

# Bayesian model choice via mixture distributions with application to epidemics and population process models

Philip D. O'Neill<sup>1</sup>,

Theodore Kypraios<sup>1</sup>

<sup>1</sup>School of Mathematical Sciences, University of Nottingham, UK

## Abstract

We consider Bayesian model choice for the setting where the observed data are partially observed realisations of a stochastic population process. A new method for computing Bayes factors is described which avoids the need to use reversible jump approaches. The key idea is to perform inference for a hypermodel in which the competing models are components of a mixture distribution. The method itself has fairly general applicability. The methods are illustrated using simple population process models and stochastic epidemics.

*Keywords:* Bayes factors; Epidemic models; Markov chain Monte Carlo methods; Model choice

## 1 Introduction

Consider an observed sequence of event times, each event being of the same type, and suppose we wish to assess whether a homogeneous Poisson process or an alternative non-homogeneous Poisson process best fits the observations. Alternatively, suppose we have case-detection times in an outbreak of infectious disease,

and wish to know which of two possible SIR (susceptible-infective-removed) disease transmission models is most plausible as a model for how the data were generated, assuming that removals correspond to case-detections. Both of these examples are special cases of a generic situation in which we wish to assess which of a number of proposed point-process models best fits the data to hand. In the first example there is one type of event, and all events are observed. In the second example there are two event types (infections and removals) but only the latter are observed. In both examples two models are compared, but in general we may have more models of interest.

In a Bayesian framework, questions of model choice can be addressed using Bayes factors, which quantify the relative likelihood of any two models given the data and within-model prior distributions. Bayes factors can suffer from two practical drawbacks, namely (i) they can be difficult to compute, and (ii) they can be highly sensitive to the choice of within-model prior distributions, and in particular apparently natural choices can give misleading results. Here our focus is towards addressing the first difficulty, but in respect of the second we briefly remark that alternative methods of Bayesian model assessment have their own difficulties in the setting we consider. For example, neither the Deviance Information Criterion (DIC) nor Bayesian Information Criterion (BIC) appear entirely natural for settings where the data are typically far from being independent observations, as is the case when the data are realisations of a stochastic process. For problems involving missing data, such as the epidemic example above, it is not even clear how suitable information criteria should best be defined (Celeux et al. (2006) give nine candidates, for instance). Finally, methods involving a comparison between the observed data and what the fitted model would predict typically involve a subjective judgement as to precisely what should be compared, and how.

In all but the simplest cases, Bayes factors must be evaluated numerically. For many problems, this can be achieved via reversible jump Markov chain Monte Carlo (RJMCMC) methods (Green, 1995). To be precise, consider two models  $M_1$  and  $M_2$  with parameters  $\theta_1$  and  $\theta_2$ , respectively, where  $\theta_j \in \Theta_j$ .

Define  $k \in \{1, 2\}$  to be a model indicator which specifies the model under consideration. RJMCMC methods proceed by defining a Markov chain with state space  $\{1\} \times \Theta_1 \cup \{2\} \times \Theta_2$  such that the proportion of time for which  $k = j$  converges to the posterior model probability  $P(M_j|x)$ , where  $x$  denotes the observed data. Given model prior probabilities  $P(M_j)$ , the Bayes factor in favour of model 1 is given by the expression  $P(M_2)P(M_1|x)/P(M_1)P(M_2|x)$ , which can be estimated from the RJMCMC output.

The main practical challenge in implementing RJMCMC algorithms is constructing efficient between-model proposal distributions, i.e. defining how the Markov chain jumps between the different components of the union of model parameter spaces. Although there have been theoretical advances which address this issue (Brooks et al., 2003), for many problems it remains a case of trial and error. In this paper we propose a method which goes some way to removing this complication. The key idea is to consider a hypermodel which is itself a mixture model whose components are the two or more competing models of interest. An MCMC algorithm can then be defined on the product space of all model parameters and mixture probabilities. Bayes factors for the models can be expressed in terms of the posterior means of the mixture probabilities, and thus estimated from the MCMC output.

Before proceeding to the details, we consider the general context. First, defining a Markov chain on a product (rather than union) of model-parameter spaces is the approach pioneered by Carlin and Chib (1995), and further developed to more general settings (Green and O'Hagan (1998), Dellaportas et al. (2002), Godsill (2001)). This approach, as for RJMCMC, involves defining a probability distribution over the set of possible models, and introduces a parameter which indicates which model is chosen. In our setting there is no chosen model as such, but instead a mixture of all possible models. The product-space approach also relies on defining so-called pseudo-priors for the within-model parameters, upon which algorithm efficiency is crucially dependent, and this can be difficult in practice. Our methods do not involve the need to introduce such pseudo-priors, although for some missing data problems we need to specify

similar prior distributions for the missing data.

Second, computational methods for the Bayesian analysis of mixture models are well-established, both when the number of components in the mixture is known (Diebolt and Robert, 1994) and when it is not (Richardson and Green, 1997). The typical situation under consideration is one in which the data are assumed to comprise independent and identically distributed observations from the proposed mixture distribution(s). In contrast, we consider the case where there is only one datum, but it consists of the realisation of a stochastic process, either fully or partially observed.

The paper is structured as follows. Section 2 contains general theory which describes the inference framework in detail, and computational matters are described in Section 3. Section 4 contains examples and we conclude with discussion in Section 5.

## 2 General Theory

In this section we introduce the underlying framework of interest. For ease of exposition we adopt the usual abuse of notation and terminology in which ‘a density  $\pi(\theta)$ ’ can refer to both the density function  $\pi$  of a random variable  $\theta$ , or the same function evaluated at a typical point  $\theta$ .

### 2.1 Mixture model with no missing data

Suppose we observe data  $x$ , and wish to consider  $n$  competing models  $M_1, \dots, M_n$ . For  $i = 1, \dots, n$  denote the probability density of  $x$  under model  $i$  by  $\pi_i(x|\theta_i)$ , where  $\theta_i$  denotes the vector of within-model parameters, and set  $\theta = (\theta_1, \dots, \theta_n)$ . We assume that all the  $\pi_i(x|\theta_i)$  are densities with respect to the same common reference measure. Define a mixture model by

$$\pi(x|\alpha, \theta) = \sum_{i=1}^n \alpha_i \pi_i(x|\theta_i), \quad (1)$$

where  $\alpha = (\alpha_1, \dots, \alpha_n)$  satisfies  $\sum_{i=1}^n \alpha_i = 1$  and  $\alpha_i \geq 0$  for  $i = 1, \dots, n$ .

## 2.2 Mixture model with missing data

In our setting, the data  $x$  may be a partial observation of a stochastic process. In consequence,  $\pi_i(x|\theta_i)$  in (1) may be intractable, meaning that it cannot be analytically or numerically evaluated in an efficient manner. We adopt data augmentation to overcome this problem, as follows. Let  $y = (y_1, \dots, y_N)$  be a vector comprising different kinds of ‘missing data’, and for  $i = 1, \dots, n$  let  $\mathcal{I}(i) \subseteq \{1, \dots, N\}$  and define  $y_{\mathcal{I}(i)}$  as the vector with components  $y_j$ ,  $j \in \mathcal{I}(i)$ . Thus  $y_{\mathcal{I}(i)}$  denotes the missing data for model  $i$ , and in practice it is chosen to make the augmented probability density  $\pi_i(x, y_{\mathcal{I}(i)}|\theta_i)$  tractable. If model  $i$  does not require missing data, then  $y_{\mathcal{I}(i)}$  is null. Note that this formulation allows different models to share common elements of missing data. Conversely, if each model has its own missing data then we simply set  $\mathcal{I}(i) = i$  for  $i = 1, \dots, n$ .

