Journal of Applied Statistics, Vol. 29, Nos. 1-4, 2002, 239-240

Discussion comments on 'Prior distributions for stratified capture-recapture models'

C. J. SCHWARZ, *Department of Statistics and Actuarial Science, Simon Fraser University, Canada*

This paper discusses a Bayesian framework for the analysis of movement models in the Arnason-Schwarz framework. The mechanics of such an analysis are known, in general, to many researchers—prior information about the unknown parameters is combined with information from actual data to give a posterior distribution. In the past, the technical difficulties of actually finding the posterior distribution in closed form has limited the use of Bayesian methods. However, computer intensive methods such as the Gibb's sampler, have removed this difficult step.

This paper presents a number of interesting ideas.

(1) Separation of the experiment into two processes—a movement process and an observation process. The observation process reveals parts of this movement process. This separation is similar to that of state-space modelling where a system is broken into two elements. The first element is the unknown, underlying movement of the system from state to state. The second element is the observation process where selected portions of the state-space are revealed to the experimenter. In the usual state-space formulation, the observation process has observation error—this is absent in the formulation by Dupuis—observations on the state of an animal are without error. The analysis of state-space models typically uses the Kalman filter approach. How is the approach by Dupuis related to the state space model?

(2) A common difficulty in using the Gibb's sampler is being able to write out the conditional distribution of each parameter in turn. This is the result of a 'difficult' looking likelihood function. I like Dupuis's 'missing data' approach, where both the missing data and the parameters form part of the Gibb's sampling process. This approach could make much of capture-recapture modelling easier in a Bayesian approach.

Correspondence: C. J. Schwarz, Department of Statistics and Actuarial Science, Simon Fraser University, Burnaby, British Columbia, Canada V5A 156. E-mail: cschwarz@stat.sfu.ca

ISSN 0266-4763 print; 1360-0532 online/02/010239-02
DOI: 10.1080/02664760120108647

(3) The role of parameter non-identifiability in a Bayesian method has always puzzled me. Because of the prior distributions, no parameter is 'non-identifiable'—at the very least, if the data contains no information about a parameter, its posterior should be identical to the prior. However, if two parameters are confounded, exactly what happens? There is some information about the pair of parameters, but it is not clear how this is parcelled out to the two parameters. Is there a way to detect these non-identifiable parameters in a Bayesian context? This could provide important model diagnostics for the experimenter.

(4) Another objection often raised to Bayesian methods is 'how is the prior determined'? Presumably expert knowledge about a parameter does not exist in the form of parameters of a Beta distribution. Dupuis suggests that the parameters for a single parameter be found by elucidating a '95% confidence interval' *a priori*, by, for example, asking experts for their best guess for the value of the parameter and perhaps a lower and upper bound for a range of plausible values. However, I am not clear on how to construct a prior for a multi-parameter distributions, e.g. the movement rates. Dupuis's method appears to have only a single 'degree of freedom' after the best guesses are determined and can be constructed from a single '95% confidence interval' about one component of the movement. Presumably, you would get different values depending on which component was chosen. Perhaps one strategy would to examine each component in turn and build a prior based on the component that gave the worse value of l.

(5) Dupuis gives an example where the standard confidence interval from MARK is so wide as to be useless, but the Bayesian intervals are much tighter. Presumably, this indicates that the data contain little information about the parameter and much of posterior is based on the prior. Is there a way of indicating the relative amount of information in the posterior from each source? Earlier in the paper, Dupuis notes that the posterior estimate is a weighted average of the MLE and the prior mean. Can this information be routinely computed for each parameter. Then, upon seeing that a particular parameter had a weight of 90% toward the prior would tell the experimenter that the particular parameter is still essentially unknown. This would also give guidance to the experimenter on which parameters have to be investigated in a sensitivity analysis to different priors.