Bayesian Analysis (2010) **5**, Number 4, pp. 671–676

Comment on Article by Vernon et al.

David Poole*

1 Introduction

The authors are to be congratulated on their presentation of a detailed and comprehensive case study. It is clear that this effort represents a large amount of work by both statisticians and cosmologists. The important topic of statistical inference from complex non-linear deterministic simulation models has received a lot of attention over the last 20 years, and this paper attempts to deal with the many sources of uncertainty in a formal and novel manner, while providing an interesting description of galaxy formation.

I would like to focus this discussion on just one aspect of the paper, that of Bayesian inference for parameters linked by a deterministic simulator (Galform in this case).

2 Pragmatic compromise in Bayesian analyses

The authors employ their Bayes linear approach in this application. Based solely on a specification of means, variances and covariances, it is simpler to implement in complicated problems where a fully Bayesian analysis may be intractable. They describe the Bayes linear approach as a "pragmatic compromise" to a fully Bayesian solution, given a) the difficulty in eliciting a full joint prior probability distribution, and b) the technical challenges in implementing a massive MCMC or similar analysis. These are very valid points, and I am reminded of my own experience with an application in marine mammal assessment that shares some similarities with the Galform problem.

2.1 A population dynamics model for bowhead whales

Population growth for marine mammals is often modeled using a deterministic simulation model, usually a highly non-linear set of differential equations. Given a set of inputs to this model (fecundity, survival rates, pre-exploitation stock size, maximum sustainable yield, etc.) and a known commercial catch history, the population is projected through time from the start of the commercial fishery to the present time where a series of abundance estimates from surveys is typically available. Similar to Galform, the challenge is to find combinations of the inputs that are scientifically plausible, do not lead to extinction, and which produce population trajectories (over time) that match the current data on abundance and rates of increase.

In the 1990's the International Whaling Commission used a model named BALEEN II for assessment of the Bering-Chukchi-Beaufort Seas stock of bowhead whales, *Balaena mysticetus*. Further background and history are given in Raftery et al. (1995), who also

*AT&T Labs-Research, 180 Park Avenue, Florham Park, NJ, mailto:poole@research.att.com

 DOI:10.1214/10-BA524A

presented an initial Bayesian approach to inference for the inputs, outputs (and functions thereof) of the BALEEN II model. The method was updated and slightly modified in Poole and Raftery (2000), where it was given the name *Bayesian melding*.

The essence of Bayesian melding is as follows: a prior distribution over the model inputs θ is transformed to the output space ϕ. The resulting induced prior for ϕ is combined with an existing independent prior for ϕ using logarithmic pooling of the two densities. This pooled prior on the outputs is then transformed backwards through the model to produce a pooled prior over the inputs that hopefully reflects subtle correlations and dependencies absent in our initial prior. This adjusted prior, which in some sense can be considered an analog of the adjusted expectations used in the Bayes linear method, is then combined with a likelihood in the normal way to achieve posterior inference. The approach was initially developed specifically for the whale application, but it has since been applied to models of tree growth (Radtke et al. 2002), disease transmission (Spear et al. 2002), and soil loss (Falk et al. 2010). An extension to stochastic simulation models is proposed in Sevcikova et al. (2007).

2.2 A simple example

To illustrate, reproduced here is Example 1 of Poole and Raftery (2000). This is a trivial deterministic transform $M : Z = Y/X$, where X and Y are the two inputs to the model while Z is the single output. Hence, $\theta = (X, Y)$ and $\phi = Z$. Assume that we are able to specify independent marginal prior distributions for each quantity as follows: $X \sim U[2, 4]$, $Y \sim U[6, 9]$, and $Z \sim U[0, 5]$. It follows that a joint prior distribution for the inputs, which we call q_1, is given by

