



Herbert Robbins

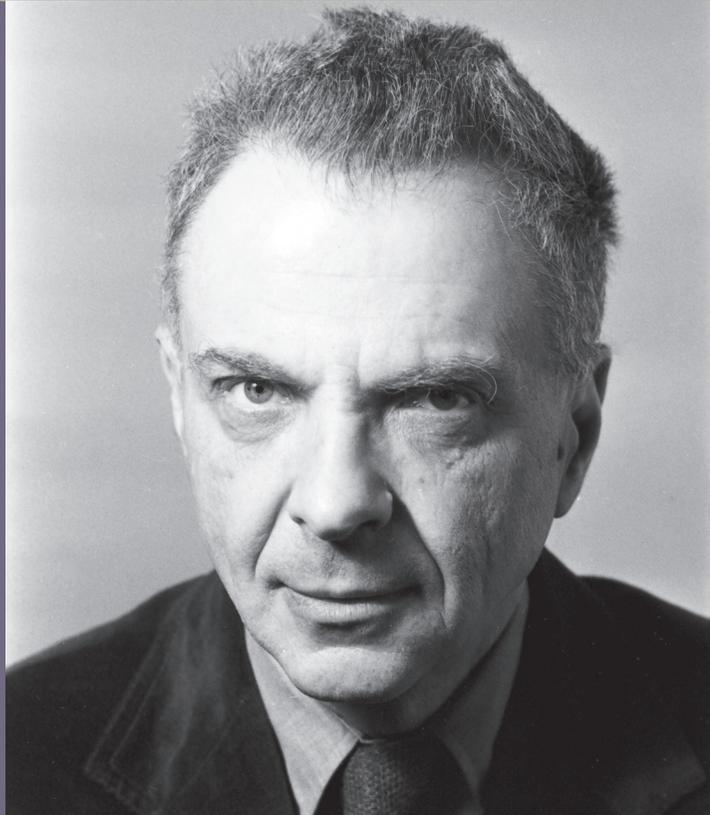
1915–2001

BIOGRAPHICAL

Memoirs

*A Biographical Memoir by
Tze Leung Lai
and David Siegmund*

©2018 National Academy of Sciences.
Any opinions expressed in this memoir are
those of the authors and do not
necessarily reflect the views of the
National Academy of Sciences.



NATIONAL ACADEMY OF SCIENCES

HERBERT ELLIS ROBBINS

January 12, 1915–February 12, 2001

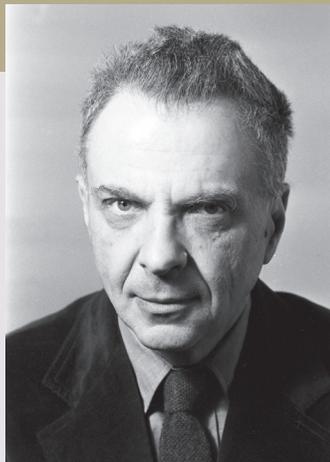
Elected to the NAS, 1974

Herbert Robbins was widely recognized as one of the most creative and influential mathematical statisticians from 1950 to the turn of the millennium. Entering the then-nascent field of statistics serendipitously in the 1940s, he showed daring originality in three major innovations, all of which occurred during 1951–1956:

- (a) Compound statistical decision theory and empirical Bayes
- (b) Sequential design of experiments and multi-armed bandits
- (c) Stochastic approximation and recursive algorithms in stochastic optimization and control.

Although the importance of these contributions was immediately recognized by Robbins' contemporaries in statistics and control engineering, their full impacts were not realized until the big-data era in science and technology that arose after his death.

Born in New Castle, PA, Robbins entered Harvard University in 1931, at the age of 16. Although his interests until then had been predominantly literary, he found himself increasingly attracted to mathematics through the influence of Marston Morse, who during many long conversations conveyed a vivid sense of the intellectual challenge of creative work in that field. Robbins received the A.B. summa cum laude in 1935, and his Ph.D. in 1938, both in mathematics, from Harvard. His thesis, in the field of combinatorial topology and written under the supervision of Hassler Whitney, was published in 1941. After graduation, Robbins worked for a year at the Institute for Advanced Study in Princeton, NJ, as Marston Morse's assistant. He then spent the next three years at New York University as an instructor of mathematics. He became nationally known in 1941 as the coauthor, with Richard Courant, of the classic *What is Mathematics?*—an important book that, in many editions and translations, has influenced generations of mathematics students.



