curriculum of many colleges and universities, the next generation of texts may have a substantially different emphasis. Devaney's book is an excellent choice for professional mathematicians to read as an introduction to the subject. There are ample exercises.

JOHN GUCKENHEIMER

BULLETIN (New Series) OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 16, Number 2, April 1987

0273-0979/87 $1.00 + $.25 per page

Markov processes: *Characterization and convergence*, by Stewart N. Ethier and Thomas G. Kurtz, John Wiley & Sons, New York, Chichester, Brisbane, Toronto, and Singapore, 1986, x + 534 pp., $47.50. ISBN 0-471-08186-8

The theoretical side of mathematical probability has long been preoccupied with limit theorems; not surprisingly, since one of the natural interpretations of *probability* is as *long-run frequency*. Traditionally, limit theorems have been divided into two categories: strong limits, where some asymptotic event is asserted to occur with probability one, and weak limits, where the distributions of a sequence of random quantities are asserted to converge to a limit distribution. The prototype weak limit result is the *central limit theorem* (CLT), which says that under mild conditions the sums $S_n = X_1 + \cdots + X_n$ of independent random variables can asymptotically be approximated by Gaussian distributions. This is the key result in elementary mathematical statistics. The average height of a random sample of people is a random quantity whose exact distribution is very complicated, depending on the entire list of heights of the population; but the CLT says that the distribution of the average height of a large sample is approximately a Gaussian distribution with two parameters which depend only on the mean and standard deviation of the population heights. This illustrates the practical purpose of weak convergence theorems, to approximate complicated exact finite distributions by simpler limiting distributions. A more sophisticated example concerns neutral genetic models. Much observed genetic variation within a species (e.g., eye color in humans) confers no apparent selective advantage (i.e., is apparently "neutral"). Is it plausible that such variation is really neutral? To study this question one needs to set up a mathematical model and compare its predictions with observations. In detail, any model will be rather arbitrary and unrealistic, but one can hope that the long-term behavior of a model is insensitive to its details and instead approximates some mathematically natural process with only a few parameters.

Returning to the CLT, a pure mathematician would regard its proof as an easy exercise in Fourier analysis. Modern probabilists look at it differently. For each n, consider not only the single random variable S_n but instead the whole process $(S_m; 0 \leqslant m \leqslant n)$, which can be regarded as a random element of function space. Under the same conditions as the CLT these processes, suitably normalized, converge to the Brownian motion process $(B_t; 0 \leqslant t \leqslant 1)$. Not