# BAYESIAN VARIABLE SELECTION FOR TIME SERIES COUNT DATA 

Joseph G. Ibrahim* ${ }^{* \dagger}$, Ming-Hui Chen\# and Louise M. Ryan* ${ }^{*}$<br>*Harvard School of Public Health, ${ }^{\dagger}$ Dana-Farber Cancer Institute and \# Worcester Polytechnic Institute


#### Abstract

We consider a parametric model for time series of counts by constructing a likelihood-based generalization of a model considered by Zeger (1988). We consider a Bayesian approach and propose a class of informative prior distributions for the model parameters that are useful for variable subset selection. The prior specification is motivated from the notion of the existence of data from similar previous studies, called historical data, which is then quantified in a prior distribution for the current study. We derive theoretical and computational properties of the proposed priors and develop novel methods for computing posterior model probabilities. To compute the posterior model probabilities, we show that only posterior samples from the full model are needed to estimate the posterior probabilities for all of the possible subset models. We demonstrate our methodology with a simulated and a real data set.


Key words and phrases: Correlated counts, Gibbs sampling, hierarchical centering, historical data, Poisson regression, posterior distribution.

## 1. Introduction

Data from similar previous studies, called historical data, is often available in applied research settings where the investigator has access to previous studies measuring the same response and covariates as the current study. From a Bayesian perspective, historical data can be very helpful in interpreting the results of the current study. However, very few methods exist for the formal incorporation of historical data into a prior distribution. There is some literature addressing this issue for the linear model and generalized linear models. See for example, Ibrahim, Ryan and Chen (1998), Chen, Ibrahim and Yiannoutsos (1999), and Bedrick, Christensen and Johnson (1996). In these papers, the authors assume a univariate independent response variable. The literature for informative prior elicitation for models with correlated responses is essentially nonexistent.

In this paper, we propose classes of informative prior distributions for time series count data. The prior specification is based on the notion of specifying an
$n_{0} \times 1$ vector $y_{0}$ of prior predictions for the response vector, $y$, of the current study, along with a covariate matrix $X_{0}$ corresponding to $y_{0}$. Then $\left(n_{0}, y_{0}, X_{0}\right)$ are used to specify an automated parametric informative prior for the regression coefficients. The quantity $y_{0}$ can be taken as the raw response vector from the historical data, a vector of fitted values based on the historical data, a vector obtained from a theoretical prediction model, or a vector specified from expert opinion or case-specific information. Thus $y_{0}$ can be viewed as a prior "prediction" for $y$, the actual data in the current study. Similarly, $X_{0}$ can be taken as the raw covariate matrix based on the historical data or it can be specified in other ways. In any case, taking $\left(n_{0}, y_{0}, X_{0}\right)$ to be the raw historical data results in a more natural, interpretable, and automated specification. The Monte Carlo methods we propose will facilitate a very fast and efficient way of computing posterior model probabilities using only a single posterior sample from a single model, that being the full model. Such a procedure has proved to be quite feasible and powerful in the model selection context (see for example, Chen, Ibrahim and Yiannoutsos, (1999)). In addition, our proposed informative prior elicitation schemes allow us to incorporate historical data in a natural way.

## 2. The Method

### 2.1. The likelihood function

Let $\mathcal{M}$ denote the model space. We enumerate the models in $\mathcal{M}$ by $m=$ $1, \ldots, \mathcal{K}$, where $\mathcal{K}$ is the dimension of $\mathcal{M}$ and model $\mathcal{K}$ denotes the full model. The full model is defined here as the model containing all of the available covariates in the study. Further, let $I$ denote a model indicator, so that $I=m$ means that model $m$ is selected. If $k$ is the number of covariates for the full model, our model space contains $2^{k}$ models. Let $\beta^{(\mathcal{K})}=\left(\beta_{0}, \ldots, \beta_{k}\right)^{\prime}$ denote the regression coefficients for the full model including an intercept, and let $\beta^{(m)}$ denote a $k_{m} \times 1$ vector of regression coefficients for model $m$ with an intercept, and a specific choice of $k_{m}-1$ covariates. We write $\beta^{(\mathcal{K})}=\left(\beta^{(m)^{\prime}}, \beta^{(-m)^{\prime}}\right)^{\prime}$, where $\beta^{(-m)}$ is $\beta^{(\mathcal{K})}$ with $\beta^{(m)}$ deleted.

Consider a time series of counts $y_{t}, t=1, \ldots, n$, where each $y_{t}$ has a corresponding $k_{m} \times 1$ covariate vector $x_{t}^{(m)}$ under model $m$. Under model $m$, conditional on $\beta^{(m)}$ and a stationary unobserved process $\epsilon_{t}$, the $y_{t}$ 's are assumed to be independent discrete random variables from a distribution in the exponential family. This leads to the conditional density

$$
\begin{align*}
& p\left(y \mid \beta^{(m)}, \epsilon, I=m\right)=\prod_{t=1}^{n} p\left(y_{t} \mid \beta^{(m)}, \epsilon_{t}\right) \\
= & \prod_{t=1}^{n} \exp \left\{\tau_{t}^{-1}\left[y_{t} \theta\left(x_{t}^{(m)}, \beta^{(m)}, \epsilon_{t}\right)-q\left(\theta\left(x_{t}^{(m)}, \beta^{(m)}, \epsilon_{t}\right)\right)\right]-c\left(y_{t}\right)\right\}, \tag{2.1}
\end{align*}
$$

indexed by the canonical parameter $\theta_{t} \equiv \theta\left(x_{t}^{(m)}, \beta^{(m)}, \epsilon_{t}\right)$ and the scale parameter $\tau_{t}$, where $y=\left(y_{1}, \ldots, y_{n}\right)^{\prime}$, and $\epsilon=\left(\epsilon_{1}, \ldots, \epsilon_{n}\right)^{\prime}$. Further suppose $\theta\left(x_{t}^{(m)}, \beta^{(m)}, \epsilon_{t}\right)$ satisfies the equation

$$
\begin{equation*}
\theta\left(x_{t}^{(m)}, \beta^{(m)}, \epsilon_{t}\right)=h\left(\left(x_{t}^{(m)}\right)^{\prime} \beta^{(m)}+\epsilon_{t}\right), \quad t=1, \ldots, n \tag{2.2}
\end{equation*}
$$

where $h$ is a monotonic differentiable function, often referred to as the link function. In (2.1), the functions $q$ and $c$ determine a particular family in the class, such as the binomial or Poisson distributions. For example, if we take $y_{t}$ to have a Poisson distribution with conditional mean $\lambda_{t}=\exp \left(\left(x_{t}^{(m)^{\prime}} \beta^{(m)}+\right.\right.$ $\left.\epsilon_{t}\right)$, then $\tau_{t}=1, h\left(\left(x_{t}^{(m)}\right)^{\prime} \beta^{(m)}+\epsilon_{t}\right)=\left(x_{t}^{(m)}\right)^{\prime} \beta^{(m)}+\epsilon_{t}, q\left(\theta\left(x_{t}^{(m)}, \beta^{(m)}, \epsilon_{t}\right)\right)=$ $\exp \left\{\left(x_{t}^{(m)}\right)^{\prime} \beta^{(m)}+\epsilon_{t}\right\}$, and $c\left(y_{t}\right)=\ln \left(y_{t}!\right)$. We emphasize here that the likelihood in (2.1) is a general exponential family model for discrete outcomes, with the Poisson model being a special case. For ease of exposition, we assume $\tau_{t}=1$ throughout, since this is in fact the case for many models in the exponential family, including the binomial and the Poisson. In addition, it will be convenient to write (2.1) in vector notation as
$p\left(y \mid \beta^{(m)}, \epsilon, I=m\right)=\exp \left\{y^{\prime} \theta\left(X^{(m)}, \beta^{(m)}, \epsilon\right)-J_{n}^{\prime} Q\left(X^{(m)}, \beta^{(m)}, \epsilon\right)-J_{n}^{\prime} C(y)\right\}$,
where $X^{(m)}$ is the $n \times k_{m}$ matrix of covariates with $t$ th row equal to $\left(x_{t}^{(m)}\right)^{\prime}, J_{n}$ is an $n \times 1$ vector of ones, $\theta\left(X^{(m)}, \beta^{(m)}, \epsilon\right), Q\left(X^{(m)}, \beta^{(m)}, \epsilon\right)$, and $C(y)$ are $n \times 1$ vectors with the $t$ th components equal to $\theta_{t}=h\left(\left(x_{t}^{(m)}\right)^{\prime} \beta^{(m)}+\epsilon_{t}\right), q_{t}=q\left(\theta_{t}\right)$, and $c_{t}=c\left(y_{t}\right)$, respectively.

