stochastic process superimposed on a tree. In the problems addressed in the paper, the parameters of the stochastic process are of primary interest whereas the underlying tree is a nuisance parameter. For evolutionary biologists these roles are reversed. The authors make a convincing case that, in some situations, importance sampling (IS) is competitive with or superior to Markov chain Monte Carlo (MCMC) sampling as a computational tool for inference. As a developer of computational methods for phylogenetic inference based on MCMC methods (see Mau *et al.* (1999), Larget and Simon (1999) and Simon and Larget (2000)), I am intrigued to compare the techniques of this paper with our own.

Phylogenetic inferences produced by our methods are based on the post burn-in portion of a dependent sample of trees. As determined by comparing results from independent long runs, the amount of phylogenetic information in dependent samples can be similar to that from independent samples hundreds or thousands of times smaller. This paper indicates that IS may be a computationally tractable alternative.

I welcome a further expansion by the authors on the general characteristics for which an IS sampling scheme may outperform analysis by MCMC sampling. The discussion in Section 6 indicates that MCMC sampling may have an advantage when there are few constraints (meaning many trees have similar likelihoods for producing the observed data) whereas IS may have an advantage in more constrained situations.

The MCMC methods that we use propose new trees without regard to the genetics data and interact with the data only through the acceptance ratio. This has advantages and disadvantages. A sampler based on the methods of this paper may be much more efficient in some situations. However, changes to the form of the likelihood models by which genetics information evolves must entail substantial recoding for sampling methods based in part on the data whereas only minimal changes may be necessary in methods that ignore the data in proposing new trees. Direct applications of the ideas of this paper to phylogenetic inference and comparisons of the computational and statistical efficiencies would be most interesting.

Lada Markovtsova, Paul Marjoram and Simon Tavaré (*University of Southern California, Los Angeles*)
The authors have presented a most inspirational paper on computational methods for the coalescent. Their suggestion that one might combine Markov chain Monte Carlo (MCMC) and importance sampling (IS) approaches is particularly intriguing. As an example they suggest using the IS proposal distribution to update a random amount of the upper part of the tree. We have previously experimented with a pure MCMC approach in which the proposal distribution worked in a similar manner, replacing a random amount of the top of the tree with a *random* topology. Perhaps not surprisingly, such changes may have a very low acceptance rate. Updates that replaced the tree from a relatively high point were accepted frequently, but when the update involved a large part of the topology the acceptance rate was very low. One can improve this naïve approach by alternating such proposals with updates that replace a random amount of the lower part of the tree (again we generated a random topology), but it is not clear how this would be accomplished in the IS framework that the authors suggest. In our approach such a scheme was very inefficient, particularly for large sample sizes, but there is reasonable hope that the improved efficiency of the proposal distribution given in this paper might circumvent the problems that we experienced.

Our experience with implementations of a fully Bayesian approach for deoxyribonucleic acid sequence data (e.g. Markovtsova *et al.* (2000a)) supports the authors' observation that Bayesian methods provide useful computational tools even when one's interest is in maximum likelihood estimation. Checking the adequacy of the estimated likelihood near a maximum can be accomplished by using different priors. The apparent simplicity of estimating relative likelihoods from marginals of the posterior distribution and the prior seems difficult to pass up. Do the authors have any thoughts on when this approach is likely to be misleading? We also note that posterior trees and rates can be used in a bootstrap approach for checking model adequacy (e.g. Markovtsova *et al.* (2000b)).

Testing IS and MCMC implementations is notoriously difficult; the development of test examples seems worthwhile. Another useful approach is to compare results with those generated by simpler schemes like the rejection methods. For instance, we have found this helpful in checking tree topology updates.

Bob Mau (*University of Wisconsin, Madison*)
Problems in population genetics are often intractable with analytical methods once the sample size approaches 10. The reason is simple: the form of the likelihood, conditioned on the genealogy, is