In order to define a mixture model using missing data, it is necessary to introduce additional terms so that each component of the mixture is a probability density function on the possible values of  $x$  and  $y$ . To this end, we assume that there exist tractable probability densities  $\pi_i(y_{-\mathcal{I}(i)}|x, y_{\mathcal{I}(i)}, \theta)$ , where  $y_{-\mathcal{I}(i)}$  denotes the vector with components  $y_j$ ,  $j \notin \mathcal{I}(i)$ . If the latter set is empty then we set  $\pi_i(y_{-\mathcal{I}(i)}|x, y_{\mathcal{I}(i)}, \theta) = 1$ . We refer to the  $\pi_i(y_{-\mathcal{I}(i)}|x, y_{\mathcal{I}(i)}, \theta)$  terms as *missing data prior densities*. In practice, they need not explicitly depend on any of  $x$ ,  $y_{\mathcal{I}(i)}$  or  $\theta$ , depending on the application at hand.

Define an augmented mixture model by

$$\pi(x, y|\alpha, \theta) = \sum_{i=1}^n \alpha_i \pi_i(x, y_{\mathcal{I}(i)}|\theta_i) \pi_i(y_{-\mathcal{I}(i)}|x, y_{\mathcal{I}(i)}, \theta). \quad (2)$$

Here we assume that each  $\pi_i(x, y_{\mathcal{I}(i)}|\theta_i) \pi_i(y_{-\mathcal{I}(i)}|x, y_{\mathcal{I}(i)}, \theta)$  term in the sum in (2) is a probability density with respect to a common reference measure, from which it follows that  $\pi(x, y|\alpha, \theta)$  is also a probability density.

## 2.3 Bayes Factors

We now show how Bayes factors can be computed directly from certain summaries of the posterior distribution of  $\alpha$  given the data  $x$ . For simplicity we

assume that  $\alpha, \theta_1, \dots, \theta_n$  are mutually independent *a priori*. By Bayes' Theorem,

$$\pi(\alpha|x) = \frac{\pi(x|\alpha)\pi(\alpha)}{\pi(x)} = \frac{\pi(\alpha) \sum_{i=1}^n \alpha_i m_i(x)}{\pi(x)},$$

where  $\pi(\alpha)$  denotes the prior density of  $\alpha$  and, for  $i = 1, \dots, n$ ,

$$m_i(x) = \int \pi_i(x, y_{\mathcal{I}(i)} | \theta_i) \pi_i(\theta_i) d\theta_i dy_{\mathcal{I}(i)},$$

where  $\pi_i(\theta_i)$  is the within-model prior density of  $\theta_i$ . Note also that

$$1 = \int \pi(\alpha|x) d\alpha = \pi(x)^{-1} \sum_{i=1}^n E[\alpha_i] m_i(x),$$

whence

$$\pi(x) = \sum_{i=1}^n E[\alpha_i] m_i(x). \quad (3)$$

Now for  $i \neq j$ , the Bayes factor in favour of  $M_i$  relative to  $M_j$  is defined to be  $B_{ij} = B_{ij}(x) = m_i(x)/m_j(x)$ . However,

$$\begin{aligned} E[\alpha_i|x] &= \int \alpha_i \pi(\alpha|x) d\alpha \\ &= \pi(x)^{-1} \int \alpha_i \left( \sum_{j=1}^n \alpha_j m_j(x) \right) \pi(\alpha) d\alpha \\ &= \pi(x)^{-1} \sum_{j=1}^n E[\alpha_i \alpha_j] m_j(x), \end{aligned}$$

which combined with (3) yields that

$$E[\alpha_i|x] = \frac{\sum_{j=1}^n E[\alpha_i \alpha_j] m_j(x)}{\sum_{j=1}^n E[\alpha_j] m_j(x)}, \quad i = 1, \dots, n. \quad (4)$$

Next, fix  $k \in \{1, \dots, n\}$ . Dividing the numerator and denominator of the fraction in (4) by  $m_k(x)$  and rearranging we obtain

$$\sum_{j=1}^n (E[\alpha_j] E[\alpha_i|x] - E[\alpha_i \alpha_j]) B_{jk}(x) = 0, \quad i = 1, \dots, n. \quad (5)$$

It remains to solve equations (5) to find  $B_{jk}(x)$ ,  $j = 1, \dots, n$ . Define  $A$  as the  $n \times n$  matrix with elements

$$A_{ij} = E[\alpha_i|x] E[\alpha_j] - E[\alpha_i \alpha_j], \quad 1 \leq i, j \leq n.$$

Note that  $A$  depends on  $x$ , although we suppress this dependence in our notation. Define  $\tilde{A}_{-k}$  as the  $(n-1) \times (n-1)$  matrix formed by removing the  $k$ th row and  $k$ th column of  $A$ . Similarly for  $j \neq k$  define  $\tilde{A}_{-jk}$  as the  $(n-1) \times (n-1)$  matrix formed from  $\tilde{A}_{-k}$  by replacing the elements  $A_{ij}$  with  $-A_{ik}$ ,  $i = 1, \dots, n$ ,  $i \neq k$ .

*Lemma 1.* (a) If  $\det \tilde{A}_{-k} \neq 0$  then

$$B_{jk}(x) = \frac{\det \tilde{A}_{-jk}}{\det \tilde{A}_{-k}}. \quad (6)$$

(b) Suppose that  $0 < m_i(x) < \infty$  for  $i = 1, \dots, n$ . Then if either (i)  $n = 2$  and  $0 < E[\alpha_1] < 1$ , or (ii)  $\alpha$  has a Dirichlet prior distribution,  $\mathcal{D}(p_1, \dots, p_n)$ , then

$$B_{jk}(x) = \frac{A_{jk}}{A_{kj}}.$$

The proof of Lemma 1 is in the Appendix. The result shows that the required Bayes factors can be expressed in terms of the prior distribution summaries  $E[\alpha_i]$  and  $E[\alpha_i \alpha_j]$  and the posterior means  $E[\alpha_i | x]$ ,  $i, j = 1, \dots, n$ .

The condition on the determinant of  $\tilde{A}_{-k}$  in Lemma 1(a) is not vacuous in general, as illustrated by the somewhat pathological case where  $\alpha_i$  has a point mass prior for all  $i = 1, \dots, n$ . Then for all  $1 \leq i, j \leq n$ ,  $E[\alpha_i | x] = E[\alpha_i]$  and  $E[\alpha_i \alpha_j] = E[\alpha_i]E[\alpha_j]$ , whence  $A_{ij} = 0$  and (5) cannot be solved to find the Bayes factors.

At first sight the need for a Dirichlet prior on  $\alpha$  to yield simple evaluation of the Bayes factors via Lemma 1 may appear restrictive. We make three remarks. First, the mixture construction is itself introduced solely as a tool for evaluation of Bayes factors, and so there is no particular need to assign an arbitrary prior distribution to  $\alpha$ . Second, in practice a Dirichlet prior is both straightforward to use and flexible enough for computational purposes as described below. Third, it may well be that (6) holds for arbitrary prior distributions on  $\alpha$ , subject to mild constraints which imply that  $\det \tilde{A}_{-k} \neq 0$ , but this does not appear straightforward to prove.

Finally, the simple form for the Bayes factor in Lemma 1 (b) does not appear to be true in general; an example for  $n = 3$  can be found in the Appendix.

## 2.4 Two competing models

We give special attention to the case  $n = 2$  since this is of practical importance. Here we have  $\alpha = (\alpha_1, 1 - \alpha_1)$  and Lemma 1 yields that

$$B_{12} = \frac{E[\alpha_1] - E[\alpha_1^2] - E[\alpha_1|x](1 - E[\alpha_1])}{E[\alpha_1]E[\alpha_1|x] - E[\alpha_1^2]}.$$

It follows that

$$\frac{E[\alpha_1] - E[\alpha_1^2]}{1 - E[\alpha_1]} \leq E[\alpha_1|x] \leq \frac{E[\alpha_1^2]}{E[\alpha_1]},$$

with the upper and lower bounds corresponding to Bayes factors entirely in favour of models 1 and 2, respectively. A practical consequence is that any numerical estimate of  $E[\alpha_1|x]$  lying outside these bounds must be incorrect.

Under the further assumption that  $\pi(\alpha)$  is a uniform density, so that  $\alpha_1 \sim U(0, 1)$ , we obtain

$$B_{12} = \frac{3E[\alpha_1|x] - 1}{2 - 3E[\alpha_1|x]},$$

$$\pi(\alpha_1|x) \propto \alpha_1 m_1(x) + (1 - \alpha_1) m_2(x),$$

$E[\alpha_1|x] = (2m_1(x) + m_2(x))/(3(m_1(x) + m_2(x)))$  and  $1/3 \leq E[\alpha_1|x] \leq 2/3$ .