The latent process $\epsilon_{t}$ is assumed to have a normal distribution with mean 0 . We assume an $\operatorname{AR}(1)$ structure for the covariance matrix of $\epsilon$. This structure is well motivated in the statistical literature and is one of the most widely used in the time series setting (see Zeger (1988)). It proves to be quite appropriate for our purposes, as demonstrated in Section 4. Thus we assume that $\epsilon$ has a multivariate normal distribution with mean 0 and covariance matrix $\sigma^{2} \Sigma$, where the $(i, j)$ th element of $\Sigma$ has the form $\sigma_{i j}=\rho^{|i-j|},-1 \leq \rho \leq 1$. The unobserved process $\epsilon_{t}$ is analogous to a "random effect" in a random effects model except for the correlation. We note that the mean and variance of $\epsilon_{t}$ do not depend on $t$. Zeger (1988) constructs a similar model for Poisson count data through the mean and covariance of the latent process, which then define the estimating equations. He does not specify a parametric distribution for the latent process as is done here, and only considers Poisson count models.

Let $\phi_{n}\left(\epsilon \mid \mu, \sigma^{2} \Sigma\right)$ denote the $n$-dimensional normal density of the latent process $\epsilon$ with mean $\mu$ and covariance matrix $\sigma^{2} \Sigma$, i.e.,

$$
\begin{equation*}
\phi_{n}\left(\epsilon \mid \mu, \sigma^{2} \Sigma\right)=\left(2 \pi \sigma^{2}\right)^{-n / 2}|\Sigma|^{-1 / 2} \exp \left\{-\frac{1}{2 \sigma^{2}}(\epsilon-\mu)^{\prime} \Sigma^{-1}(\epsilon-\mu)\right\} \tag{2.4}
\end{equation*}
$$

We note that $|\Sigma|=\left(1-\rho^{2}\right)^{n-1}$. Then the joint density of $(y, \epsilon)$ can be written as

$$
\begin{align*}
& p\left(y, \epsilon \mid \beta^{(m)}, \sigma^{2}, \rho, I=m\right) \\
= & \exp \left\{y^{\prime} \theta\left(X^{(m)}, \beta^{(m)}, \epsilon\right)-J_{n}^{\prime} Q\left(X^{(m)}, \beta^{(m)}, \epsilon\right)-J_{n}^{\prime} C(y)\right\} \times \phi_{n}\left(\epsilon \mid 0, \sigma^{2} \Sigma\right) . \tag{2.5}
\end{align*}
$$

To induce the correlation structure on $y$ we integrate out $\epsilon$ from (2.5), leading to the "marginal" likelihood of $\beta^{(m)}$ given by

$$
\begin{equation*}
p\left(y \mid \beta^{(m)}, \sigma^{2}, \rho, I=m\right)=\int p\left(y, \epsilon \mid \beta^{(m)}, \sigma^{2}, \rho, I=m\right) d \epsilon \tag{2.6}
\end{equation*}
$$

where $p\left(y, \epsilon \mid \beta^{(m)}, \sigma^{2}, \rho, I=m\right)$ is given by (2.5). The marginal likelihood of $\beta^{(m)}$ in (2.6) does not have a closed form.

The implications of the process $\epsilon_{t}$ on the correlation structure in the $y_{t}$ 's and the regression model is as follows. Note first that $\epsilon_{t}^{*}=\exp \left(\epsilon_{t}\right)$ has a log-normal distribution with mean $\alpha=\exp \left(\frac{1}{2} \sigma^{2}\right)$ and variance $\nu^{2}=\exp \left(2 \sigma^{2}\right)-\exp \left(\sigma^{2}\right)$. The unobserved process $\epsilon_{t}$ allows for overdispersion and autocorrelation in $y_{t}$. In addition, the degree of overdispersion depends on the marginal mean of $y_{t}$. For the Poisson model, the autocorrelation in $y_{t}$ must be less than or equal to that in $\epsilon_{t}$, and the degree of autocorrelation in $y_{t}$ relative to $\epsilon_{t}$ decreases as the marginal mean of $y_{t}$ and $\nu^{2}$ decrease.

### 2.2. The prior distributions

Informative prior elicitation is an important part of a Bayesian analysis. This is especially true for the problem of variable subset selection, since proper prior distributions are required for the computation of posterior model probabilities. We propose a class of informative priors for the regression coefficients $\beta^{(m)}$, since these parameters are of primary inferential interest in the variable selection problem. Our prior construction for $\beta^{(m)}$ is based on the availability of historical data, as suggested in Section 1. Suppose there are $N$ historical data sets and the sample size of the $i$ th historical study is $n_{0 i}$. Let $y_{0 i}$ denote the $n_{0 i} \times 1$ vector of time series counts for the $i$ th historical study, and let $X_{0 i}^{(m)}$ denote the $n_{0 i} \times k_{m}$ matrix of covariates corresponding to the $i$ th historical study. In addition let $\epsilon_{0 i}$ denote the latent process for the $i$ th historical study, where $\epsilon_{0 i}$ is an $n_{0 i} \times 1$ vector, $i=1, \ldots, N$, and $\epsilon_{0 i}$ has an $n_{0 i}$ dimensional multivariate normal distribution with mean 0 and covariance matrix $\sigma^{2} \Sigma_{0 i}$, where $\Sigma_{0 i}$ is an $n_{0 i} \times n_{0 i}$ matrix with $\left(j, j^{*}\right)$ th element equal to $\rho^{\left|j-j^{*}\right|}$. Finally let $y_{0}=\left(y_{01}, \ldots, y_{0 N}\right)$ denote the response vector for all the historical studies. Throughout, we assume that given the model parameters, $y_{0}$ and $y$ are independent.

We propose a prior distribution for $\beta^{(m)}$ taking the form

$$
\begin{align*}
& \pi\left(\beta^{(m)} \mid \sigma^{2}, \rho, y_{0 i}, a_{0 i}, I=m\right) \\
\propto & \prod_{i=1}^{N} \int p\left(y_{0 i} \mid \beta^{(m)}, \epsilon_{0 i}, I=m\right)^{a_{0 i}} \phi_{n_{0 i}}\left(\epsilon_{0 i} \mid 0, \sigma^{2} \Sigma_{0 i}\right) d \epsilon_{0 i} \tag{2.7}
\end{align*}
$$

where $a_{0 i}$ is a scalar prior parameter that controls the weight of the $i$ th historical study relative to the likelihood of the current study. Small values of $a_{0 i}$ give less weight whereas large values give more weight. It is most sensible to restrict $a_{0 i}$ to $0 \leq a_{0 i} \leq 1$, since we do not want to weight the historical data more than the current data. The parameter $a_{0 i}$ can also be interpreted as a precision parameter which takes into account the between and within study variability in the historical data sets.