Herbert Robbins

By Tze Leung Lai
and David Siegmund

Robbins enlisted in the Navy in 1941 and was demobilized four years later as a lieutenant commander. His interest in probability theory and mathematical statistics began during the war and was itself something of a chance phenomenon, arising from an overheard conversation between two senior naval officers about the effect of random scatter on bomb impacts (Page 1984, pp. 8–10). Because he lacked the appropriate security clearance, Robbins was prevented from pursuing the officers' problem during his service. Nevertheless, his eventual work on this problem led to fundamental papers in the field of geometric probability, published in the *Annals of Mathematical Statistics* in 1944 and 1945.

In 1946, Harold Hotelling was setting up a department of mathematical statistics at the University of North Carolina at Chapel Hill. Having read these two papers and greatly impressed by Robbins' mathematical skills, Hotelling called Robbins and offered him the position of associate professor in the new department. Surprised by the phone call because he "knew nothing about statistics," Robbins replied at first that he could not be the person whom the caller sought, and he offered to get out the American Mathematical Society directory to find the actual Robbins. Hotelling insisted that there was no mistake, as the position was to teach "measure theory, probability, analytic methods, etc., to the department's graduate students" (Page 1984, p. 11). Robbins accepted the position and spent the next six years at Chapel Hill.

After a Guggenheim Fellowship at the Institute for Advanced Study during 1952–1953, Robbins moved from Chapel Hill to Columbia University as professor and chairman of the Department of Mathematical Statistics. With the exception of 1965–1968, spent between Minnesota, Purdue, Berkeley, and Michigan, Robbins remained at Columbia until his retirement at the age of 70, when he became Higgins Professor Emeritus of Mathematical Statistics. Robbins had been president of the Institute of Mathematical Statistics in 1965–1966, Rietz Lecturer in 1963, Wald Lecturer in 1969, and Neyman Lecturer in 1982.

We have reviewed Robbins' influential research in Lai and Sigmund (1985, 1986). The *Annals of Statistics* published a memorial section on Robbins' major contributions to statistics in April 2003, consisting of:

His publications and writings

An invited paper, “Robbins, Empirical Bayes, and Microarrays,” by Bradley Efron

Another invited paper, “Compound Decision Theory and Empirical Bayes Methods,” by Cun-hui Zhang

Siegmund’s invited paper, “Herbert Robbins and Sequential Analysis”

Lai’s invited paper, “Stochastic Approximation.”

These invited papers complemented Lai and Siegmund (1985)—in which we highlighted Robbins’ work and its impact up to 1984—especially regarding the important developments, many of which occurred in the period 1990–2000, in his specialty areas.

Below we focus on two of the research areas, mentioned previously, that Robbins created in 1951–1956, when he was relatively new to the field of statistics. These innovations, which Robbins, his students, and subsequent generations of statistical scientists then further explored, have remained vibrant to this day.

Sequential experimentation and multi-armed bandits

We begin with some highlights of the paper (Robbins 1952), in which Robbins formulated new ways to apply sequential methods to the design and analysis of experiments. He noted that because statisticians were typically “consulted, if at all, only after the experiment was over,” and also because of the “mathematical difficulty of working with anything but a fixed number of independent random variables,” until recently the size and composition of the sample had been “completely determined” before the sampling experiment. This situation changed during World War II, when a double sampling inspection was introduced at Bell Labs and Abraham Wald developed his sequential probability ratio test.

Although these developments had “freed statistics from the restriction to samples of fixed size,” Robbins argued that sequential experimentation could lead to other efficiencies as well. In particular, he introduced the k -armed bandit problem in the case $k = 2$. The name derives from an imagined slot machine with $k \geq 2$ arms. When an arm is pulled, the player wins a random reward. For each arm j , there is an unknown probability distribution Π_j , with mean μ_j , of the reward; and the player’s problem is to choose N pulls on the k arms so as to maximize $E(S_N)$, where $S_N = y_1 + \cdots + y_N$ and y_n is the reward of the n th pull.