Finally, if  $\alpha$  is assigned a Dirichlet prior distribution, bounds for  $E[\alpha_i|x]$  for any value of  $n$  can be obtained. Full details can be found in the proof of Lemma 1 in the Appendix.

## 3 Computation

We now describe how to use the mixture framework in practice, specifically via MCMC methods. Our objective is to sample from the target density

$$\pi(\alpha, \theta, y|x) \propto \pi(x, y|\alpha, \theta)\pi(\alpha)\pi(\theta), \quad (7)$$

and the first issue is that of assigning any missing data prior density terms in  $\pi(x, y|\alpha, \theta)$ .



### 3.1 Missing data prior densities

Although the desired Bayes factors are invariant to the choice of any missing data prior densities, this choice is important in practice for computations. This is largely a problem-specific issue, but we make two general remarks. First, if all models share the same missing data ( $y_1$ , say) then no missing data prior densities are required, and (7) becomes

$$\pi(\alpha, \theta, y|x) \propto \sum_{i=1}^n \alpha_i \pi_i(x, y_1|\theta_i) \pi(\alpha) \pi(\theta).$$

Second, it can be beneficial to assign missing data priors which mimic the marginal density of the  $y_{-\mathcal{I}(i)}$  components in other models. As discussed below, the mixing properties of suitable MCMC algorithms are improved if the chains can easily move between different models, and such movement is hindered if the density of the missing data in one model is very different to the missing prior density assigned in another.

### 3.2 MCMC methods

Sampling from the target density defined at (7) will typically be possible via a range of standard MCMC methods, but here we offer some observations on practical aspects. The fact that the target density is a sum will usually make direct Gibbs sampling infeasible, but the approach of Diebolt and Robert (1994), which relies on the introduction of allocation variables which indicate the ‘true’ model as described in Dempster et al. (1977), can be adapted as follows.

Introduce  $z = (z_1, \dots, z_n)$  such that  $z_i \in \{0, 1\}$  and  $\sum_{i=1}^n z_i = 1$ . Thus  $z$  can take  $n$  possible values, each of which is a vector of zeroes other than a 1 at one position. Define the augmented likelihood

$$\pi(z, x, y|\alpha, \theta) = \prod_{i=1}^n (\alpha_i \pi_i(x, y_{\mathcal{I}(i)}|\theta_i) \pi_i(y_{-\mathcal{I}(i)}|x, y_{\mathcal{I}(i)}, \theta))^{z_i},$$

so that the augmented likelihood at (2) is recovered by summing over  $z$ . If the prior distribution on  $\alpha$  is Dirichlet,  $\mathcal{D}(p_1, \dots, p_n)$ , it follows that  $\alpha$  has full

conditional distribution

$$\alpha | \dots \sim \mathcal{D}(p_1 + z_1, \dots, p_n + z_n).$$

The full conditional distribution of  $z$  is multinomial  $\mathcal{M}(1; q_1, \dots, q_n)$ , where the probabilities are given by

$$q_i \propto \alpha_i \pi_i(x, y_{\mathcal{I}(i)} | \theta_i) \pi_i(y_{-\mathcal{I}(i)} | x, y_{\mathcal{I}(i)}, \theta), \quad i = 1, \dots, n.$$

For  $i = 1, \dots, n$ ,  $\theta_i$  has full conditional distribution given by

$$\pi(\theta_i | \dots) \propto \begin{cases} \pi_i(\theta_i) & z_i = 0, \\ \pi_i(x, y_{\mathcal{I}(i)} | \theta_i) \pi_i(\theta_i) & z_i = 1. \end{cases}$$

Finally, any missing data component  $y_j$ ,  $j = 1, \dots, N$ , has full conditional distribution given by

$$\pi(y_j | \dots) \propto \begin{cases} \pi_i(x, y_{\mathcal{I}(i)} | \theta_i) & j \in \mathcal{I}(i), \\ \pi_i(y_{-\mathcal{I}(i)} | x, y_{\mathcal{I}(i)}, \theta_i) & j \notin \mathcal{I}(i), \end{cases}$$

where  $i$  denotes the current model, i.e.  $z_i = 1$ .

The prior distribution for  $\alpha$  can often be chosen to improve the mixing of the MCMC algorithm above. In particular this can be achieved by trying to make the multinomial distribution of  $z$  as close to uniform as possible. To illustrate this, consider the trivial example with two models in which  $\pi_1(x)/\pi_2(x) = m_1(x)/m_2(x) = B_{12} = 50$ . The full conditional distributions for  $z_1$  and  $\alpha_1$  are, respectively,  $Bern(50\alpha_1/(50\alpha_1 + (1 - \alpha_1)))$  and  $Beta(z_1 + p_1, 1 - z_1 + p_2)$ , where  $Bern$  and  $Beta$  respectively denote Bernoulli and Beta distributions. Setting  $p_1 = p_2 = 1$  produced wildly different estimates for  $B_{12}$  (87.8, 40.9, 547.1) for three MCMC runs of  $10^6$  iterations, while repeating the exercise with  $p_1 = 1$  and  $p_2 = 50$  yielded estimates 50.3, 50.3 and 50.7.

It is of course not necessary to use allocation variables, and one can equally use any suitable MCMC scheme for the target density. However, the above illustrates the fact that the full conditional distributions of  $\theta_i$  and any missing data will be of mixture form, which has implications for the design of efficient algorithms. Note that this also illustrates that the marginal densities  $\pi(\theta_i | x)$

are mixtures, and in particular are not the same as those obtained from a single-model analysis, which are proportional to  $\pi_i(x|\theta_i)\pi_i(\theta_i)$ . The marginal densities can be either be explored via a standard single-model MCMC algorithm, or by using the allocation variables approach and conditioning the output on  $z_i = 1$  to obtain within-model posterior density samples for  $\theta_i$ .

### 3.3 Connections with other approaches

The framework we adopt is related to that described in Carlin and Chib (1995) and Godsill (2001), in which the target distribution of interest is defined over a product space of models and their parameters. In order to clarify the differences in our approach, consider the simplest possible setting in which we have two models defined by densities  $\pi_1(x|\theta_1)$  and  $\pi_2(x|\theta_2)$ , and within-model prior densities  $\pi_1(\theta_1)$  and  $\pi_2(\theta_2)$ . The framework of Carlin and Chib (1995) and Godsill (2001) introduces a model indicator  $k \in \{1, 2\}$  to denote the ‘current’ model. The target density of interest is specified by

$$\pi(k, \theta_1, \theta_2|x) \propto \pi_k(x|\theta_k)\pi_k(\theta_k|k)\pi(\theta_{3-k}|\theta_k, k)\pi(k),$$

where it is necessary to specify  $\pi(\theta_{3-k}|\theta_k, k)$ , i.e. the ‘prior’ for the non-current model parameter. Assuming  $\theta_1$  and  $\theta_2$  to be independent of each other and  $k$  gives that  $\pi(\theta_{3-k}|\theta_k, k) = \pi_{3-k}(\theta_{3-k})$ .

Conversely, our formulation has target density

$$\pi(\alpha, \theta_1, \theta_2|x) \propto \pi(\alpha)\pi_1(\theta_1)\pi_2(\theta_2)[\alpha_1\pi_1(x|\theta_1) + \alpha_2\pi_2(x|\theta_2)].$$

If we adopt the allocation-variable approach, the target density becomes

$$\pi(z, \alpha, \theta_1, \theta_2|x) \propto \pi(\alpha)\pi_1(\theta_1)\pi_2(\theta_2)[\alpha_1\pi_1(x|\theta_1)]^{z_1}[\alpha_2\pi_2(x|\theta_1)]^{z_2},$$

from which we see that it is the existence of the  $\alpha$  parameter which distinguishes our formulation from that of Carlin and Chib (1995) and Godsill (2001). Of course, posterior estimation of  $\alpha$  is what enables us to estimate Bayes factors, so this difference is an important one.