Let $a_{0}=\left(a_{01}, \ldots, a_{0 N}\right)$. The prior specification is completed by specifying priors for $\left(\sigma^{2}, \rho, a_{0}\right)$. We take these parameters to be independent a priori. We specify an inverse gamma prior for $\sigma^{2}$, denoted $\operatorname{IG}\left(\delta_{0}, \gamma_{0}\right)$, a scaled beta prior for $\rho$, denoted $\operatorname{scbeta}\left(\nu_{0}, \psi_{0}\right)$, and independent identically distributed beta priors for each $a_{0 i}$, denoted $\operatorname{beta}\left(\alpha_{0}, \lambda_{0}\right)$. Here, $\left(\delta_{0}, \gamma_{0}, \nu_{0}, \psi_{0}, \alpha_{0}, \lambda_{0}\right)$ are specified prior hyperparameters. Thus, we propose a joint prior distribution of the form

$$
\begin{equation*}
\pi\left(\beta^{(m)}, \sigma^{2}, \rho, a_{0} \mid y_{0}, I=m\right) \propto p_{0}^{*}\left(\beta^{(m)}, \sigma^{2}, \rho, a_{0} \mid y_{0}, I=m\right) \tag{2.8}
\end{equation*}
$$

where $p_{0}^{*}\left(\beta^{(m)}, \sigma^{2}, \rho, a_{0} \mid y_{0}, I=m\right)$ is the unnormalized prior density defined by

$$
\begin{align*}
& p_{0}^{*}\left(\beta^{(m)}, \sigma^{2}, \rho, a_{0} \mid y_{0}, I=m\right) \\
= & \prod_{i=1}^{N}\left(\int p\left(y_{0 i} \mid \beta^{(m)}, \epsilon_{0 i}, I=m\right)^{a_{0 i}} \phi_{n_{0 i}}\left(\epsilon_{0 i} \mid 0, \sigma^{2} \Sigma_{0 i}\right) d \epsilon_{0 i}\right) \\
& \times\left(\prod_{i=1}^{N} a_{0 i}^{\alpha_{0}-1}\left(1-a_{0 i}\right)^{\lambda_{0}-1}\right) \times\left(\sigma^{2}\right)^{-\left(\delta_{0}+1\right)} \exp \left(-\sigma^{-2} \gamma_{0}\right)(1+\rho)^{\nu_{0}-1}(1-\rho)^{\psi_{0}-1} .( \tag{2.9}
\end{align*}
$$

Our joint prior for $\left(\beta^{(m)}, \sigma^{2}, \rho, a_{0}\right)$ clearly does not have a closed form in general. However, it has a natural motivation and several appealing interpretations. One motivation for the prior in (2.9) is that, by taking $a_{0}$ random, the tails of the marginal prior distribution for $\beta^{(m)}$ are heavier than those obtained by taking $a_{0 i}$ to be a fixed hyperparameter. In addition, a prior on $a_{0}$ provides flexibility and allows us to express our uncertainty about it. By allowing different $a_{0 i}$ 's for different historical studies, we are able to develop a more flexible prior that can weight each historical study differently. This would certainly be desirable if one historical study has a much larger sample size than others. Another motivation for $(2.9)$ is that it mimics the marginal likelihood function of $\beta^{(m)}$ based on the historical data. If for example $a_{0 i}=1$, then (2.9) is precisely the marginal
likelihood function of $\beta^{(m)}$ based on the historical data. Thus, our prior can be viewed as a weighted marginal likelihood of $\beta^{(m)}$. This seems like a natural prior when historical data is available.

To show the propriety of the prior distribution given by (2.9), we first introduce a useful lemma.

Lemma 2.1. Let $\alpha_{0}>0, \lambda_{0}>0$. There exists $K=K\left(\alpha_{0}, \lambda_{0}\right)>0$ such that $\forall 0 \leq \xi \leq 1$,

$$
\begin{equation*}
\int_{0}^{1} \xi^{a_{0 i}} a_{0 i}^{\alpha_{0}-1}\left(1-a_{0 i}\right)^{\lambda_{0}-1} d a_{0 i} \leq K(1+\ln (1 / \xi))^{-\alpha_{0}} \tag{2.10}
\end{equation*}
$$

The proof of the lemma is given in the Appendix.
Let $y_{0 i t}$ denote the $t$ th component of $y_{0 i}$ and let $\left(x_{0 i t}^{(m)}\right)^{\prime}$ denote the $t$ th row of $X_{0 i}^{(m)}$. Using Lemma 2.1, we obtain the following result, which ensures the propriety of the joint prior distribution $\pi\left(\beta^{(m)}, \sigma^{2}, \rho, a_{0} \mid y_{0}, I=m\right)$.
Theorem 2.1. Assume that

$$
\begin{equation*}
\exp \left\{\left(y_{0 i t} \theta_{0 i t}-q\left(\theta_{0 i t}\right)\right)-c\left(y_{0 i t}\right)\right\} \leq M \tag{2.11}
\end{equation*}
$$

for $t=1, \ldots, n_{0 i}, i=1, \ldots, N$, where $M$ is some finite constant. Suppose there exist $y_{0 i t_{i 1}}, y_{0 i t_{i 2}}, \ldots, y_{0 i t_{i k_{m}}}\left(1 \leq t_{i 1} \leq t_{i 2} \leq \cdots \leq t_{i k_{m}}\right)$ such that

$$
\begin{equation*}
\int_{-\infty}^{\infty} e^{d_{0}|\eta|} \exp \left\{\left(y_{0 i t_{j}} h(\eta)-q(h(\eta))\right\} d \eta<\infty\right. \tag{2.12}
\end{equation*}
$$

for some $d_{0}>0$ and $j=1, \ldots, k_{m}$, and that the corresponding design matrix $\left(x_{0 i t_{1}}^{(m)}, x_{0 i t_{2}}^{(m)}, \ldots, x_{0 i t_{k_{m}}}^{(m)}\right)^{\prime}$ has full rank $k_{m}$. Then if $\alpha_{0}>k_{m} / N, \lambda_{0}>0$, and (2.12) holds, the joint prior distribution $\pi\left(\beta^{(m)}, \sigma^{2}, \rho, a_{0} \mid y_{0}, I=m\right)$ is proper.

The proof of Theorem 2.1 is given in the Appendix.
For elicitation purposes, it is easier to work with the prior mean and variance of $a_{0 i}$, given by $\mu_{a_{0}}=\alpha_{0} /\left(\alpha_{0}+\lambda_{0}\right)$ and $\sigma_{a_{0}}^{2}=\mu_{a_{0}}\left(1-\mu_{a_{0}}\right)\left(\alpha_{0}+\lambda_{0}+1\right)^{-1}$. From Theorem 2.1, a sufficient condition for the propriety of the prior distribution is that $\alpha_{0}>(k+1) / N$ for the full model. Therefore, a reasonable starting point for the analysis is to choose $\alpha_{0}=\lambda_{0}=(k+2) / N$, which gives $\mu_{a_{0}}=1 / 2$. Then we conduct several sensitivity analyses within a suitable range of the uniform prior, using various values of $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)$. We do not recommend doing an analysis based on one set of proposed values of $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)$.

