology* **4** 271–298.
- FELLER, W. (1946). The law of the iterated logarithm for identically distributed variables. *Ann. Math.* **47** 631–638.
- FELLER, W. (1970). On the oscillations of sums of independent random variables. *Ann. Math.* **91** 402–418.
- FLEHINGER, B., LOUIS, T., ROBBINS, H. and SINGER, B. (1972). Reducing the number of inferior treatments in clinical trials. *Proc. Nat. Acad. Sci. U.S.A.* **69** 2993–2994.
- FREEMAN, P. R. (1983). The secretary problem and its extensions: A review. *Internat. Statist. Rev.* **51** 189–206.
- GIANINI, J. (1977). The infinite secretary problem as the limit of the finite problem. *Ann. Probab.* **5** 636–644.
- GIANINI, J. and SAMUELS, S. (1976). The infinite secretary problem. *Ann. Probab.* **4** 418–432.
- GITTINS, J. (1989). *Multi-armed Bandit Allocation Indices*. Wiley, New York.
- GITTINS, J. and JONES, D. M. (1974). A dynamic allocation index for the sequential design of experiments. In *Progress in Statistics* (J. Gani, K. Sarkardi and I. Vince, eds.) 241–266. North-Holland, Amsterdam.
- GLESER, L. (1965). On the asymptotic theory of fixed-size sequential confidence bounds for linear regression parameters. *Ann. Math. Statist.* **36** 463–467. [Corrigendum (1966) *Ann. Math. Statist.* **37** 1053–1055.]
- HALL, P. (1981). Asymptotic theory of triple sampling for sequential estimation of a mean. *Ann. Statist.* **9** 1229–1238.
- HAYRE, L. S. (1979). Two-population sequential tests with three hypotheses. *Biometrika* **66** 465–474.

- HELLMAN, M. and COVER, T. M. (1970). Learning with finite memory. *Ann. Math. Statist.* **41** 765–782.
- HU, I. (1988). Repeated significance tests for exponential families. *Ann. Statist.* **16** 1643–1666.
- KLASS, M. J. (1973). Properties of optimal extended-valued stopping rules for S_n/n . *Ann. Probab.* **1** 719–757.
- LAI, T. L. (2001). Sequential analysis: Some classical problems and new challenges (with discussion). *Statist. Sinica* **11** 303–408.
- LAI, T. L., LEVIN, B., ROBBINS, H. E. and SIEGMUND, D. (1980). Sequential medical trials. *Proc. Nat. Acad. Sci. U.S.A.* **77** 3135–3138.
- LAI, T. L. and SIEGMUND, D. (1977). A nonlinear renewal theory with applications to sequential analysis I. *Ann. Statist.* **5** 946–954.
- LAI, T. L. and SIEGMUND, D. (1979). A nonlinear renewal theory with applications to sequential analysis II. *Ann. Statist.* **7** 60–76.
- LAI, T. L. and SIEGMUND, D. (1983). Fixed accuracy estimation of an autoregressive parameter. *Ann. Statist.* **11** 478–485.
- LALLEY, S. (1983). Repeated likelihood ratio tests for curved exponential families. *Z. Wahrsch. Verw. Gebiete* **62** 293–321.
- LINDLEY, D. V. (1961). Dynamic programming and decision theory. *Appl. Statist.* **10** 39–51.
- LORDEN, G. (1971). Procedures for reacting to a change in distribution. *Ann. Math. Statist.* **42** 1897–1908.
- LORENZEN, T. (1979). Generalizing the secretary problem. *Adv. in Appl. Probab.* **11** 384–396.
- LOUIS, T. (1975). Optimal allocation in sequential tests comparing the means of two Gaussian populations. *Biometrika* **62** 359–369.
- MORIGUTI, S. and ROBBINS, H. E. (1962). A Bayes test of $p \leq 1/2$ versus $p > 1/2$. *Rep. Statist. Appl. Res. Un. Japan. Sci. Engrs.* **9** 39–60.
- MUCCI, A. G. (1973). On a class of secretary problems. *Ann. Probab.* **1** 417–427.
- PAGE, E. S. (1954). Continuous inspection schemes. *Biometrika* **41** 100–115.
- PETRUCELLI, J. D. (1988). Secretary problem. In *Encyclopedia of Statistical Sciences* (S. Kotz and N. Johnson, eds.) **8** 326–329. Wiley, New York.
- POCOCK, S. (1977). Group sequential methods in the design and analysis of clinical trials. *Biometrika* **64** 191–200.
- POLLAK, M. (1985). Optimal detection of a change in distribution. *Ann. Statist.* **13** 206–227.
- POLLAK, M. (1987). Average run lengths of an optimal method of detecting a change in distribution. *Ann. Statist.* **15** 749–779.
- POLLAK, M. and SIEGMUND, D. (1975). Approximations to the expected sample size of certain sequential tests. *Ann. Statist.* **3** 1267–1282.
- ROBBINS, H. E. (1951). Asymptotically subminimax solutions of compound statistical decision problems. *Proc. Second Berkeley Symp. Math. Statist. Probab.* **1** 131–148. Univ. California Press, Berkeley.
- ROBBINS, H. E. (1952). Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.* **58** 527–535.
- ROBBINS, H. E. (1956). A sequential decision problem with finite memory. *Proc. Nat. Acad. Sci. U.S.A.* **42** 920–923.
- ROBBINS, H. E. (1970). Statistical methods related to the law of the iterated logarithm. *Ann. Math. Statist.* **41** 1397–1409.
- ROBBINS, H. E. and MONRO, S. (1951). A stochastic approximation method. *Ann. Math. Statist.* **22** 400–407.
- ROBBINS, H. E. and SIEGMUND, D. (1970). Boundary crossing probabilities for the Wiener process and sample sums. *Ann. Math. Statist.* **41** 1410–1429.