For $k = 2$, Robbins (1952) showed how sequential sampling can give an asymptotically optimal solution such that $\lim_{N \rightarrow \infty} N^{-1} E(S_N) = \max(\mu_1, \mu_2)$. The solution is an adaptive sampling scheme that samples from the apparent leader (with the larger sample mean), except at a prescribed sparse set of times during which one has to sample from the arm that has been sampled less frequently. This simple strategy soon became a prototype of a class of stochastic adaptive control rules, called “certainty equivalence control with forcing” in the control engineering literature; see Chapter 12 of Kumar and Varaiya (1986).

The multi-armed bandit problem has become a classic problem in the fields of stochastic control, reinforcement learning, and artificial intelligence, as it addresses the dilemma between exploration (to generate information about unknown system parameters) and exploitation (to choose inputs that attempt to maximize the expected rewards from the outputs). Major advances in this problem were made by Gittins (1979), Whittle (1980), Lai and Robbins (1985), Lai (1987), and Anantharam, Varaiya, and Walrand (1987), culminating in what is now commonly known in the engineering literature as the UCB (upper confidence bound) rule with logarithmic regret; see Choi, Kim, and Lai (2018).

Whereas these advances in the 1980s represented achievements in classical bandit theory, which is referred to as “context-free multi-armed bandits” in reinforcement learning, the main theme of Choi, Kim, and Lai (2018) is “contextual” multi-armed bandit theory, which provides a mathematical framework in which to guide dynamic data-driven decisions for personalized recommendation technology. “Personalized” or “individualized” decisions are those that involve the individual’s characteristics as covariates in arm selection. An example is adaptive-enrichment designs of confirmatory clinical trials; “enrichment” refers to choosing the “right” patient subgroup for the new treatment.

In their commentary on precision medicine, Collins and Varmus (2015) argued that the time was ripe for its broad application because of recent development of large-scale biologic databases (such as the human genome sequence); powerful methods using proteomics, metabolomics, cellular arrays, and even mobile health technology to characterize patients; and computational and statistical tools for analyzing massive amounts of data. They pointed out the need for “more clinical trials with novel designs conducted in adult and pediatric patients and more reliable methods for preclinical testing.” After a review of context-free bandit theory, Choi, Kim, and Lai (2018) describe recent advances for contextual bandits, giving a complete parallel to the context-free theory. They also describe applications of this theory to personalized recommendation technology. In

particular, for web-based advertising, maximizing the click-through rate (CTR) is the corresponding contextual bandit problem, and it is an active area of research for personalized recommender systems in web services. These recommender systems strive to adapt services (such as news articles and advertisements) to individual users by exploiting both content and user information. Hence the vast potential of sequential experimentation for efficient design that Robbins (1952) envisioned in 1952 has been realized in the new millennium.

Compound decisions and empirical Bayes

By using the framework of game theory developed by John von Neumann, Wald introduced statistical decision theory in the 1940s to generalize the Neyman-Pearson optimality theory of hypothesis testing. Given that the minimax theorem was one of the major results of game theory, the minimax property also became an essential element of admissible statistical decision rules. But in 1951, Robbins showed that “the minimax solution may not be the best, since there may exist solutions which are asymptotically subminimax” for the compound statistical decision problem of testing k simple hypotheses for large k , in which the loss function for the compound problem is the sum of the component losses (Robbins 1951).

Robbins’ basic insight was that by allowing the decision rules for the individual component problems to depend on the observations from the other problems, one might be able to reduce the total risk. In 1956, he considered k similar statistical decision problems in a Bayesian context, with density function $f(\cdot; \theta_i)$ for the i th problem and with squared error loss for estimating θ_i and a common prior distribution G for the θ_i . Letting $d_G(y)$ be the Bayes decision rule when $Y_i = y$ is observed, he noted in Robbins (1956) the possibility of estimating d_G consistently from Y_1, \dots, Y_k . Specifically, the k structurally similar problems can be pooled to provide information about unspecified hyperparameters in the prior distribution, thereby yielding \hat{G} and the decision rules $d_{\hat{G}}(Y_i)$ for the independent problems. In particular, Robbins (1956) considered Poisson Y_i with mean θ_i , as in the case of the number of accidents by the i th driver in a sample of size n (in a given year) from a population of drivers, with distribution G for the accident-proneness parameter θ . In this case, the Bayes estimate (with respect to squared error loss) of θ_i , when $Y_i = y$ is observed is