### 2.3. Prior distribution on the model space

Let

$$
\begin{equation*}
p_{0}^{*}\left(\beta^{(m)} \mid y_{0}, I=m\right)=\int p_{0}^{*}\left(\beta^{(m)}, \sigma^{2}, \rho, a_{0} \mid y_{0}, I=m\right) d \sigma^{2} d \rho d a_{0} \tag{2.13}
\end{equation*}
$$

where $p_{0}^{*}\left(\beta^{(m)}, \sigma^{2}, \rho, a_{0} \mid y_{0}, I=m\right)$ is given by (2.9). We see that $p_{0}^{*}\left(\beta^{(m)} \mid y_{0}, I=\right.$ $m$ ) is proportional to the marginal prior of $\beta^{(m)}$. We propose to take the prior probability of model $m$, denoted $p(I=m)$, as

$$
\begin{equation*}
p(I=m)=\frac{\int p_{0}^{*}\left(\beta^{(m)} \mid y_{0}, I=m\right) d \beta^{(m)}}{\sum_{j=1}^{\mathcal{K}} \int p_{0}^{*}\left(\beta^{(j)} \mid y_{0 j}, I=j\right) d \beta^{(j)}} \tag{2.14}
\end{equation*}
$$

The choice of $p(I=m)$ in (2.14) is a natural one since the numerator is just the normalizing constant of the joint prior of $\left(\beta^{(m)}, a_{0}, \sigma^{2}, \rho\right)$ under model $m$. The prior model probabilities in (2.14) are based on coherent Bayesian updating and this fact has several attractive interpretations. First $p(I=m)$ in (2.14) corresponds to the posterior probability of model $m$ based on the data $y_{0}$ under model $m$, using a uniform prior for the previous study, i.e., $p_{0}(I=m)=2^{-k}$ for $m \in \mathcal{M}$ as $\alpha_{0} \rightarrow \infty$. We also note that as $\alpha_{0} \rightarrow \infty, a_{0 i} \rightarrow 1$ with probability 1. Second as $\lambda_{0} \rightarrow \infty, p(I=m)$ reduces to a uniform prior on the model space. Therefore, as $\lambda_{0} \rightarrow \infty$, the historical data $y_{0}$ have a minimal impact in determining $p(m)$. In addition, as $\lambda_{0} \rightarrow \infty, a_{0 i} \rightarrow 0$ with probability 1. Finally, we mention that the choice of $p(I=m)$ in (2.14) greatly eases the computational burden for calculating posterior model probabilities. There is more detailed discussion in Section 3.

## 3. Posterior Model Probabilities

In this section we explore the theoretical properties of posterior model probabilities based on the choice of the prior model probabilities of (2.14), and then propose novel Monte Carlo implementation procedures to compute posterior model probabilities. A key result is a formula for the posterior model probability that does not depend directly on $p(I=m)$.

The posterior probability of model $m$ is given by

$$
\begin{equation*}
p(I=m \mid y)=\frac{p(y \mid I=m) p(I=m)}{\sum_{j=1}^{\mathcal{K}} p(y \mid I=j) p(I=j)} \tag{3.1}
\end{equation*}
$$

where $p(y \mid I=m)$ denotes the marginal distribution of the data under model $m$ for the current study, and $p(I=m)$ denotes the prior probability of model $m$ in (2.14).

Theorem 3.1. $p(I=m \mid y)$ in (3.1) is given by

$$
\begin{equation*}
p(I=m \mid y)=\frac{p\left(\beta^{(-m)}=0 \mid y, y_{0}, I=\mathcal{K}\right)}{\sum_{j=1}^{\mathcal{K}} p\left(\beta^{(-j)}=0 \mid y, y_{0}, I=\mathcal{K}\right)}, \tag{3.2}
\end{equation*}
$$

$m=1, \ldots, \mathcal{K}$, where $p\left(\beta^{(-m)}=0 \mid y, y_{0}, I=\mathcal{K}\right)$ is the marginal posterior density of $\beta^{(-m)}$ evaluated at $\beta^{(-m)}=0$.

In (3.2), for notational convenience, we assume $p\left(\beta^{(-\mathcal{K})}=0 \mid y, y_{0}, I=\mathcal{K}\right)=$ 1. The proof of Theorem 3.1 is given in the Appendix. We mention here that the derivation of (3.2) assumes that $y_{0}$ and $y$ are independent given the model parameters. The result in (3.2) is very attractive since it shows that the posterior model probability $p(I=m \mid y)$ is simply a function of the marginal posterior density functions of $\beta^{(-m)}$ for the full model evaluated at $\beta^{(-m)}=0$. This is an important feature since it allows us to compute the posterior model probabilities directly without numerically computing the prior model probabilities. We note that this computational device works best if all of the covariates are standardized to have mean 0 and variance 1 . This is not restrictive since this transformation is used quite often in practice to numerically stabilize the Gibbs sampler and adaptive rejection algorithms.

Due to the complexity of our model, the analytical evaluation of $p\left(\beta^{(-m)}=\right.$ $\left.0 \mid y, y_{0}, I=\mathcal{K}\right)$ does not appear possible. Therefore, we propose a novel Monte Carlo method to compute posterior model probabilities using a single MCMC sample from the full model. The hierarchical centering reparameterization technique of Gelfand, Sahu and Carlin (1996) is particularly suitable for the implementation of MCMC sampling for our problem. This is due to the fact that (2.9) leads to (3.2). From (3.2), it is easy to see that the posterior model probability is proportional to the marginal posterior density evaluated at 0 . Thus, computing the posterior model probability is essentially equivalent to estimating the marginal posterior density. It is well known that the hierarchical centering technique is very useful in developing an efficient Monte Carlo method for estimating the marginal posterior density $p\left(\beta^{(-m)}=0 \mid y, y_{0}, I=\mathcal{K}\right)$ and this leads to the efficient computation of (3.2). To the best of our knowledge, this is the first time that the hierarchical centering reparameterization technique has been used to ease the computational burden in Bayesian variable selection.

To this end, consider the following reparameterization:

$$
\begin{equation*}
\eta=\epsilon+X^{(\mathcal{K})} \beta^{(\mathcal{K})} \tag{3.3}
\end{equation*}
$$

and

$$
\begin{equation*}
\eta_{0 i}=\epsilon_{0 i}+X_{0 i}^{(\mathcal{K})} \beta^{(\mathcal{K})} \tag{3.4}
\end{equation*}
$$

for $i=1, \ldots, N$. Let $\eta_{0}=\left(\eta_{01}, \ldots, \eta_{0 N}\right)$. Then we write the reparameterized posterior for the full model, $p\left(\beta^{(\mathcal{K})}, \sigma^{2}, \rho, a_{0}, \eta, \eta_{0} \mid y, y_{0}, I=\mathcal{K}\right)$. Using the hierarchical centering technique, we obtain an MCMC sample $\left\{\left(\beta_{(l)}^{(\mathcal{K})}, \sigma_{(l)}^{2}, \rho_{(l)}, a_{0(l)}, \eta_{(l)}\right.\right.$, $\left.\left.\eta_{0(l)}\right), \quad l=1, \ldots, L\right\}$ from this reparameterized posterior. Then, following the
lines of Chen (1994), $p\left(\beta^{(-m)}=0 \mid y, y_{0}, I=\mathcal{K}\right)$ can be estimated by the conditional marginal density estimation (CMDE) method. Gelfand, Smith, and Lee (1992), Chen (1994), and Chen and Shao (1997) have shown that the CMDE is the most efficient Monte Carlo method for estimating marginal posterior densities when a joint posterior density is known up to a normalizing constant. It directly follows from Chen and Shao (1997) that a simulation consistent estimator of $p\left(\beta^{(-m)}=0 \mid y, y_{0}, I=\mathcal{K}\right)$ is given by

$$
\begin{equation*}
\hat{p}\left(\beta^{(-m)}=0 \mid y, y_{0}, I=\mathcal{K}\right)=\frac{1}{L} \sum_{l=1}^{L} N_{k+1-k_{m}}\left(\beta^{(-m)}=0 \mid \beta_{(l)}^{(-m)}, \hat{\beta}_{(l)},\left(B_{(l)}\right)^{-1}\right) \tag{3.5}
\end{equation*}
$$

where $N_{k+1-k_{m}}\left(\beta^{(-m)}=0 \mid \beta_{(l)}^{(-m)}, \hat{\beta}_{(l)}^{(\mathcal{K})},\left(B_{(l)}\right)^{-1}\right)$ is the $\left(k+1-k_{m}\right)$-dimensional conditional normal density function of $N_{k+1}\left(\hat{\beta}_{(l)}^{(\mathcal{K})},\left(B_{(l)}\right)^{-1}\right)$ given $\beta_{(l)}^{(m)}$ evaluated at $\beta^{(-m)}=0$,

$$
\begin{gathered}
B_{(l)}=\frac{1}{\sigma_{(l)}^{2}}\left(\left(X^{(\mathcal{K})}\right)^{\prime} \Sigma_{(l)}^{-1} X^{(\mathcal{K})}+\sum_{i=1}^{N}\left(X_{0 i}^{(\mathcal{K})}\right)^{\prime} \Sigma_{0 i(l)}^{-1} X_{0 i}^{(\mathcal{K})}\right) \\
\hat{\beta}^{(\mathcal{K})}=\frac{1}{\sigma_{(l)}^{2}}\left\{B_{(l)}^{-1}\left(\left(X^{(\mathcal{K})}\right)^{\prime} \Sigma_{(l)}^{-1} \eta_{(l)}+\sum_{i=1}^{N}\left(X_{0 i}^{(\mathcal{K})}\right)^{\prime} \Sigma_{0 i(l)}^{-1} \eta_{0 i(l)}\right)\right\}
\end{gathered}
$$