The general formulation in Godsill (2001) also allows each model to potentially share parameters with other models. Specifically, the parameters of model  $k$  will be some subset of a set of parameters  $\{\theta_1, \dots, \theta_N\}$ . This is similar to the way we have dealt with missing data, although the two set-ups are not technically equivalent, and in particular one cannot simply treat our missing data as model parameters. The fundamental difference is that missing data may not always require a prior, whereas model parameters always do. For instance, if a density  $\pi(x|\theta)$  is intractable, then our missing data approach uses Bayes' Theorem in the form  $\pi(\theta, y|x) \propto \pi(x, y|\theta)\pi(\theta)$  whereas augmenting with extra model parameter  $\psi$  gives  $\pi(\theta, \psi|x) \propto \pi(x|\theta, \psi)\pi(\theta, \psi)$ .

## 4 Examples

In this section we illustrate the theory with three examples featuring population processes or epidemics. First, however, we consider a simpler classical example which briefly compares our methods with alternatives. This example illustrates that our methods have wider applicability than population processes, and moreover appear to be competitive against alternative methods.

### 4.1 Pines data

We consider the well-known model choice problem of assigning non-nested linear regression models to the pines data set in Williams (1959). These data have been analyzed by several authors (see, for example, Han and Carlin, 2001; Carlin and Chib, 1995; Friel and Pettitt, 2008) in order to compare methods for estimating Bayes factors. The data describe the maximum compression strength parallel to the grain  $y_i$ , the density  $x_i$ , and the resin-adjusted density  $z_i$  for 42 specimens of radiata pine. The two competing models we consider are

$$\begin{aligned} M_1 : \quad & y_i = \alpha + \beta(x_i - \bar{x}) + \epsilon_i, & \epsilon_i & \sim N(0, \sigma^2); \\ M_2 : \quad & y_i = \gamma + \delta(z_i - \bar{z}) + \eta_i, & \eta_i & \sim N(0, \tau^2). \end{aligned}$$

Table 1: Pines data set: Comparison of Bayes factors from different methods

Method	Bias	Standard Error
RJMCMC	67	2678
RJ corrected	9	124
Power posterior (serial MCMC)	10	132
Power posterior (population MCMC)	22	154
Mixture method	10	39

We assigned identical prior distributions for the parameters  $\alpha, \beta, \gamma, \delta, \sigma^2$  and  $\tau^2$  as the papers cited above. We assigned a Beta(100, 1) prior distribution for  $\alpha_1$  and carried out 100 MCMC runs of our method, this being the same as the number of MCMC runs used for the methods described in Friel and Pettitt (2008).

Furthermore, Friel and Pettitt (2008) compared conventional RJMCMC methods (with each model *a priori* equally likely), ‘corrected’ RJMCMC methods (model priors chosen to improve mixing) and two power posterior methods (Serial and Population MCMC). Full details can be found in Friel and Pettitt (2008), and for convenience we simply quote the results obtained in Table 1, along with our result. The bias is calculated by comparison with the estimate of 4862 obtained by numerical integration in Green and O’Hagan (1998). It can be seen that our method is certainly competitive.

## 4.2 Poisson process vs. linear birth process

Our first population process example is analytically tractable and illustrates that our methods produce results in agreement with the known true values. Consider data given by the vector of event times  $x = (x_1, \dots, x_n)$  observed during a time interval  $[0, T]$ , where  $0 \leq x_1 \leq x_2 \leq \dots \leq x_n \leq T$ . We will compare two models, namely a homogeneous Poisson process of rate  $\lambda$  ( $M_1$ ) and a linear birth process  $\{X(t) : t \in [0, T]\}$  with per-capita birth rate  $\mu$  and  $X(0) = 1$  ( $M_2$ ). Suppose further that  $\lambda$  and  $\mu$  are assigned independent exponential prior distributions

with mean  $\theta^{-1}$ . The model likelihoods, which we write as densities with respect to the reference measure induced by a unit rate Poisson process on  $[0, T]$ , are

$$\pi_1(x|\lambda) = \lambda^n \exp\{-(\lambda - 1)T\}, \quad \pi_2(x|\mu) = n!\mu^n \exp\{-\mu[(n + 1)T - S(x)] + T\},$$

where  $S(x) = \sum_{j=1}^n x_j$ . In this setting no missing data are required so we use the model defined at (1). The Bayes factor in favour of  $M_1$  relative to  $M_2$  is

$$\begin{aligned} B_{12} &= \frac{\int \pi_1(x|\lambda)\pi(\lambda) d\lambda}{\int \pi_2(x|\mu)\pi(\mu) d\mu} = \frac{\int_0^\infty \theta \lambda^n \exp\{-\lambda(T + \theta)\} d\lambda}{\int_0^\infty \theta n! \mu^n \exp\{-\mu[(n + 1)T - S(x) + \theta]\} d\mu} \\ &= \frac{[(n + 1)T - S(x) + \theta]^{n+1}}{(T + \theta)^{n+1} n!}. \end{aligned}$$

Assuming that  $\alpha_1 \sim U(0, 1)$  *a priori*, a simple Gibbs sampler for the target density consists of parameter updates as follows:

$$\begin{aligned} \alpha_1 | \dots &\sim \text{Beta}(z_1 + 1, 2 - z_1), \\ z_1 | \dots &\sim \text{Bern}\left(\frac{\alpha_1 \pi_1(x|\lambda)}{\alpha_1 \pi_1(x|\lambda) + (1 - \alpha_1) \pi_2(x|\mu)}\right), \\ \lambda | \dots &\sim \begin{cases} \Gamma(1, \theta) & z_1 = 0, \\ \Gamma(n + 1, T + \theta) & z_1 = 1, \end{cases} \\ \mu | \dots &\sim \begin{cases} \Gamma(n + 1, (n + 1)T - S(x) + \theta) & z_1 = 0, \\ \Gamma(1, \theta) & z_1 = 1, \end{cases} \end{aligned}$$

where  $\Gamma(m, \xi)$  denotes a Gamma distribution with density  $f(x) \propto x^{m-1} \exp(-\xi x)$ .

Typical results from MCMC runs are given in Table 2, illustrating that the Gibbs sampler recovers the true known values. We found that the algorithm mixing was good in all cases.

Finally, we comment on the relationship between the above algorithm and standard reversible jump methods. The RJMCMC requires a way of proposing a value of  $\mu$  given  $\lambda$  for jumps from  $M_1$  to  $M_2$ , and vice versa. In practice it is not immediately obvious how best to do this, but an approach suggested in Green (2003) is to propose  $\mu$  independently of  $\lambda$ , ideally according to the within-model density  $\pi(\mu|x)$ . This is similar to what we obtain above.

Table 2: Example 2: Bayes factors from MCMC output ( $\hat{B}_{12}$ ) compared to true values ( $B_{12}$ ).

$n$	$T$	$S(x)$	$\theta$	$\hat{B}_{12}$	$B_{12}$
5	10	36	1	1.15	1.148
5	10	36	0.01	1.58	1.587
5	10	25	1	10.25	10.239
10	20	150	1	0.18	0.181

### 4.3 SIR model with two different infection periods

Recall the standard SIR (Susceptible-Infective-Removed) epidemic model (see e.g. Andersson and Britton (2000), Chapter 2), defined as follows. A closed population contains  $N + a$  individuals of whom  $N$  are initially susceptible and  $a$  initially infective. Each infective remains so for a period of time distributed according to a specified random variable  $T_I$ , known as the infectious period, after which it becomes removed and plays no further part in the epidemic. During its infectious period an infective makes contact with each other member of the population at times given by a homogeneous Poisson process of rate  $\beta/N$ , and any contact occurring with a susceptible individual results in that individual immediately becoming infective. The infectious periods of different individuals and the Poisson processes between different pairs of individuals are assumed to be mutually independent. The epidemic ends when there are no infectives left in the population.