$$q_1(x, y) = \frac{1}{6}, \text{ for } 2 < x < 4, \quad 6 < y < 9$$

and for the output, labelled q_2, it is

$$q_2(z) = \frac{1}{5}, \text{ for } 0 < z < 5.$$

Using the Bayesian melding machinery the general form of the pooled prior over the inputs θ is given by

$$\tilde{q}^{[\theta]}(\theta) \propto q_1(\theta) \left(\frac{q_2(M(\theta))}{q_1^*(M(\theta))} \right)^{1-\alpha} \tag{1}$$

where $q_1^*(\cdot)$ is the distribution of the output ϕ induced by the deterministic model, and $0 \leq \alpha \leq 1$ is the pooling weight. Note that the pooled prior $\tilde{q}^{[\theta]}(\cdot)$ is simply the original prior, $q_1(\cdot)$, weighted by the ratio of two possibly lower-dimensional densities in ϕ-space, $q_2(\cdot)$ and $q_1^*(\cdot)$, evaluated at $M(\theta)$, where M is the model function. The magnitude of the weight is dictated by the value of α, and the ratio is never evaluated at arbitrary values of ϕ, only at $M(\theta)$ for a given value of θ. The interested reader is referred to Poole and Raftery (2000) for more details on the derivation of equation (1). In particular, it is shown that, conditional on the output ϕ, the distribution in equation (1) is a choice that minimizes Kullback-Leibler distance to the original prior $q_1(\cdot)$.

In the simple example above we can use standard change-of-variable methods to obtain an analytic form of $q_1^*(\cdot)$. This can be plugged into equation (1), with α chosen to be 0.5 here, to yield

$$\tilde{q}^{[\theta]}(x,y) = \begin{cases} \frac{ky}{2\sqrt{15(4y^2-9x^2)}} & : y < 2.25x \\ \frac{k\sqrt{3}y}{45x} & : 2.25x < y < 3x \\ \frac{ky}{\sqrt{15(81x^2-4y^2)}} & : y > 3x \end{cases}$$

where $k \approx 1.4$ is the appropriate normalizing constant. This pooled prior on $\theta = (X, Y)$ is shown graphically in Figure 1, where we can see the that the original flat prior on the inputs (a product of 2 independent marginals) has been modified to reflect information derived from the model output. The three superimposed solid lines represent sets of points in the (X,Y)-plane that map to three single points in Z-space. In this case the prior for Z specifies that each z-value is equally likely, and the pooled input prior accounts for this by increasing the density where such lines are shorter.

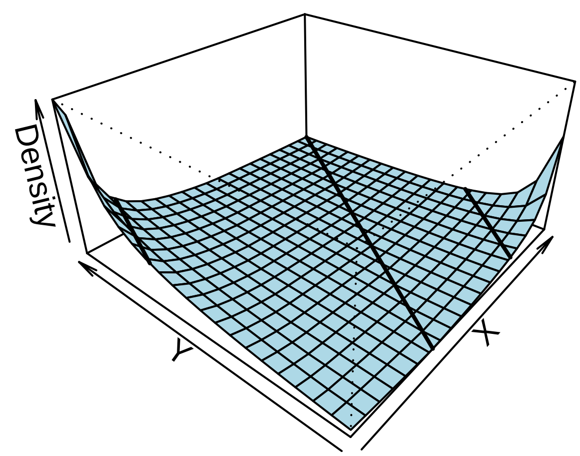

Figure 1: The pooled joint prior distribution of X and Y in the simple example. The shape reflects prior information about the output Z projected through the deterministic relationship.

Naturally, in any real application, a functional form of the induced prior $q_1^*(\cdot)$ will be impossible to obtain, so we approximate using Monte Carlo density estimation and then employ sampling-importance-resampling methods to sample from the posterior distribution. The **key** point is this: we use information implicit in the deterministic model, in which we presumably have some belief, in conjunction with marginally-specified prior distributions for inputs and outputs, to induce a coherent adjusted joint prior distribution over all the inputs to the model. This joint prior, when combined with a likelihood for observed data, results in a joint posterior over inputs, and then via the model transform, in a similarly coherent joint posterior distribution for the outputs.