$\Sigma_{(l)}$ is an $n \times n$ matrix with $\left(j, j^{*}\right) t h$ element equal to $\rho_{(l)}^{\left|j-j^{*}\right|}$, and $\Sigma_{0 i(l)}$ is an $n_{0 i} \times n_{0 i}$ matrix with $\left(j, j^{*}\right) t h$ element equal to $\rho_{(l)}^{\left|j-j^{*}\right|}$.

There are several advantages of the above Monte Carlo procedure. First, as previously mentioned, it is not required to compute $p(I=m)$ for each model. Second, we need only one random draw from $p\left(\beta^{(\mathcal{K})}, \sigma^{2}, \rho, a_{0}, \eta, \eta_{0} \mid y, y_{0}, I=\mathcal{K}\right)$. Third, after we obtain an MCMC sample from the posterior distribution of the full model, calculating $\hat{p}\left(\beta^{(-m)}=0 \mid y, y_{0}, I=\mathcal{K}\right)$ given by (3.5) is straightforward and almost free of computational time. Fourth, for the purposes of computing posterior model probabilities, one need only store a $(k+1)$-dimensional vector $\hat{\beta}_{(l)}$ and a $(k+1) \times(k+1)$ matrix $B_{(l)}$ for each MCMC sampling iteration. This becomes even more advantageous for cases where multiple previous studies are available and each $n_{0 i}$ is large. The above features of our Monte Carlo procedure make Bayesian variable selection feasible in the presence of a large number of covariates (say, $k>20$ ).

We note here that Bayesian inference for non-Gaussian time series models has been considered by Shephard and Pitt (1997), but their model and computational development is quite different from what we do here. Similarly, there are several differences between our methodology and that of Chen, Ibrahim and Yiannoutsos (1999) .

## 4. Examples

## Example 1. Simulation study

In this example, we demonstrate variable subset selection with our proposed methodology using a simulated data set. We also demonstrate the computational feasibility of our methods.

The data for the current study is generated as follows. We generate $n=50$ independent observations from a Poisson distribution each with mean $\lambda_{t}=$ $\exp \left(\beta_{0}+\beta_{1} x_{t 1}+\beta_{2} x_{t 2}+\epsilon_{t}\right), t=1, \ldots, 50$, and $\beta=\left(\beta_{0}, \beta_{1}, \beta_{2}\right)=(1,1,-1)$. Also, $\left(x_{t 1}, x_{t 2}\right)^{\prime}$ are generated as i.i.d. bivariate normal random vectors with mean $(0.5,0.5)^{\prime}$ and covariance matrix $\left(\begin{array}{cc}0.5 & 0.25 \\ 0.25 & 0.5\end{array}\right)$. We take $\epsilon=\left(\epsilon_{1}, \ldots, \epsilon_{50}\right)^{\prime}$ to have a 50 -dimensional multivariate normal distribution with mean 0 and covariance matrix $\sigma^{2} \Sigma$, where $\sigma^{2}=1$ and the $(i, j)$ th element of $\Sigma$ is $(0.5)^{|i-j|}$. In addition, we generate two other covariates $\left(x_{t 3}, x_{t 4}\right)$ which are i.i.d. normal random variables each with mean 0 and variance .5 , and independent of ( $x_{t 1}, x_{t 2}$ ). Thus, the true model contains the two covariates $\left(x_{1}, x_{2}\right)$, and the "full" model contains the four covariates $\left(x_{1}, \ldots, x_{4}\right)$. Therefore, our model space $\mathcal{M}$ contains 16 models, with an intercept included in each. The historical data were generated in a similar fashion. We take $n_{0}=50$, and $\sigma^{2}=1.2$, with all other parameters the same as for the current data. In addition, $\left(x_{0 t 1}, \ldots, x_{0 t 4}\right)$ are generated in exactly the same way as the current data.

Table 1 shows posterior probabilities for the top model under various choices of $\mu_{a_{0}}, \sigma_{a_{0}}$ and $N$. The posterior model probability is denoted by $p(I=m \mid y)$ in Table 1. From Table 1, under each choice of ( $\mu_{a_{0}}, \sigma_{a_{0}}$ ) and $N$, the true model $\left(x_{1}, x_{2}\right)$ obtains the largest posterior probability. Although not shown in Table 1 , the ( $x_{1}, x_{2}, x_{3}$ ) model is consistently chosen as the second best model under the choices of ( $\mu_{a_{0}}, \sigma_{a_{0}}$ ) and $N$ given in Table 1. In addition, the order of the models with respect to their posterior probabilities is preserved as $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)$ are varied according to Table 1. We do not see a dramatic change in the posterior model probabilities as $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)$ are varied for a given $N$. It is apparent that the posterior model probabilities are quite robust with respect to changes in $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)$. In addition, we observe a monotonic increase in the posterior model probabilities as we assign more weight to the historical data, i.e., as $\mu_{a_{0}}$ increases and/or $\sigma_{a_{0}}$ decreases. Specifically, for $\mu_{a_{0}}=0.5$ a monotonic increase in the posterior model probability is observed as $\sigma_{a_{0}}$ is decreased. This is a desirable feature since it shows that a heavier weight given to the historical data results in an increase in the posterior probability of the true model. Also, a low weight for the historical data with $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)=(0.09,0.04)$ still yields $\left(x_{1}, x_{2}\right)$ as the top model with posterior probability of 0.349 . In addition, more extreme values
of $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)$ were used. For example, using an extremely low weight for the historical data with $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)=(0.05,0.02)$, the true model, $\left(x_{1}, x_{2}\right)$, obtains the largest posterior probability of 0.347 and the $\left(x_{1}, x_{2}, x_{3}\right)$ model obtains the second largest posterior probability. We note that $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)=(0.05,0.02)$ represents the smallest weight we can place on the historical data and still obtain proper priors (see Theorem 2.1). Similar results are obtained for other small values of $\mu_{a_{0}}$ and moderate to large values of $\sigma_{a_{0}}$. We note here that we cannot do Bayesian variable selection using $a_{0 i}=0$ with probability 1 , since this would result in an improper prior for $\beta$ and, in this case, the posterior model probabilities would not be well defined. When historical data is not available, one can specify a normal prior for $\beta$, and the specification of the hyperparameters would be based on elicitation from expert opinion or case-specific information.

Table 1 also indicates that a monotonic increase in the posterior probability of the true model occurs as $N$ is increased. This is a solid feature of our methodology since it shows that increasing the number of historical studies provides more precise estimates of the posterior model probabilities. Similar results were obtained for other combinations for $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)$. Finally, an analysis using $\rho=0$ and $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)=(0.5,0.06)$ was conducted, and results very similar to those of Table 1 were obtained. Specifically, the model obtaining the largest posterior probability is the $\left(x_{1}, x_{2}\right)$ model with value 0.414 , which is similar to the one using $\rho=0.5$. In addition, the model with the second largest posterior probability is $\left(x_{1}, x_{2}, x_{4}\right)$ with value 0.243 , and the model with the third largest posterior probability is $\left(x_{1}, x_{2}, x_{3}\right)$ with value 0.215 .