$$d_G(y) = (y+1)g(y+1)/g(y), \quad y = 0, 1, \dots,$$

where

$$g(y) = \int_0^\infty \theta^y \exp(-\theta) dG(\theta)/(y!).$$

Using

$$\hat{g}(j) = n^{-1} \sum_{i=1}^n I_{\{Y_i=j\}}$$

to replace $g(j)$ above yields the empirical Bayes estimate $d_{\hat{G}}(y)$.

Compound decision theory and empirical Bayes (EB) methodology were acclaimed by Neyman (1962) as “two breakthroughs in the theory of statistical decision making.” Subsequently, Robbins (1964) provided a general framework for the empirical Bayes approach to statistical decision problems in which “the same decision problem presents itself repeatedly and independently with a fixed but unknown a priori distribution of the parameter.” In this framework, the decision problem can be hypothesis testing or estimation.

In 2003, Efron described EB as “an idea of great practical potential” that is “realized in the analysis of microarrays—a new biogenetic technology for the simultaneous measurement of thousands of gene-expression levels.” He concluded that modern science is “poised for an avalanche of empirical Bayes applications” (Efron 2003). Efron further developed the EB approach to large-scale simultaneous hypothesis testing in Efron (2004 and 2007), in which he introduced the “local false discovery rate” to connect the EB approach to FDR (false discovery rate) control in simultaneous testing of many hypotheses. Chen, Heyse, and Lai (2018) use these ideas, together with compound decision theory, to address multiplicity issues in the evaluation of medical-product safety. Zhang (2003) gave a list of references on the applications of EB methods to “numerous real-life problems,” among which were insurance issues. Lai, Su, and Sun (2014) described fundamental applications of EB methods, based on Robbins’ ideas, to actuarial science; they started with credibility theory (specifically, setting the premium of an insurance policy), and presented recent advances in evolutionary credibility theory (which included time-series effects of the underlying latent factors). Lai and Xing (2018) provide further applications of risk analytics to finance and insurance.

In his seminal 1951 paper on asymptotically subminimax compound decision theory on n independent normal observations x_i with unknown means θ_i and known common

variance 1, Robbins said: “ x_1 could be an observation on a butterfly in Ecuador, x_2 on an oyster in Maryland, x_3 the temperature of a star, and so on, all observations being taken at different times.” The compound decision problem is to determine, “on the basis of the observed values x_1, \dots, x_n , for every $i = 1, \dots, n$, whether $\theta_i = 1$ or -1 , in such a way as to minimize the expected total number of errors.” This illustrative example led many of Robbins’ contemporaries to regard compound decision theory as elegant but useless, as they could not envision such questions being asked in practice. In those days, they did not have microarray technology, flow cytometry, fMRI imaging, or electronic marketing. Robbins was far ahead of his time when he introduced compound decision theory, EB, and adaptive treatment selection in 1951, 1956, and 1952, respectively. As Zhang (2003) has noted, Robbins’ contributions, “characterized by his great originality and power” and his collaboration with successive generations of younger statistical scientists (including ourselves and Zhang), have had resounding impacts not only on statistical methodology but also throughout modern science and technology.