- ROBBINS, H. E. and SIEGMUND, D. (1974a). Sequential tests involving two populations. *J. Amer. Statist. Assoc.* **69** 132–139.
- ROBBINS, H. E. and SIEGMUND, D. (1974b). The expected sample size of some tests of power one. *Ann. Statist.* **2** 415–436.
- ROBBINS, H. E. and SIEGMUND, D. (1974c). Sequential estimation of p in Bernoulli trials. In *Studies in Probability and Statistics* (E. J. Williams, ed.) 103–107. Jerusalem Acad. Press.
- RUBIN, H. (1966). The “secretary” problem (abstract). *Ann. Math. Statist.* **37** 544.
- SAMUELS, S. M. (1991). Secretary problems. In *Handbook of Sequential Analysis* (B. K. Ghosh and P. K. Sen, eds.) 381–405. Dekker, New York.
- SEN, P. K. and GHOSH, M. (1971). On bounded length sequential confidence intervals based on one-sample rank order statistics. *Ann. Math. Statist.* **42** 189–203.
- SHEPP, L. (1969). Explicit solutions to some problems of optimal stopping. *Ann. Math. Statist.* **40** 993–1010.
- SHEWHART, W. A. (1931). *Economic Control of Quality of Manufactured Product*. Van Nostrand, Princeton.
- SHIRYAYEV, A. N. (1963). On optimum methods in quickest detection problems. *Theory Probab. Appl.* **8** 22–46.
- SHIRYAYEV, A. N. (1978). *Optimal Stopping Rules*. Springer, New York.
- SIEGMUND, D. (1967). Some problems in the theory of optimal stopping rules. *Ann. Math. Statist.* **38** 1627–1640.
- SIEGMUND, D. (1982). A sequential confidence interval for the odds ratio. *Probab. Math. Statist.* **2** 149–156.
- SIEGMUND, D. (1985). *Sequential Analysis: Tests and Confidence Intervals*. Springer, Heidelberg.
- SIEGMUND, D. (1986). Boundary crossing probabilities and statistical applications. *Ann. Statist.* **14** 361–404.
- SIEGMUND, D. (1988). Approximate tail probabilities for the maxima of some random fields. *Ann. Probab.* **16** 487–501.
- SIEGMUND, D. and VENKATRAMAN, E. S. (1995). Using the generalized likelihood ratio statistic for sequential detection of a change-point. *Ann. Statist.* **23** 255–271.
- SIEGMUND, D. and YAKIR, B. (2000a). Tail probabilities for the null distribution of scanning statistics. *Bernoulli* **6** 191–213.
- SIEGMUND, D. and YAKIR, B. (2000b). Approximate p -values for local sequence alignments. *Ann. Statist.* **28** 657–680.
- SIMONS, G. (1968). On the cost of not knowing the variance when making a fixed-width confidence interval for the mean. *Ann. Math. Statist.* **39** 1946–1952.
- SNELL, J. L. (1952). Applications of martingale system theorems. *Trans. Amer. Math. Soc.* **73** 293–312.
- SRIVASTAVA, M. S. (1971). On fixed-width confidence bounds for regression parameters. *Ann. Math. Statist.* **42** 1403–1411.
- STEIN, C. (1945). A two sample test for a linear hypothesis whose power is independent of the variance. *Ann. Math. Statist.* **16** 243–258.
- STEIN, C. (1949). Some problems in sequential estimation. *Econometrica* **17** 77–78.
- WALD, A. (1947). *Sequential Analysis*. Wiley, New York.
- WALD, A. and WOLFOWITZ, J. (1948). Optimum character of the sequential probability ratio test. *Ann. Math. Statist.* **19** 326–339.
- WARE, J. (1989). Investigating therapies of potentially great benefit: ECMO (with discussion). *Statist. Sci.* **4** 298–340.
- WHITTLE, P. (1982). *Optimization Over Time: Dynamic Programming and Stochastic Control 1*. Wiley, New York.

- WHITTLE, P. (1983). *Optimization Over Time: Dynamic Programming and Stochastic Control* **2**. Wiley, New York.
- WOODROOFE, M. (1976). Frequentist properties of Bayesian sequential tests. *Biometrika* **63** 101–110.
- WOODROOFE, M. (1977). Second order approximations for sequential point and interval estimation. *Ann. Statist.* **5** 984–995.
- WOODROOFE, M. (1978). Large deviations of the likelihood ratio statistic with applications to sequential testing. *Ann. Statist.* **6** 72–84.
- WOODROOFE, M. (1982). *Nonlinear Renewal Theory in Sequential Analysis*. SIAM, Philadelphia.
- WOODROOFE, M. (1986). Asymptotic optimality in sequential interval estimation. *Adv. in Appl. Math.* **7** 70–79.
- YAKIR, B. and POLLAK, M. (1998). A new representation for a renewal-theoretic constant appearing in asymptotic approximations of large deviations. *Ann. Appl. Probab.* **8** 749–774.

DEPARTMENT OF STATISTICS
STANFORD UNIVERSITY
STANFORD, CALIFORNIA 94305
E-MAIL: dos@stat.stanford.edu