A distinguishing characteristic of infectious disease data is that the infection process itself is rarely observed, and so we suppose that the data  $r$  consist of  $n$  observed removal times  $r_1 \leq \dots \leq r_n$ . We consider two competing models, namely that  $T_I \sim \Gamma(1, \gamma)$  ( $M_1$ ) and  $T_I \sim \Gamma(m, \lambda)$  ( $M_2$ ), where the shape parameter  $m$  will be assumed known. Both model likelihoods  $\pi_1(r|\gamma)$  and  $\pi_2(r|\lambda)$  are intractable in practice since their evaluation relies on integrating over all possible realisations of the infection process, and so we introduce missing data as follows.

For  $j = 1, \dots, n$  define  $i_j$  as the infection time of the individual removed at time  $r_j$ . We assume that there is  $a = 1$  initial infective, denoted by  $p$ , so that  $i_p \leq i_j$  for all  $j \neq p$ . For simplicity we assume *a priori* that  $p$  is equally likely to be any of the  $n$  infected individuals and that  $i_p$  has an improper uniform prior density on  $(-\infty, r_1)$ . Finally, define  $i = \{i_j : i \neq p\}$  to be the  $n - 1$  non-initial infection times. For  $k = 1, 2$  the augmented model likelihoods, which we write here with respect to Lebesgue measure on  $\mathbb{R}^{2n-1}$ , are

$$\pi_k(i, r | p, i_p, \beta_k, \eta_k) = \left( \prod_{j=1; j \neq p}^n (\beta_k/N) I(i_j -) \right) \exp \left\{ -(\beta_k/N) \int_{i_p}^{r_n} S(t) I(t) dt \right\} \left( \prod_{j=1}^n f_k(r_j - i_j | \eta_k) \right),$$

where  $S(t)$  and  $I(t)$  denote respectively the numbers of susceptibles and infectives at time  $t$ ,  $I(t-) = \lim_{s \uparrow t} I(s)$ ,  $\beta_k$  denotes the parameter  $\beta$  under  $M_k$ ,  $f_k$  denotes the infectious period density under  $M_k$ ,  $\eta_1 = \gamma$  and  $\eta_2 = \lambda$  (see e.g. O'Neill and Roberts (1999), Streftaris and Gibson (2004), Höhle and O'Neill (2005)). Note that in this formulation, the missing data  $i$ ,  $p$  and  $i_p$  are assumed common to both models, although these quantities could also be model-specific.

The target density of interest is

$$\pi(\alpha_1, \beta_1, \beta_2, \gamma, \lambda | r) \propto [\alpha_1 \pi_1(i, r | p, i_p, \beta_1, \gamma) + (1 - \alpha_1) \pi_2(i, r | p, i_p, \beta_2, \lambda)] \pi(\beta_1) \pi(\beta_2) \pi(\gamma) \pi(\lambda).$$

Note that here we need no missing data priors densities because the the missing data appear in both model likelihoods. Prior distributions for  $\beta_1, \beta_2, \gamma$  and  $\lambda$  were all set as  $\Gamma(1, 1)$ , and  $\alpha_1 \sim \text{Beta}(p_1, p_2)$ .

Introducing the allocation variable  $z_1$  yields the full conditional distributions



below, each of which yields a simple Gibbs update for the parameter in question.

$$\begin{aligned}
\alpha_1 | \dots &\sim \text{Beta}(z_1 + p_1, 1 - z_1 + p_2), \\
z_1 | \dots &\sim \text{Bern} \left( \frac{\alpha_1 \pi_1(i, r | p, i_p, \beta_1, \gamma)}{\alpha_1 \pi_1(i, r | p, i_p, \beta_1, \gamma) + (1 - \alpha_1) \pi_2(i, r | p, i_p, \beta_2, \lambda)} \right), \\
\beta_1 | \dots &\sim \begin{cases} \Gamma(n, N^{-1} \int_{i_p}^{r_n} S(t) I(t) dt + 1) & z_1 = 0, \\ \Gamma(1, 10^{-3}) & z_1 = 1, \end{cases} \\
\beta_2 | \dots &\sim \begin{cases} \Gamma(1, 1) & z_1 = 0, \\ \Gamma(n, N^{-1} \int_{i_p}^{r_n} S(t) I(t) dt + 1) & z_1 = 1, \end{cases} \\
\gamma | \dots &\sim \begin{cases} \Gamma(n + 1, \sum_{j=1}^n (r_j - i_j) + 1) & z_1 = 0, \\ \Gamma(1, 1) & z_1 = 1, \end{cases} \\
\lambda | \dots &\sim \begin{cases} \Gamma(1, 1) & z_1 = 0, \\ \Gamma(nm + 1, \sum_{j=1}^n (r_j - i_j) + 1) & z_1 = 1, \end{cases}
\end{aligned}$$

Finally, the infection time parameters  $i$ ,  $i_p$  and  $p$  are updated using a Metropolis-Hastings step as follows. One of the  $n$  infected individuals,  $j$  say, is chosen uniformly at random. A proposed new infection time for  $j$  is defined as  $i_j^* = r_j - x$ , where  $x$  is sampled from a  $\Gamma(1, \delta)$  distribution. Note that this may also result in proposed new values for  $p$  and  $i_p$ ; either way, proposed values are denoted  $i^*$ ,  $i_p^*$  and  $p^*$  and accepted with probability

$$1 \wedge \frac{\pi_k(i^*, r | p^*, i_p^*, \beta_k, \eta_k)}{\pi_k(i, r | p, i_p, \beta_k, \eta_k)} \exp(\delta(i_j - i_j^*)),$$

where  $k = 2 - z_1$  denotes the ‘current’ model.

To illustrate the algorithm, we considered the SIR model with  $N = 50$ ,  $a = 1$ , various  $\beta$  values and  $\Gamma(\tilde{m}, \tilde{\lambda})$  infectious periods with three different choices for  $(\tilde{m}, \tilde{\lambda})$ . For each scenario we simulated 100 data sets, and evaluated the Bayes factor using the above MCMC algorithm for each data set. For two of the  $(\tilde{m}, \tilde{\lambda})$  pairs we set the shape parameter  $m$  in  $M_2$  equal to  $\tilde{m}$ , and for one we did not. In practice, one is rarely interested in data from epidemics with few cases, so we also evaluated the Bayes factors using a subset of each of the 100 simulations in which the epidemic had clearly ‘taken off’, evaluated by eye, which we refer

Table 3: Example 3: Bayes factors from MCMC output.

Scenario	True model	$\beta$	$M_2$	$E[B_{12}](st.dev.)$ (all simulations)	$E[B_{12}](st.dev.)$ (major epidemics)
A	$\Gamma(5, 5)$	2	$\Gamma(5, \lambda)$	0.06 (0.06)	0.008 (0.006)
B	$\Gamma(1, 0.75)$	1	$\Gamma(1, \lambda)$	1.03 (0.17)	1.05 (0.22)
C	$\Gamma(1, 1)$	3	$\Gamma(2, \lambda)$	3022 (3969)	2291 (3428)

to as major epidemics. The numbers of major epidemics were 63, 56 and 65 for scenarios A, B and C, respectively.

Table 3 contains a summary of the Bayes factors estimated from the simulated data sets. In scenarios A, B and C the true models are  $M_2$ , both  $M_1$  and  $M_2$ , and  $M_1$  respectively. The estimated Bayes factors behave as we might expect, giving clear evidence in favour of models  $M_2$  and  $M_1$  for scenarios A and C respectively, whilst for scenario B the mean of  $B_{12}$  is close to the true value of 1. In scenario A there is a marked difference in the Bayes factors when using all simulations compared to using only major epidemics. A possible explanation is that major epidemics contain more data, and so any difference between the models becomes easier to detect. There is a less pronounced difference in Bayes factors in scenario C, although the large posterior standard deviations suggest there is no compelling evidence for a clear difference in this case.

#### 4.4 SIR epidemic model vs. Poisson process

Our final example is motivated by the situation in which we wish to decide whether observed cases of disease are the result of an epidemic (with transmission between individuals) or simply sporadic events. Specifically, suppose we observe  $n$  events at times  $0 < r_1 < \dots < r_n < T$ , and let  $r = (r_1, \dots, r_n)$ . Under model  $M_1$ ,  $r$  is a vector of event times of a homogeneous Poisson process of rate  $\lambda$  observed during the time interval  $[0, T]$ . Under model  $M_2$ ,  $r$  is a vector of removal times in an SIR epidemic model with exponentially distributed infectious periods, again observed during  $[0, T]$ .

