Comment

Jeffrey S. Rosenthal

The papers by Geyer and Gelman-Rubin discuss the question of whether one long run or several shorter runs are preferable when using Markov chain Monte Carlo (MCMC). We feel that there are good arguments on both sides. We agree with Gelman-Rubin that multiple runs allow for more effective use of diffuse, well-chosen starting distributions. Furthermore, they can help identify situations (Section 4.8) where a Markov chain happens to get "stuck" in a low-probability region for too long. On the other hand, they require throwing away more initial samples to allow for multiple "burn-in periods." Furthermore, we share Geyer's concern (Section 4) about the difficulty of obtaining useful starting distributions in the first place. The debate will continue to rage, as these pages show. We merely wish to observe that spectacular successes have been achieved with each strategy.

Each of these papers presents useful methods for monitoring convergence of estimates to their correct values. We appreciate Geyer's careful discussion (Section 3) about estimating the variance of a function $\hat{\mu}_n$, using a version (Section 2) of the central limit theorem (CLT). (Unfortunately, Geyer does not appear to have considered the rate of convergence of this CLT; in other words, how long does the MCMC have to be run before the given normal approximation is valid?) We also appreciate the methods described by Gelman-Rubin (Section 2) for estimating convergence by using multiple runs with an overdispersed starting distribution and then monitoring the "potential scale reduction." However, despite the authors' admirable efforts to present their methods straightforwardly and directly, each method still seems to involve certain "heuristics," such as difficult choices about weight functions and bandwidths (Geyer, Section 3) or "simply discarding" certain undesirable sequences (Gelman-Rubin, Section 4.8). In general, we feel that MCMC can be used with greater confidence if it is more automated and requires less "poking around."

We agree with Geyer (Section 5) that "guarantees can only come from theoretical calculations. . . ." In this spirit, in Rosenthal (1991, 1991a), ideas related to Harris Recurrence are used to get specific, sharp theoretical bounds on time to convergence for MCMC for certain specific models, including (1991) the Gibbs sampler applied to a standard variance components model (where the required run length is shown to increase only logarithmically with increasing numbers of parameters). These bounds allow an MCMC to be run for a prespecified number of iterations, without any need for difficult or controversial monitoring techniques. Similarly, in Diaconis and Hanlon (1992), a Metropolis algorithm on the set of permutations is explicitly diagonalized, giving results on convergence rate. We feel that further theoretical results such as these could provide solid, quantitative bases for running MCMCs, thereby eliminating some of the difficulties and heuristics that are often encountered.

Jeffrey S. Rosenthal is Assistant Professor, School of Mathematics, University of Minnesota, Minneapolis, Minnesota 55455.

Comment

Bruce Schmeiser

Estimating high-dimensional volumes is analogous to estimating steady-state performance of computer, communications or manufacturing systems. The issues now attracting wide attention in the statistics community—alleviating initial bias, estimating precision and improving point-estimator quality—are long-standing problems in the operations-research literature. Based on the operations-research community's decades-old, continuing debate of one long replication versus many shorter replications, I doubt that the statistics community will soon reach a consensus. Having little hope of aiding a consensus, I only briefly discuss the number of

Bruce Schmeiser is Professor, School of Industrial Engineering, Purdue University, West Lafayette, Indiana 47907-1287.