## Example 2. Pollen data

Ragweed pollen data were collected daily in Kalamazoo, Michigan from 1991 to 1994. Frequentist analyses of these data using standard Poisson regression methods have been conducted by Stark, Ryan, McDonald and Burge (1997). Our aim here is to demonstrate our Bayesian methodology for variable selection. The response variable $y$ is the pollen count for a particular day in the season for a given year. Initially, we take the 1991, 1992, and 1993 data $(N=3)$ as the historical data and the 1994 data as the current data. The data for each year was collected roughly over a 3 month interval between the months of July and October. However, for each year, the first and last observations were collected on different days. For example, in 1991, the first observation was collected on July 28 and the last was collected on October 27th. In 1992, the first observation was collected on August 6th and the last observation on October 26th.

The full model contains an intercept and seven covariates, extensively discussed and motivated by Stark et al. (1997). These are $x_{1}=$ rain, (which is a binary variable taking the value 0 if there were at least three hours of steady rain,
and 1 otherwise), $x_{2}=$ day in the pollen season , $x_{3}=\log$ (day). Two covariates are functions of temperature. These are $x_{4}$ which is the lowess smoothed function of temperature constructed from a non-parametric estimate of the regression of pollen count on average temperature, and $x_{5}$, which denotes the deviation from the daily averages temperature to the lowess line. The final two covariates are $x_{6}=$ windspeed and $x_{7}=$ cold (a binary variable taking the value 0 if the overnight temperature dropped below 50 degrees Fahrenheit, and 1 otherwise).

The model space $\mathcal{M}$ contains $2^{7}$ models. We specify noninformative priors for $\rho$ and $\sigma^{2}$. Specifically, we take a uniform prior for $\rho$ on $[-1,1]$ (i.e. $\nu_{0}=\psi_{0}=1$ ) and take $\sigma^{2} \sim I G(.005, .005)$. Table 2 give results for the model with the largest posterior probability based on several values of $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)$. The top model in each case is $\left(x_{1}, x_{2}, x_{3}, x_{4}, x_{5}\right)$, for all combinations of $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)$ (and $\left.N\right)$. In addition, we see that the posterior model probabilities increase monotonically as more weight is given to the historical data. For example, using $N=3$, when we put very small weight on the historical data, such as $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)=(0.009,0.003)$, the $\left(x_{1}, x_{2}, x_{3}, x_{4}, x_{5}\right)$ model still obtains the largest posterior probability, with value 0.117 . When we put extremely small weight on the historical data such as $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)=(0.0009,0.0003)$, the $\left(x_{2}, x_{3}, x_{4}, x_{5}, x_{7}\right)$ model obtains the largest posterior probability, with value 0.122 and the $\left(x_{1}, x_{2}, x_{3}, x_{4}, x_{5}\right)$ model obtains the fourth largest posterior probability with value 0.101 . When we put a small weight on the historical data using a moderate variance, $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)=(0.09,0.027)$, the $\left(x_{1}, x_{2}, x_{3}, x_{4}, x_{5}\right)$ model obtains the largest posterior probability with value 0.130 , and the $\left(x_{2}, x_{3}, x_{4}, x_{5}, x_{7}\right)$ model obtains the second largest posterior probability, with value 0.127 . Thus we see that the model choice is reasonably robust to the choice of $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)$, consistently yielding the $\left(x_{1}, x_{2}, x_{3}, x_{4}, x_{5}\right)$ model as the top model for a suitable range of $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)$. Based on these analyses, it does not appear that the variables $x_{6}$ (windspeed) and $x_{7}$ (coldness of temperature) are important predictors of pollen counts.

An analysis was also conducted using $\rho=0$ and $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)=(0.5,0.05)$. In this case, the model that obtains the largest posterior probability is $\left(x_{1}, x_{2}, x_{3}, x_{5}, x_{6}\right)$ with value 0.272 , and the model that obtains the second largest posterior probability is $\left(x_{1}, x_{2}, x_{3}, x_{4}, x_{5}, x_{6}\right)$ with value 0.267 . These results are a bit different from those of Table 2. This can be partially explained by the fact that, for $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)=(0.5,0.05)$, the posterior mean of $\rho$ equals 0.87 , implying a strong degree of correlation between the time measurements. Thus, posterior model probabilities can be sensitive to the choice of $\rho$ if there is a high correlation in the data.

We also did a sensitivity analysis on the choice of $N$. We computed the posterior model probabilities for $N=1,2$ and these are shown in Table 2. The top model for $N=1$ and $N=2$ is $\left(x_{1}, x_{2}, x_{3}, x_{4}, x_{5}\right)$ for all combinations of
$\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)$. Here, $N=1$ corresponds to using the 1993 data as historical data, and $N=2$ corresponds to using the 1992 and 1993 data as historical data. We see the same behavior as in Example 1. For a given $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)$, there is a monotonic increase in the posterior model probability as $N$ is increased.

In this example, as well as in Example 1, 50,000 Gibbs iterations were used in all of the computations after a burn-in of 1,000 iterations. Convergence was checked using the methods discussed in Cowles and Carlin (1996). Specifically, trace plots, autocorrelations, and psr's were computed, and convergence was observed to occur before 500 iterations.

## 5. Discussion

The examples presented in Section 4 had posterior model probabilities quite robust to various choices of ( $\mu_{a_{0}}, \sigma_{a_{0}}$ ), including choices that give high or low weight to the historical data. In Example 1, posterior model probabilities were not sensitive to the choice of $\rho, \rho=0$ and $\rho=0.5$ gave nearly identical results. This suggests that when there is low to moderate correlation in the data, the posterior model probabilities are not sensitive to the choice of $\rho$. However, in Example 2, the results for $\rho=0$ were different from those based on a posterior mean of $\rho$ equal to 0.87 . This suggests that when there is high correlation in the data, the posterior model probabilities can be sensitive to the choice of $\rho$.

Table 1. Posterior model probabilities for simulated data.

| $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)$ | $N=1$ | $N=2$ | $N=3$ |
| :---: | :---: | :---: | :---: |
| $(0.5,0.15)$ | 0.356 | 0.357 | 0.359 |
| $(0.5,0.11)$ | 0.381 | 0.383 | 0.403 |
| $(0.5,0.08)$ | 0.411 | 0.427 | 0.467 |
| $(0.5,0.06)$ | 0.422 | 0.436 | 0.502 |
| $(0.98,0.02)$ | 0.443 | 0.474 | 0.571 |

Table 2. Posterior model probabilities for pollen data.

| $\left(\mu_{a_{0}}, \sigma_{a_{0}}\right)$ | $N=1$ | $N=2$ | $N=3$ |
| :---: | :---: | :---: | :---: |
| $(0.5,0.11)$ | 0.116 | 0.121 | 0.142 |
| $(0.5,0.08)$ | 0.154 | 0.206 | 0.290 |
| $(0.5,0.06)$ | 0.211 | 0.261 | 0.385 |
| $(0.5,0.05)$ | 0.262 | 0.279 | 0.420 |
| $(0.98,0.02)$ | 0.274 | 0.279 | 0.421 |

## Acknowledgements

The authors wish to thank the Editor, Associate Editor and Referees for several suggestions which have greatly improved the paper. We would like to
thank Dr. Harriet Burge for providing us with the data in Example 2, whose work was partially supported by the Center Grant Es0002 from the NIEHS. Dr. Ibrahim's research was supported by NIH grants CA 70101 and CA 74015, and Dr. Chen's research was supported by NSF grant DMS-9702172 and NIH grant CA 74015.