As for the previous example, we proceed by adding unobserved infection times in order to obtain a tractable likelihood for  $M_2$ . For simplicity we assume that there is one initial infective at time zero, and furthermore that there is a population of  $N$  individuals in total, where  $N \geq n$ . Unlike the previous example, in which we unobserved infection times with observed removal times, we here define  $i = (i_2, \dots, i_m)$  to be a vector of  $m$  ordered infection times, so that  $0 = i_1 < i_2 < \dots < i_m$ , where  $n \leq m \leq N$ . The reason for this approach is that it appears to be easier when it comes to assigning missing data prior densities, as described below. Note also that under  $M_2$  we allow the possibility that the epidemic is still in progress at time  $T$ .

The likelihood for  $M_1$  and augmented likelihood for  $M_2$  are respectively given by

$$\begin{aligned}\pi_1(r|\lambda) &= \lambda^n \exp\{-\lambda T\}, \\ \pi_2(i, r|\beta, \gamma) &= \left( \prod_{j=2}^m \beta S(i_{j-}) I(i_{j-}) \right) \left( \prod_{j=1}^n \gamma I(r_{j-}) \right) \exp \left\{ - \int_0^T \beta S(t) I(t) + \gamma I(t) dt \right\}.\end{aligned}$$

To proceed we require a missing data prior density  $\pi_1(i|r, \lambda, \beta, \gamma)$ . Now for a given ordered vector of event times  $r$ ,  $\pi_2(i, r|\beta, \gamma) > 0$  if and only if  $i \in \mathcal{F}(r)$ , where

$$\mathcal{F}(r) = \{i : i_1 < i_2 < \dots < i_m < T; i_k < r_{k-1}, k = 1, \dots, n+1; i_k < T, k = n+2, \dots, m\}.$$

One way to define  $\pi_1(i|r, \lambda, \beta, \gamma)$  is via the following construction, which simulates an element of  $\mathcal{F}(r)$ . First, select  $m$  according to some probability mass function  $f$  on  $\{n, n+1, \dots, N\}$ . Next, sequentially set  $i_2 \sim \text{TrExp}(\mu; i_1, r_1)$ ,  $i_3 \sim \text{TrExp}(\mu; i_2, r_2), \dots, i_{n+1} \sim \text{TrExp}(\mu; i_n, r_n), i_{n+2} \sim \text{TrExp}(\mu; i_{n+1}, T), \dots, i_m \sim \text{TrExp}(\mu; i_{m-1}, T)$ , where  $\text{TrExp}(\mu; a, b)$  denotes an exponential random variable with rate  $\mu$ , truncated to the interval  $(a, b)$ . This in turn induces a probability distribution with density

$$\pi_1(i|r) = f(m) \prod_{j=1}^{m-1} \frac{\mu \exp(-\mu i_j)}{\exp(-\mu i_{j-1}) - \exp(-\mu s_{j-1})}, \quad i \in \mathcal{F}(r),$$

where  $s_j = r_j$  for  $j = 1, \dots, n$  and  $s_j = T$  for  $n < j \leq m$ , and we set  $\pi_1(i|r, \lambda, \beta, \gamma) = \pi_1(i|r)$ .

We remark that it is not necessary to define the missing data prior density in this manner. For instance, one could proceed by choosing  $m$  as before and then assigning a uniform density to the set  $\{i : i_2 < i_3 < \dots < i_m\}$ . The practical drawback with this is that, if using allocation variables, the Markov chain can never leave model  $M_1$  if  $i$  is such that  $\pi_2(i, r|\beta, \gamma) = 0$ .

Prior distributions were assigned as  $\beta \sim \Gamma(\nu_\beta, \mu_\beta)$ ,  $\gamma \sim \Gamma(\nu_\gamma, \mu_\gamma)$ ,  $\lambda \sim \Gamma(\nu_\lambda, \mu_\lambda)$  and  $\alpha_1 \sim \text{Beta}(p_1, p_2)$ .

An MCMC algorithm is easily developed in a similar manner to the previous example. Specifically we have the following full conditional distributions:

$$\begin{aligned} \alpha_1 | \dots &\sim \text{Beta}(z_1 + p_1, 1 - z_1 + p_2), \\ z_1 | \dots &\sim \text{Bern} \left( \frac{\alpha_1 \pi_1(r|\lambda) \pi_1(i|r)}{\alpha_1 \pi_1(r|\lambda) \pi_1(i|r) + (1 - \alpha_1) \pi_2(i, r|\beta, \gamma)} \right), \\ \lambda | \dots &\sim \begin{cases} \Gamma(\nu_\lambda, \mu_\lambda) & z_1 = 0, \\ \Gamma(n + \nu_\lambda, T + \mu_\lambda) & z_1 = 1, \end{cases} \\ \beta | \dots &\sim \begin{cases} \Gamma(m - 1 + \nu_\beta, \int_0^T S(t)I(t) dt + \mu_\beta) & z_1 = 0, \\ \Gamma(\nu_\beta, \mu_\beta) & z_1 = 1, \end{cases} \\ \gamma | \dots &\sim \begin{cases} \Gamma(n + \nu_\gamma, \int_0^T I(t) dt + \mu_\gamma) & z_1 = 0, \\ \Gamma(\nu_\gamma, \mu_\gamma) & z_1 = 1. \end{cases} \end{aligned}$$

Updates for  $i$  are achieved as follows. If  $z_1 = 1$  then  $i$  has full conditional density  $\pi_1(i|r)$  which can be sampled as described above. If  $z_1 = 0$  then  $i$  can be updated by moving, adding or deleting infection times as described in O'Neill and Roberts (1999).

To illustrate this algorithm we considered a data set taken from an outbreak of Gastroenteritis described in Britton and O'Neill (2002) which take the form of 28 case detection times among a population of 89 individuals. The daily numbers of cases on days 0 to 7 are given respectively by

$$1, 0, 4, 2, 3, 3, 10, 5.$$

Strictly speaking, such data should be analysed by allowing the unknown time of the initial infection,  $i_1$ , to be estimated (see e.g. O'Neill and Roberts (1999)).

Since our main objective here is to illustrate our methodology, we instead make the simplifying assumption that day 0 actually corresponded to the start of the outbreak, and then consider the remaining 27 case detection times.

For the missing data prior density  $\pi_1(i|r)$  we set

$$f(m) = \frac{(1 - \theta)^{m-n}\theta}{1 - (1 - \theta)^{N-n+1}}, \quad m = n, \dots, N,$$

so that  $m$  has a truncated Geometric distribution with parameter  $\theta$ , and set  $\mu = 4$  in the truncated exponential distribution.

We ran the algorithm with two choices of  $T$ , the time of observation, with  $\beta$ ,  $\gamma$  and  $\mu$  all given  $Exp(1)$  prior distributions. First, with  $T = 10$  we estimated the Bayes factor in favour of the Poisson model to be 0.003, here using  $p_1 = 400$ ,  $p_2 = 1$  to obtain reasonable mixing in the MCMC algorithm. So in this case there appears to be overwhelming evidence to suggest that the case detection times are better described by an epidemic model than a Poisson process. Second, we set  $T = 3.5$  and used only the case observation times up until day 3. We estimated the Bayes factor in favour of the Poisson model to be 21.1. In comparison to the  $T = 10$  case we would certainly expect a value closer to 1, since there are less data, and equally it is intuitively reasonable that there are insufficient data to provide evidence in favour of an epidemic.

## 5 Discussion

We have presented a new method for evaluating Bayes factors. Although motivated by epidemic models and population processes, our approach is clearly applicable in more general settings, as illustrated by the pines data set example in Section 4.

The methods we propose are not without drawbacks. First, in common with the product-space methods it seems likely that they are best suited to situations in which there are only a small number of competing models, although we have not investigated this issue in this paper. Second, constructing missing data prior densities, when required, seems likely to require problem-specific insights

in order to obtain reasonably efficient algorithms. Intuitively we expect that it is best to choose missing data prior distributions to mimic the true distribution of missing data in competing models. These aspects, as well as the method in general, appear worthy of more detailed exploration.