It is not at all clear that an approach like this would be suited to a problem as complex as the Galform application. For instance, density estimation in higher dimensions can be difficult, the selection of α is an issue, and the basic method assumes that the deterministic model is known exactly. These would all be problematic in this situation, and are discussed generally in Poole and Raftery (2000). That said, the melding approach does provide a fully Bayesian solution with full posterior inference available from the Monte Carlo samples. It also helps to cut down the input space in a manner similar to that which the authors describe in their application. Indeed, in the whale case the melded prior was key in eliminating a large region of the input space before combination with the likelihood function for abundance data.

I view Bayesian melding as a pragmatic compromise brought on, in our case, by an inability to specify anything other than marginal prior probabilities independent of the simulation model. If we could reasonably elicit complex joint priors, we would do so. Thus in effect we have a partial specification of our prior beliefs. In this sense, it mirrors the partial specifications of the Bayes linear approach, and can be considered a compromise to a more complicated analysis.

2.3 What does one lose?

Approaches such as these beg the further question: what does one lose by adopting the compromises inherent in melding or the Bayes linear method? The preface to Goldstein and Wooff (2007) suggests that the Bayes linear method achieves "90% of the answer for 10% of the effort", but I wonder how verifiable that really is in an application such as Galform. Is the result really 90% as good as it would be under full Bayes? How do we know it is not only 50%? Expressed differently, if a full Bayesian solution were hypothetically possible, would we learn a lot more about galaxy formation from the observed data? Could one get much sharper confidence statements about key parameters? It seems that cautionary statements of this type should perhaps accompany the declaration of results when, as here, one is forced to stray from the fully Bayesian path. In the case of Bayesian melding, while it provides fully Bayesian inference and forces a joint prior elicitation, it requires additional subjective judgements about pooling and the fixed form of the simulation model. It is hard to quantify the effects of this extra subjectivity.

3 Conclusion

A discussion of this type could carry on at length and would likely devolve into another foundational debate about what it means to be a subjective Bayesian, or whether the various Bayesian approaches are preferable to other methods. That is not the aim here.

All that remains for me is to congratulate the authors again on an impressive and detailed piece of work, one that contains a great deal of novelty tailored to the complexities of the application. My hope is that it will stimulate interest and further discussion among the readers of Bayesian Analysis.

References

Falk, M. G., Denham, R. J., and Mengersen, K. L. (2010). "Estimating uncertainty in the revised universal soil loss equation via Bayesian melding." *Journal of Agricultural, Biological and Environmental Statistics*, 15: 20–37. 672

Goldstein, M. and Wooff, D. A. (2007). *Bayes Linear Statistics: Theory and Methods*. Chichester: Wiley. 674

Poole, D. and Raftery, A. E. (2000). "Inference for deterministic simulation models: the Bayesian melding approach." *Journal of the American Statistical Association*, 95: 1244–1255. 672, 674

Radtke, P. J., Burk, T. E., and Bolstad, P. V. (2002). "Bayesian melding of a forest ecosystem model with correlated inputs." *Forest Science*, 48: 701–711. 672

Raftery, A. E., Givens, G. H., and Zeh, J. E. (1995). "Inference From a Deterministic Population Dynamics Model for Bowhead Whales (with discussion)." *Journal of the American Statistical Association*, 90: 402–430. 671

Sevcikova, H., Raftery, A. E., and Waddell, P. A. (2007). "Assessing uncertainty in urban simulations using Bayesian melding." *Transportation Research Part B-Methodological*, 41: 652–669. 672

Spear, R. C., Hubbard, A., Liang, S., and Seto, E. (2002). "Disease transmission models for public health decision making: Toward an approach for designing intervention strategies for *Schistosomiasis japonica*." *Environmental Health Perspectives*, 110: 907–915. 672