## Appendix

Proof of Lemma 2.1. When $1 / 2 \leq \xi \leq 1$, (2.10) is obviously true. It is easy to see that, for any $0<\xi<1 / 2$,

$$
\begin{aligned}
& \int_{0}^{1} \xi^{a_{0 i}} a_{0 i}^{\alpha_{0}-1}\left(1-a_{0 i}\right)^{\lambda_{0}-1} d a_{0 i} \\
& =\int_{0}^{1 / 2} \xi^{a_{0 i}} a_{0 i}^{\alpha_{0}-1}\left(1-a_{0 i}\right)^{\lambda_{0}-1} d a_{0 i}+\int_{1 / 2}^{1} \xi^{a_{0 i}} a_{0 i}^{\alpha_{0}-1}\left(1-a_{0 i}\right)^{\lambda_{0}-1} d a_{0 i} \\
& \leq K^{*}\left[\int_{0}^{1 / 2} \exp \left(-a_{0 i} \ln (1 / \xi)\right) a_{0 i}^{\alpha_{0}-1} d a_{0 i}+\xi^{1 / 2}\right] \\
& =K^{*}\left[(\ln (1 / \xi))^{-\alpha_{0}} \int_{0}^{\ln (1 / \xi) / 2} \exp \left(-a_{0 i}\right) a_{0 i}^{\alpha_{0}-1} d a_{0 i}+\xi^{1 / 2}\right] \\
& \leq K^{*} K^{* *}(\ln (1 / \xi))^{-\alpha_{0}} \leq K(1+\ln (1 / \xi))^{-\alpha_{0}}
\end{aligned}
$$

where $K^{*}>0, K^{* *}>0$, and $K>0$ are constants. This proves the lemma.
Proof of Theorem 2.1. We first show that if (2.12) holds, for $i=1, \ldots, N$,

$$
\begin{equation*}
\int \exp \left(d_{0}^{*}\left\|\beta^{(m)}\right\|\right) p\left(y_{0 i} \mid \beta^{(m)}, \epsilon_{0 i}, I=m\right) d \beta^{(m)}<K_{1} \tag{A.1}
\end{equation*}
$$

where $d_{0}^{*}>0, K_{1}>0$ is a finite constant independent of $\epsilon_{0 i}$, and $\left\|\beta^{(m)}\right\|=$ $\sqrt{\left(\beta^{(m)}\right)^{\prime} \beta^{(m)}}$.

We have

$$
\begin{align*}
p\left(y_{0 i} \mid \beta^{(m)}, \epsilon_{0 i}, I=m\right) \leq & M^{*} \prod_{j=1}^{k_{m}} \exp \left\{\left(y_{0 i t_{j}} h\left(x_{0 i t_{j}}^{\prime} \beta^{(m)}+\epsilon_{0 i t_{j}}\right)\right.\right. \\
& \left.-q\left(h\left(x_{0 i t_{j}}^{\prime} \beta^{(m)}+\epsilon_{0 i t_{j}}\right)\right)\right\} \tag{A.2}
\end{align*}
$$

where $M^{*}>0$ is a finite constant, and $\epsilon_{0 i t_{j}}$ is the $t_{j}^{t h}$ component of $\epsilon_{0 i}$. Now we make the transformation $u=\left(u_{1}, u_{2}, \ldots, u_{k_{m}}\right)^{\prime}=\left(x_{0 i t_{1}}^{(m)}, x_{0 i t_{2}}^{(m)}, \ldots, x_{0 i t_{k_{m}}}^{(m)}\right)^{\prime} \beta^{(m)}+$ $\left(\epsilon_{0 i t_{1}}, \epsilon_{0 i t_{2}}, \ldots, \epsilon_{0 i t_{k_{m}}}\right)^{\prime}$. This is a one-to-one linear transformation in $k_{m}$ dimensions, since $\left(x_{0 i t_{1}}^{(m)}, x_{0 i t_{2}}^{(m)}, \ldots, x_{0 i t_{k_{m}}}^{(m)}\right)^{\prime}$ has full rank $k_{m}$. Thus,

$$
\begin{equation*}
\left\|\beta^{(m)}\right\| \leq c_{1} \sum_{j=1}^{k_{m}}\left|u_{j}\right| \tag{A.3}
\end{equation*}
$$

where $c_{1}>0$ is a constant. It is easy to see that (A.2) and (A.3) lead to

$$
\begin{aligned}
& \int \exp \left(d_{0}^{*}| | \beta^{(m)} \|\right) p\left(y_{0 i} \mid \beta^{(m)}, \epsilon_{0 i}, I=m\right) d \beta^{(m)} \\
\leq & M^{* *} \prod_{j=1}^{k_{m}}\left\{\int_{-\infty}^{\infty} \exp \left(d_{0}^{*} c_{1}\left|u_{j}\right|\right) \exp \left(\left(y_{0 i t_{j}} h\left(u_{j}\right)-q\left(h\left(u_{j}\right)\right)\right) d u_{j}\right\}=K_{1}<\infty\right. \text { (A.4) }
\end{aligned}
$$

by (2.12), where $M^{* *}>0$ is a constant. This proves (A.1).
Since (2.11) is true, without loss of generality, we assume that $p\left(y_{0 i} \mid \beta^{(m)}\right.$, $\left.\epsilon_{0 i}, I=m\right) \leq 1$ for $j=1, \ldots, N$. Using Lemma 2.1, (A.1), and $\alpha_{0}>k_{m} / N$, we have

$$
\begin{align*}
\int & \left\{\prod_{i=1}^{N} \int_{0}^{1}\left[p\left(y_{0 i} \mid \beta^{(m)}, \epsilon_{0 i}, I=m\right)\right]^{a_{0 i}} a_{0 i}^{\alpha_{0}-1}\left(1-a_{0 i}\right)^{\lambda_{0}-1} d a_{0 i}\right\} d \beta^{(m)} \\
\leq & \int K^{N} \prod_{i=1}^{N}\left(1-\ln \left[p\left(y_{0 i} \mid \beta^{(m)}, \epsilon_{0 i}, I=m\right)\right]\right)^{-\alpha_{0}} d \beta^{(m)} \\
= & K^{N} \int \prod_{i=1}^{N}\left(1-\ln \left[p\left(y_{0 i} \mid \beta^{(m)}, \epsilon_{0 i}, I=m\right)\right]\right)^{-\alpha_{0}} \\
& \times 1\left\{\max _{1 \leq i \leq N} p\left(y_{0 i} \mid \beta^{(m)}, \epsilon_{0 i}, I=m\right)>e^{-d_{0}^{*}\left(\left\|\mid \beta^{(m)}\right\|+1\right)}\right\}^{d \beta^{(m)}} \\
& +K^{N} \int \prod_{i=1}^{N}\left(1-\ln \left[p\left(y_{0 i} \mid \beta^{(m)}, \epsilon_{0 i}, I=m\right)\right]\right)^{-\alpha_{0}} \\
\leq & \left.\left.\times K^{N} \int \max _{1 \leq i \leq N}^{1} p\left(y_{0 i} \mid \beta^{(m)}, \epsilon_{0 i}, I=m\right) \leq e^{-d_{0}^{*}\left(\left\|\mid \beta^{(m)}\right\|+1\right)}\right\}_{1 \leq i \leq N}^{d \beta^{(m)}} p\left(y_{0 i} \mid \beta^{(m)}, \epsilon_{0 i}, I=m\right)>e^{-d_{0}^{*}\left(\left\|\beta^{(m)}\right\|+1\right)}\right\}^{d \beta^{(m)}} \\
& +K^{N} \int \prod_{i=1}^{N}\left[d_{0}^{*}\left(\left\|\beta^{(m)}\right\|+1\right)\right]^{-\alpha_{0}} d \beta^{(m)} \\
\leq & K^{N} \sum_{1 \leq i \leq N} \int p\left(y_{0 i} \mid \beta^{(m)}, \epsilon_{0 i}, I=m\right) e^{d_{0}^{*}\left(\left\|\mid \beta^{(m)}\right\|+1\right)} d \beta^{(m)} \\
& +K^{N}\left(d_{0}^{*}\right)^{-\alpha_{0} N} \int\left(\left\|\beta^{(m)}\right\|+1\right)^{-N \alpha_{0}} d \beta^{(m)} \leq K_{2}<\infty,
\end{align*}
$$