**Acknowledgement** We thank Christian Robert and Andy Wood for helpful discussions about this work. We acknowledge funding from UK EPSRC grant EP/F017006/1.

## References

- Andersson, H. and T. Britton (2000). *Stochastic Epidemic Models and their Statistical Analysis*. New York: Springer.
- Britton, T. and P. D. O’Neill (2002). Bayesian inference for stochastic epidemics in populations with random social structure. *Scand. J. Statist.* 29(3), 375–390.
- Brooks, S., P. Giudici, and G. Roberts (2003). Efficient construction of reversible jump Markov chain Monte Carlo proposal distributions. *J. R. Statist. Soc. B* 65, 3–39.
- Carlin, B. P. and S. Chib (1995). Bayesian model choice via Markov chain Monte Carlo methods. *J. R. Statist. Soc. B* 57(3), 473–484.
- Celeux, G., F. Forbes, C. Robert, and D. Titterton (2006). Deviance Information Criteria for missing data models. *Bayesian Analysis* 1, 651–674.
- Dellaportas, P., J. Forster, and I. Ntzoufras (2002). On Bayesian model and variable selection using MCMC. *Statistics and Computing* 12(1), 27–36.
- Dempster, A., N. Laird, and D. Rubin (1977). Maximum likelihood from incomplete data via the EM algorithm. *J. R. Statist. Soc. B* 39(1), 1–38.

- Diebolt, J. and C. P. Robert (1994). Estimation of finite mixture distributions through Bayesian sampling. *J. R. Statist. Soc. B* 56(2), 363–375.
- Friel, N. and A. N. Pettitt (2008). Marginal likelihood estimation via power posteriors. *J. R. Stat. Soc. Ser. B* 70(3), 589–607.
- Godsill, S. J. (2001). On the relationship between markov chain monte carlo methods for model uncertainty. *Journal of Computational and Graphical Statistics* 10(2), 1–19.
- Green, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* 82(4), 711–732.
- Green, P. J. (2003). Transdimensional Markov chain Monte Carlo. In P. J. Green, N. L. Hjort, and S. Richardson (Eds.), *Highly Structured Stochastic Systems*, pp. 179–98. Oxford: OUP.
- Green, P. J. and A. O’Hagan (1998). Model choice with MCMC on product spaces without using pseudo-priors. *Technical Report, Department of Mathematics, University of Nottingham*.
- Han, C. and B. P. Carlin (2001). Markov Chain Monte Carlo Methods for Computing Bayes Factors: A Comparative Review. *Journal of the American Statistical Association* 96(455), 1122–1132.
- Höhle, M. and P. D. O’Neill (2005). Inference in disease transmission experiments by using stochastic epidemic models. *Appl. Statist.* 54(2), 349–366.
- O’Neill, P. D. and G. O. Roberts (1999). Bayesian inference for partially observed stochastic epidemics. *J. Roy. Statist. Soc. Ser. A* 162, 121–129.
- Richardson, S. and P. J. Green (1997). On Bayesian analysis of mixtures with an unknown number of components. *J. R. Statist. Soc. B* 59(4), 731–792.
- Streftaris, G. and G. J. Gibson (2004). Bayesian inference for stochastic epidemics in closed populations. *Statistical Modelling* 4, 1–13.

Williams, E. J. (1959). *Regression analysis*. A Wiley Publication in Applied Statistics. New York: John Wiley & Sons Inc.

## 6 Appendix

### 6.1 Proof of Lemma 1

(a) Define

$$\mathbf{b} = [B_{1k}(x) \cdots B_{nk}(x)]^T,$$

so that (5) can be written as the matrix equation  $\mathbf{A}\mathbf{b} = \mathbf{0}$ . Since  $B_{kk}(x) = 1$ , we can rewrite (5) as

$$\sum_{j \neq k} A_{ij} B_{jk}(x) = -A_{ik}, \quad i = 1, \dots, n. \quad (8)$$

Now, for  $1 \leq l, j \leq n$ ,

$$\begin{aligned} A_{lj} &= E[\alpha_l | x] E[\alpha_j] - E[\alpha_l \alpha_j] \\ &= E[\alpha_j] \left( 1 - \sum_{i \neq l} E[\alpha_i | x] \right) - E \left[ \left( 1 - \sum_{i \neq l} \alpha_i \right) \alpha_j \right] \\ &= - \sum_{i \neq l} (E[\alpha_j] E[\alpha_i | x] - E[\alpha_i \alpha_j]) \\ &= - \sum_{i \neq l} A_{ij}. \end{aligned} \quad (9)$$

Summing (8) over  $i \neq k$  and using (9) now yields

$$\begin{aligned} \sum_{i \neq k} \sum_{j \neq k} A_{ij} B_{jk}(x) &= - \sum_{i \neq k} A_{ik} \\ \text{so } \sum_{j \neq k} \left( \sum_{i \neq k} A_{ij} \right) B_{jk}(x) &= - \sum_{i \neq k} A_{ik} \\ \text{so } \sum_{j \neq k} A_{kj} B_{jk}(x) &= -A_{kk}, \end{aligned}$$

which is the equation obtained from (8) when  $i = k$ . In other words, at least the  $k$ th equation in (8) is redundant. It is therefore sufficient to consider the



system of equations defined by

$$\tilde{A}_{-k}\tilde{\mathbf{b}} = \tilde{\mathbf{c}}, \quad (10)$$

where  $\tilde{\mathbf{b}}$  is the  $(n-1) \times 1$  column vector formed by removing  $B_{kk}(x) = 1$  from  $\mathbf{b}$ , and  $\tilde{\mathbf{c}}$  is the  $(n-1) \times 1$  column vector with components  $-A_{ik}$  for  $i = 1, \dots, n$ ,  $i \neq k$ . Application of Cramer's rule to solve (10) now yields part (a).

(b) For the second part, we require some preliminary results. Write  $E[\alpha_i|x] = f(m_i(x))$ , say. From (4) we have

$$\begin{aligned} f(m_i(x)) &= \frac{\sum_{j=1}^n E[\alpha_i \alpha_j] m_j(x)}{\sum_{j=1}^n E[\alpha_j] m_j(x)} \\ &= \frac{\sum_{j \neq i} E[\alpha_i \alpha_j] m_j(x) + E[\alpha_i^2] m_i(x)}{\sum_{j \neq i} E[\alpha_j] m_j(x) + E[\alpha_i] m_i(x)}. \end{aligned}$$

Differentiation yields that  $f'(m_i(x)) \geq 0$  if and only if  $C \geq 0$ , where

$$C = \sum_{j \neq i} (E[\alpha_i^2] E[\alpha_j] - E[\alpha_i] E[\alpha_i \alpha_j]) m_j(x).$$

Thus if  $C \geq 0$  we obtain the bounds

$$\frac{\sum_{j \neq i} E[\alpha_i \alpha_j] m_j(x)}{\sum_{j \neq i} E[\alpha_j] m_j(x)} \leq E[\alpha_i|x] \leq \frac{E[\alpha_i^2]}{E[\alpha_i]}, \quad (11)$$

and moreover the lower and upper bounds are attained when  $m_i(x) = 0$  and  $m_i(x) \rightarrow \infty$ , respectively. In particular, for  $C > 0$  and  $0 < m_i(x) < \infty$  then both inequalities are strict. If  $C \leq 0$  then the inequalities in (11) are simply reversed. From now on we assume that  $0 < m_i(x) < \infty$  for all  $i = 1, \dots, n$ .

Now if  $n = 2$  then (11) yields that for  $i \neq j$ ,

$$\frac{E[\alpha_1 \alpha_2]}{E[\alpha_j]} \neq E[\alpha_i|x],$$

from which it follows that  $A_{ij} = E[\alpha_i|x]E[\alpha_j] - E[\alpha_1 \alpha_2] \neq 0$ . The result for  $n = 2$  now follows directly from part (a).

For the final part, in which  $\alpha$  has a Dirichlet prior distribution, we first show that  $\det \tilde{A}_{-k} \neq 0$ , so that (10) has a unique solution. Secondly we show that this solution is given by  $B_{jk}(x) = A_{jk}/A_{kj}$ .

