Bayes or likelihood: Convergence or Divergence
Organizer: D.A.S. Fraser

What model information is appropriate for the Bayesian paradigm?
Don Fraser
dfraser@utstat.toronto.edu
Department of Statistics
University of Toronto
100 St George St
Toronto ON M5S 3G3 Canada

In its elemental form the Bayesian paradigm modulates an observed Fisher likelihood to give a modified likelihood that is treated as a relative density on the parameter space; the modulating factor is called a prior. We address the origins of this prior, in particular the use of information coming from characteristics of the model.

In some contexts there may be an acknowledged random source for the true parameter value; if the applicable probabilities for this true value are available, with or without a hyperparameter, we call the prior objective, a more restrictive usage than usual.

In other cases without an acknowledged random source, the prior can represent the personal views of an individual analyst; we then call the prior subjective. Or the prior can represent relevant characteristics of the model and how it provided information concerning the actual parameter value. Our focus here is on this later default usage of the Bayesian paradigm, that is, let's use a powerful likelihood-based paradigm and get on with the job.

Recent analysis (Fraser & Yuan, 2004) leads to neutral priors, based on the following desiderata: If the model is location relative to $\beta(\theta)$, then the neutral prior is given by $\pi(\theta)d\theta = d\beta(\theta)$; if the model is an identifiable mixture of models, then the neutral prior is that for the identfied component model, as available; if the model has at some order an approximate model with a neutral prior, then that prior is the neutral prior at the given order.

We show in wide generality that neutral priors are available, that they produce Bayesian survivor values equal to the frequentist p-values at the given order, and that if an individual has a personal prior it can be placed in juxtaposition with the default survivor value, for use where deemed appropriate.

Elimination of Nuisance Parameters by Integration
Tom Severini
severini@northwestern.edu
Department of Statistics
Northwestern University
2006 Sheridan Road
Evanston, IL 60208 USA

Consider a model with parameter θ and component parameter of interest $\psi = \psi(\theta)$ and consider the problem of constructing a likelihood function for ψ. Many non-Bayesian methods of constructing such a likelihood have been proposed, leading to marginal and conditional likelihoods, profile likelihoods, and modified versions of the profile likelihood, among others. From the Bayesian point-of-view, construction of a likeihood for ψ is easily achieved by 'integrating out' the nuisance parameter with respect to a given prior distribution, although selection of this prior distribution may be difficult. In this talk, I argue that elimination of nuisance parameters by integration is generally preferable to non-Bayesian methods; the use of an integrated likelihood function in non-Bayesian inference will also be discussed.

Developing p-values: A Bayesian-frequentist convergence
Judith Rousseau
rousseau@ceremade.dauphine.fr
6 ter Bd Orloff
77 300 Fontainebleau France

Various p-values for a composite null hypothesis have had extensive attention in the Bayesian literature, with some preference shown for two versions designated p_{ppost} and p_{cpred} (Bayarri and Berger, 2001); and it has been indicated that certain candidate p-values can be upgraded to the preceding preferred p-values by the parametric bootstrap. Also recent likelihood theory gives a factorization of a statistical model into a marginal density for a full dimensional ancillary and a conditional density for the maximum likelihood variable. Using these results we construct a frequentist p-value, namely p_{anc}. We then prove that the Bayesian p-value, p_{cpred} based on the maximum likelihood estimator and p_{anc} are equivalent under the null hypothesis to third order. We also prove that under regularity conditions on the model and hardly any conditions on the test statistic they are second order uniform under the null hypothesis and under more stringent conditions on the test statistic, such as asymptotic normality, we obtain that they are third order uniform. We also compare these two p-values with the iterative bootstrap p-values and prove that they are equivalent, to third order, to the fourth iterative bootstrap p-value.