REFERENCES

- Anantharam, V., P. Varaiya, and J. Walrand. 1987. Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays. Part I: I.I.D. rewards. *IEEE Trans. Automat. Control* 32:968–976.
- Chen, J., J. Heyse, and T. L. Lai. 2018. *Medical Product Safety Evaluation: Biological Models and Statistical Methods*. Boca Raton, FL: Chapman and Hall/CRC.
- Choi, A. L., D. W. Kim, and T. L. Lai. 2018. *Personalized Recommendation Technology and Health Analytics*. Forthcoming from Wiley.
- Collins, F. S., and H. Varmus. 2015. Perspective: A new initiative on precision medicine. *New Engl. J. Med.* 372:793–795.
- Efron, B. 2007. Size, power, and false discovery rates. *Ann. Statist.* 35:1351–1377.
- Efron, B. 2004. Large-scale simultaneous hypothesis testing: The choice of a null hypothesis. *J. Amer. Statist. Assoc.* 99:96–104.
- Efron, B. 2003. Robbins, empirical Bayes, and microarrays. *Ann. Statist.* 31:366–378.
- Gittins, J. C. 1979. Bandit processes and dynamic allocation indices. *J. Roy. Statist. Soc. Ser. B* 41:148–177.
- Kumar, P. R., and P. P. Varaiya. 1986. *Stochastic Systems: Estimation, Identification, and Adaptive Control*. Englewood Cliffs, NJ: Prentice Hall.
- Lai, T. L. 1987. Adaptive treatment allocation and the multi-armed bandit problem. *Ann. Statist.* 15:1091–1114.
- Lai, T. L. and H. Robbins. 1985. Asymptotically efficient adaptive allocation rules. *Adv. Appl. Math.* 6:4–22.
- Lai, T. L. and D. Siegmund. 1986. The contributions of Herbert Robbins to mathematical statistics. *Statist. Sci.* 1:276–284.
- Lai, T. L. and D. Siegmund. 1985. *Herbert Robbins: Selected papers*. Berlin: Springer-Verlag.
- Lai, T. L., Y. Su, and K. H. Sun. 2014. Dynamic empirical Bayes models and their applications to longitudinal data analysis and prediction. *Statistica Sinica* 24:1505–1528.

Lai, T. L., and H. Xing. 2018. *Data Analytics and Risk Management in Finance and Insurance*. Boca Raton, FL: Chapman and Hall/CRC.

Neyman, J. 1962. Two breakthroughs in the theory of statistical decision making. (French summary) *Rev. Inst. Internat. Statist.* 30:11–27.

Page, W. 1984. An interview with Herbert Robbins. *College Math. J.* 15:2–24.

Robbins, H. 1964. The empirical Bayes approach to statistical decision problems. *Ann. Math. Statist.* 35:1–20.

Robbins, H. 1956. An empirical Bayes approach to statistics. *Proceedings of the third Berkeley symposium on mathematical statistics and probability, 1954–1955, Vol. I.* Pp. 157–163. Berkeley and Los Angeles: University of California Press.

Robbins, H. 1952. Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.* 58:527–535.

Robbins, H. 1951. Asymptotically subminimax solutions of compound statistical decision problems. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, 1950*. Pp. 131–148. Berkeley and Los Angeles: University of California Press.

Whittle, P. 1980. Multi-armed bandits and the Gittins index. *J. Roy. Statist. Soc. Ser. B* 42:143–149.

Zhang, C.-H. 2003. Compound decision theory and empirical Bayes methods: Dedicated to the memory of Herbert E. Robbins. *Ann. Statist.* 31:379–390.

SELECTED BIBLIOGRAPHY

- 1941 With Richard Courant. *What is Mathematics?* London: Oxford University Press.
- 1951 Asymptotically subminimax solutions of compound statistical decision problems. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, 1950*. Pp. 131-148. Berkeley and Los Angeles: University of California Press.
- 1952 Some aspects of the sequential design of experiments. *Bull. Amer. Math. Soc.* 58:527–535.
- 1956 An empirical Bayes approach to statistics. *Proceedings of the third Berkeley symposium on mathematical statistics and probability, 1954–1955, Vol. I*. Pp. 157–163. Berkeley and Los Angeles: University of California Press.
- 1964 The empirical Bayes approach to statistical decision problems. *Ann. Math. Statist.* 35:1–20.
- 1985 With T. L. Lai. Asymptotically efficient adaptive allocation rules. *Adv. Appl. Math.* 6:4–22.

Published since 1877, *Biographical Memoirs* are brief biographies of deceased National Academy of Sciences members, written by those who knew them or their work. These biographies provide personal and scholarly views of America's most distinguished researchers and a biographical history of U.S. science. *Biographical Memoirs* are freely available online at www.nasonline.org/memoirs.