We start with conditions under which  $C > 0$ . Specifically, if  $\text{Cov}(\alpha_i, \alpha_j) < 0$  for all  $i \neq j$  and  $\text{Var}(\alpha_i) > 0$  then

$$\begin{aligned} E[\alpha_i]E[\alpha_j] &> E[\alpha_i\alpha_j] \\ \text{so } E[\alpha_i^2]E[\alpha_i]E[\alpha_j] &> E[\alpha_i]^2E[\alpha_i\alpha_j], \end{aligned}$$

from which it follows that  $C > 0$ .

Next, suppose that  $\alpha \sim \mathcal{D}(p_1, \dots, p_n)$  and set  $p_0 = \sum_{i=1}^n p_i$ . Thus for  $i \neq j$ ,  $E[\alpha_i\alpha_j] = p_i p_j / (p_0(p_0 + 1))$ ,  $E[\alpha_i] = p_i / p_0$ ,  $E[\alpha_i^2] = p_i(p_i + 1) / p_0(p_0 + 1)$ ,  $\text{Cov}(\alpha_i, \alpha_j) < 0$  and  $\text{Var}(\alpha_i) > 0$ . It follows that  $C > 0$  and that (11) simplifies to

$$\frac{p_i}{p_i + 1} < E[\alpha_i|x] < \frac{p_i + 1}{p_0 + 1}. \quad (12)$$

Next, note that for  $i \neq j$  we have

$$\begin{aligned} A_{ij} &= E[\alpha_j]E[\alpha_i|x] - E[\alpha_i\alpha_j] \\ &= \frac{p_j}{p_0} \left( E[\alpha_i|x] - \frac{p_i}{p_0 + 1} \right) \\ &= b_j a_i(x), \end{aligned}$$

say, where  $b_j = p_j / p_0$ . It follows from (12) that  $A_{ij} > 0$ . Similarly

$$\begin{aligned} A_{ii} &= \frac{p_i}{p_0} \left( E[\alpha_i|x] - \frac{p_i + 1}{p_0 + 1} \right) \\ &= b_i \tilde{a}_i(x), \end{aligned}$$

say. Recall that  $\tilde{A}_{-k}$  is the matrix  $A$  with the  $k$ th row and column deleted. It now follows that

$$\det(\tilde{A}_{-k}) = \left( \prod_{i \neq k} b_i \right) \det(D + E),$$

where  $D$  is an  $(n-1) \times (n-1)$  diagonal matrix with entries  $\tilde{a}_i(x) - a_i(x) = -1/(p_0 + 1)$ ,  $i \neq k$ , and  $E$  is an  $(n-1) \times (n-1)$  matrix consisting of  $(n-1)$  identical columns, each of which contains the  $(n-1)$  entries  $a_i$ ,  $i \neq k$ . Moreover we can write  $E$  as the product  $uv^T$ , where  $u$  is an  $(n-1) \times 1$  column vector with entries  $a_i$ ,  $i \neq k$ , and  $v$  is the  $(n-1) \times 1$  column vector of 1's. In particular,  $\det \tilde{A}_{-k} \neq 0$  if and only if  $\det(D + E) = \det(D + uv^T) \neq 0$ .

Now from the matrix determinant lemma,

$$\det(D + uv^T) = (1 + v^T D^{-1}u)\det(D),$$

and since  $\det(D) = (-1/(p_0 + 1))^{n-1} \neq 0$ , we focus on  $1 + v^T D^{-1}u$ . Now

$$\begin{aligned} 1 + v^T D^{-1}u &= 1 - \sum_{i \neq k} (p_0 + 1)a_i(x) \\ &= 1 - \sum_{i \neq k} [(p_0 + 1)E[\alpha_i|x] - p_i] \\ &= 1 - (p_0 + 1)(1 - E[\alpha_k|x]) + (p_0 - p_k) \\ &= E[\alpha_k|x](p_0 + 1) - p_k > 0, \end{aligned}$$

where the last inequality follows from (12). Hence  $\det \tilde{A}_{-k} \neq 0$  as required.

Finally, we show that for  $i \neq k$ , (8) is satisfied by  $B_{jk}(x) = A_{jk}/A_{kj}$ . First, it is straightforward to show that for  $i \neq k$ ,

$$\tilde{a}_i(x) + \sum_{j \neq k; j \neq i} a_j(x) = -a_k(x). \quad (13)$$

Now,

$$\begin{aligned} \sum_{j \neq k} A_{ij} \frac{A_{jk}}{A_{kj}} &= b_i \tilde{a}_i(x) \frac{b_k a_i(x)}{b_i a_k(x)} + \sum_{j \neq k; j \neq i} b_j a_i(x) \frac{b_k a_j(x)}{b_j a_k(x)} \\ &= \frac{a_i(x) b_k}{a_k(x)} \left( \tilde{a}_i(x) + \sum_{j \neq k; j \neq i} a_j(x) \right) \\ &= \frac{a_i(x) b_k}{a_k(x)} (-a_k(x)) \\ &= -A_{ik}, \end{aligned}$$

using (13). Hence for  $i \neq k$ , (8) is satisfied by  $B_{jk}(x) = A_{jk}/A_{kj}$  as required.

**Example 1** To illustrate the calculations in part (a), consider the case  $n = 3$ ,  $k = 1$ . The equation  $A\mathbf{b} = \mathbf{0}$  is

$$\begin{bmatrix} A_{11} & A_{12} & A_{13} \\ A_{21} & A_{22} & A_{23} \\ A_{31} & A_{32} & A_{33} \end{bmatrix} \begin{bmatrix} 1 \\ B_{21} \\ B_{31} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix},$$

and so

$$\tilde{A}_{-1} = \begin{bmatrix} A_{22} & A_{23} \\ A_{32} & A_{33} \end{bmatrix}, \quad \tilde{\mathbf{b}} = \begin{bmatrix} B_{21} \\ B_{31} \end{bmatrix}, \quad \mathbf{c} = \begin{bmatrix} -A_{21} \\ -A_{31} \end{bmatrix}.$$

Applying Lemma 1 yields

$$B_{21} = \frac{\det \tilde{A}_{-21}}{\det \tilde{A}_{-1}} = \frac{\det \begin{bmatrix} -A_{21} & A_{23} \\ -A_{31} & A_{33} \end{bmatrix}}{\det \tilde{A}_{-1}}, \quad B_{31} = \frac{\det \tilde{A}_{-31}}{\det \tilde{A}_{-1}} = \frac{\det \begin{bmatrix} A_{22} & -A_{21} \\ A_{32} & -A_{31} \end{bmatrix}}{\det \tilde{A}_{-1}}.$$

**Example 2** To show that  $B_{jk}(x)$  does not equal  $A_{jk}/A_{kj}$  in general, suppose that  $n = 3$ , that  $m_i(x) = i$  for  $i = 1, 2, 3$ , and that  $\alpha$  has a mixed Dirichlet prior distribution given by

$$\alpha \sim (0.5)\mathcal{D}(1, 1, 1) + (0.5)\mathcal{D}(1, 2, 1).$$

Direct calculation then yields that  $E[\alpha_1] = E[\alpha_3] = 7/24$ ,  $E[\alpha_2] = 10/24$  and

$$\begin{bmatrix} E[\alpha_1^2] & E[\alpha_1\alpha_2] & E[\alpha_1\alpha_3] \\ E[\alpha_2\alpha_1] & E[\alpha_2^2] & E[\alpha_2\alpha_3] \\ E[\alpha_3\alpha_1] & E[\alpha_3\alpha_2] & E[\alpha_3^2] \end{bmatrix} = \frac{1}{20} \begin{bmatrix} 2 & 2 & 1 \\ 2 & 6 & 2 \\ 1 & 2 & 2 \end{bmatrix},$$

whence  $E[\alpha_1|x] = 31/120$ ,  $E[\alpha_2|x] = 50/120$  and  $E[\alpha_3|x] = 39/120$ . Thus  $A_{21}/A_{12} = 86/46$  while  $B_{21} = m_2(x)/m_1(x) = 2$ .