where $K_{2}>0$ is a constant, independent of $\epsilon_{0 i}$ for $i=1, \ldots, N$. Finally, it follows directly from (A.5) that the normalizing constant of the prior is less than or equal to

$$
K_{2} \int_{-1}^{1} \int_{0}^{\infty} \prod_{i=1}^{N}\left(\int \phi_{n 0 i}\left(\epsilon_{0 i} \mid 0, \sigma^{2} \Sigma_{0 i}\right) d \epsilon_{0 i}\right)
$$

$$
\begin{aligned}
& \times\left(\sigma^{2}\right)^{-\left(\delta_{0}+1\right)} \exp \left(-\sigma^{-2} \gamma_{0}\right) \times(1+\rho)^{\nu_{0}-1}(1-\rho)^{\psi_{0}-1} d \sigma^{2} d \rho \\
= & K_{2} \int_{-1}^{1} \int_{0}^{\infty}\left(\sigma^{2}\right)^{-\left(\delta_{0}+1\right)} \exp \left(-\sigma^{-2} \gamma_{0}\right) \times(1+\rho)^{\nu_{0}-1}(1-\rho)^{\psi_{0}-1} d \sigma^{2} d \rho<\infty
\end{aligned}
$$

since $\left(\sigma^{2}\right)^{-\left(\delta_{0}+1\right)} \exp \left(-\sigma^{-2} \gamma_{0}\right)$ and $(1+\rho)^{\nu_{0}-1}(1-\rho)^{\psi_{0}-1}$ are proper priors. This proves the theorem.
Proof of Theorem 3.1. Let $\pi\left(\beta^{(-m)} \mid y_{0}, I=\mathcal{K}\right)$ and $p\left(\beta^{(-m)} \mid y, y_{0}, I=\mathcal{K}\right)$ denote the respective marginal prior and posterior distributions of $\beta^{(-m)}$ obtained from the full model. The Savage-Dicky density ratio (see, for example, Verdinelli and Wasserman (1995)) directly yields

$$
\begin{equation*}
\frac{p(y \mid I=m)}{p(y \mid I=\mathcal{K})}=\frac{p\left(\beta^{(-m)}=0 \mid y, y_{0}, I=\mathcal{K}\right)}{\pi\left(\beta^{(-m)}=0 \mid y_{0}, I=\mathcal{K}\right)}, \quad m=1, \ldots, \mathcal{K} . \tag{A.6}
\end{equation*}
$$

Using (A.6) and (3.1), it suffices to show that $p(I=m) \propto \pi\left(\beta^{(-m)}=0 \mid y_{0}, I=\right.$ $\mathcal{K})$ 。

It can be easily observed that $\int p_{0}^{*}\left(\beta^{(\mathcal{K})}, \sigma^{2}, \rho \mid y_{0}, I=\mathcal{K}\right) d \beta^{(\mathcal{K})} d \sigma^{2} d \rho=$ $\frac{p_{0}^{*}\left(\beta^{(m)}, \beta^{(-m)}=0, \sigma^{2}, \rho \mid y_{0}, I=\mathcal{K}\right)}{\pi\left(\beta^{(m)}, \beta^{(-m)}=0, \sigma^{2}, \rho \mid y_{0}, I=\mathcal{K}\right)}$, and $\int p_{0}^{*}\left(\beta^{(m)} \mid y_{0}, I=m\right) d \beta^{(m)}=\frac{p_{0}^{*}\left(\beta^{(m)}, \sigma^{2}, \rho \mid y_{0}, I=m\right)}{\pi\left(\beta^{(m)}, \sigma^{2}, \rho \mid y_{0}, I=m\right)}$. Then we are led to $p_{0}^{*}\left(\beta^{(m)}, \beta^{(-m)}=0, \sigma^{2}, \rho \mid y_{0}, I=\mathcal{K}\right)=p_{0}^{*}\left(\beta^{(m)}, \sigma^{2}, \rho \mid y_{0}, I=m\right)$ and $\pi\left(\beta^{(m)}, \beta^{(-m)}=0, \sigma^{2}, \rho \mid y_{0}, I=\mathcal{K}\right)=\pi\left(\beta^{(-m)}=0 \mid y_{0}, I=\mathcal{K}\right) \pi\left(\beta^{(m)}, \sigma^{2}, \rho \mid y_{0}\right.$, $I=m)$. The above two identities yield

$$
\int p_{0}^{*}\left(\beta^{(\mathcal{K})}, \sigma^{2}, \rho \mid y_{0}, I=\mathcal{K}\right) d \beta^{(\mathcal{K})} d \sigma^{2} d \rho=\frac{\int p_{0}^{*}\left(\beta^{(m)} \mid y_{0}, I=m\right) d \beta^{(m)}}{\pi\left(\beta^{(-m)}=0 \mid y_{0}, I=\mathcal{K}\right)}
$$

Now note that $\frac{\int p_{0}^{*}\left(\beta^{(\mathcal{K})}, \sigma^{2}, \rho \mid y_{0}, I=\mathcal{K}\right) d \beta^{(\mathcal{K})} d \sigma^{2} d \rho}{\sum_{j=1}^{\mathcal{K}} \int p_{0}^{*}\left(\beta^{(j)} \mid y_{0}, I=j\right) d \beta^{(j)}}$ is independent of the model index $m$. This completes the proof.

## References

Bedrick, E. J., Christensen, R. and Johnson, W. (1996). A new perspective on priors for generalized linear models. J. Amer. Statist. Assoc. 91, 1450-1460.
Chen, M.-H. (1994). Importance-weighted marginal Bayesian posterior density estimation. J. Amer. Statist. Assoc. 89, 818-824.
Chen, M.-H., Ibrahim, J. G. and Yiannoutsos, C. (1999). Prior elicitation, variable selection, and and Bayesian computation for logistic regression models. J. Roy. Statist. Soc. Ser. B61, 223-242.
Chen, M.-H. and Shao, Q.-M. (1997). Performance study of marginal posterior density estimation via Kullback-Leibler divergence. Test 6, 321-350.
Cowles, M. K. and Carlin, B. P. (1996). Markov chain Monte Carlo convergence diagnostics: a comparative review. J. Amer. Statist. Assoc. 91, 883-904.
Gelfand, A. E., Sahu, S. K. and Carlin, B. P. (1996). Efficient parametrisations for generalized linear mixed models (with discussion). In Bayesian Statistics 5 (Edited by J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith), 165-180. Oxford University Press, Oxford.

Gelfand, A. E., Smith, A. F. M. and Lee, T. M. (1992). Bayesian analysis of constrained parameter and truncated data problems using Gibbs sampling. J. Amer. Statist. Assoc. 87, 523-532.
Ibrahim, J. G., Ryan, L. M. and Chen, M.-H. (1998). Use of historical controls to adjust for covariates in trend tests for binary data. J. Amer. Statist. Assoc. 93, 1282-1293.
Shephard, N. and Pitt, M. K. (1997). Likelihood analysis of non-Gaussian measurement time series. Biometrika 84, 653-668.
Stark, P. C., Ryan, L. M., McDonald, J. L. amd Burge, H. A. (1997). Using meteorologic data to predict daily ragweed pollen levels. Aerobiologia 13, 177-184.
Verdinelli, I. and Wasserman, L. (1995). Computing Bayes factors using a generalization of the Savage-Dickey density ratio. J. Amer. Statist. Assoc. 90, 614-618.
Zeger, S. L., (1988). A regression model for time series of counts. Biometrika 75, 621-629.

Department of Biostatistics, Harvard University, School of Public Health, 44 Binney St., Boston MA 02115, U.S.A.
E-mail: ibrahim@jimmy.harvard.edu
(Received June 1998; accepted November 1999)

