Deviance Information Criterion for Comparing Stochastic Volatility Models

Andreas BERG
Department of Statistics, University of Auckland, Private Bag 92019, Auckland, New Zealand
(*andreas@stat.auckland.ac.nz*)

Renate MEYER
Department of Statistics, University of Auckland, Private Bag 92019, Auckland, New Zealand
(*meyer@stat.auckland.ac.nz*)

Jun YU
Department of Economics, University of Auckland, Private Bag 92019, Auckland, New Zealand
(*j.yu@auckland.ac.nz*)

Bayesian methods have been efficient in estimating parameters of stochastic volatility models for analyzing financial time series. Recent advances made it possible to fit stochastic volatility models of increasing complexity, including covariates, leverage effects, jump components, and heavy-tailed distributions. However, a formal model comparison via Bayes factors remains difficult. The main objective of this article is to demonstrate that model selection is more easily performed using the deviance information criterion (DIC). It combines a Bayesian measure of fit with a measure of model complexity. We illustrate the performance of DIC in discriminating between various different stochastic volatility models using simulated data and daily returns data on the Standard & Poors (S&P) 100 index.

KEY WORDS: Bayesian deviance; Jumps; Leverage effect; Markov chain Monte Carlo; Model complexity; Model selection.

1. INTRODUCTION

Progress in Bayesian posterior computation due to Markov chain Monte Carlo (MCMC) methods has made it possible to fit increasingly complex statistical models and entailed the wish to determine the best fitting model in a potentially huge class of candidates. Thus, it has become more and more important to develop efficient model selection criteria. A recent proposal by Spiegelhalter, Best, Carlin, and van der Linde (2002) was the deviance information criterion (DIC), a Bayesian version or generalization of the well-known Akaike information criterion (AIC) (Akaike 1973), related also to the Bayesian (or Schwarz) information criterion (BIC) (Schwarz 1978). Similar to AIC and BIC, it trades off a measure of model adequacy against a measure of complexity. DIC is easy to calculate and applicable to a wide range of statistical models. It is based on the posterior distribution of the log-likelihood or the deviance, following the original suggestion of Dempster (1974) for model choice in the Bayesian framework. This model comparison criterion has already been applied successfully to complex models in the field of medical statistics (Zhu and Carlin 2000). In this article, we demonstrate its usefulness in the model selection process for financial time series. The aim of this article is, therefore, to introduce DIC to the financial modeling community and show how to use it for the family of stochastic volatility (SV) models.

Indeed, many model-checking criteria (Carlin and Louis 1996; Gelman, Carlin, Stern, and Rubin 1996; Gilks, Richardson, and Spiegelhalter 1996; Key, Pericchi, and Smith 1999) have been proposed and discussed before the development of DIC. Whereas Bayes factors (e.g., Kass and Raftery 1995) have been viewed for many years as the only correct way to carry out Bayesian model comparison, they have come under increasing criticism of late (Kass and Raftery 1995; Lavine and Schervish 1999). One serious drawback is that they are not well defined when using improper priors, which is typically the case in practice when employing noninformative priors. This led to modifications, such as the *partial Bayes factor* (O'Hagan 1991), the *intrinsic Bayes factor* (Berger and Pericchi 1996), and the *fractional Bayes factor* (O'Hagan 1994). These modifications suffer from more or less arbitrary choices of training samples, weights for averaging training samples, and fractions, respectively. For specifying Bayesian stochastic volatility (SV) models, however, informative and, thus, *proper* prior distributions are usually employed and Bayes factors are well defined. Nonetheless, the number of unknown parameters in Bayesian SV models is large (exceeding the number of observations) because of the latent volatilities. Calculation of the Bayes factor for comparing any two models requires the marginal likelihoods and, thus, a marginalization over the parameter vectors in each model. If the dimension of the parameter space is large, these implicit, extremely high-dimensional integration problems pose a formidable computational challenge. In the context of SV models, Kim, Shephard, and Chib (1998) and Chib, Nardari, and Shephard (2002) showed how to compute Bayes factors using the marginal likelihood approach of Chib (1995) and evaluating the marginal likelihood at the posterior mean using *particle filtering* (Kitagawa 1996; Pitt and Shephard 1999a; Doucet, de Freitas, and Gordon 2001). Still, it remains a computationally intensive task and is not a particularly user-friendly tool for practicing statisticians. In their review of MCMC methods

Journal of Business & Economic Statistics
January 2004, Vol. 22, No. 1
DOI 10.1198/073500103288619430

for computing Bayes factors, Han and Carlin (2001, p. 29) concluded that "all of the methods. . .discussed require substantial time and effort (both human and computer) for a rather modest payoff, namely a collection of posterior model probability estimates.... As a result, one might conclude that none of the methods herein is appropriate for everyday, 'rough and ready' model comparison, and instead search for more computationally realistic alternatives."

A well-known estimate of the marginal likelihood developed by Newton and Raftery (1994) is the harmonic mean of the likelihood values. It is easy to compute and simulation consistent but not stable because the inverse likelihood does not possess a finite variance (Chib 1995). Other shortcuts to the calculation of Bayes factors that avoid multidimensional integration through large sample approximations of $-2\ln(\text{Bayes factor})$ include the familiar BIC, also referred to as the Schwarz criterion (Schwarz 1978), and the related penalized likelihood ratio model choice criterion, AIC. Either criterion requires the specification of the number of free parameters in each model. If we consider a nonhierarchical Bayesian model with parameter θ, a flat prior would correspond to a flexible and, thus, complex model, whereas a tight prior constrains the model. The classical definition of model complexity as the "number of unknown parameters" could thus be considered as a special case corresponding to a noninformative prior. However, for a complex hierarchical model the specification of its dimensionality is rather arbitrary. This is typically the case for an SV model, where the parameters are augmented by the n latent volatilities, with n being the sample size. As these are not independent but exhibit a Markovian dependence structure, they cannot be counted as n additional free parameters. Thus, neither BIC nor AIC is applicable for SV model comparison. As detailed in Section 3, DIC avoids this dilemma by using a complexity measure for the effective number of parameters that is based on an information-theoretic argument. This quantity is readily obtained from an MCMC analysis, which makes algebraic forms and large sample approximations obsolete.

It is usually hard to specify prior model probabilities necessary for the calculation of posterior model probabilities. By using DIC as a formal approach to model selection, combining a measure of fit and complexity, we can avoid this need. However, we caution in general against basing model choice solely on information criteria, as many other factors such as the model's inherent plausibility and the robustness of its inferences and model diagnostics (as, for instance, outlined in Kim et al. 1998, sec. 4.2; Spiegelhalter et al. 2002, sec. 6) need to be taken into account. In many instances, when none of the models is clearly superior, model averaging (Hoeting, Madigan, Raftery, and Volinsky 1999) might be more appropriate. Whether DIC can be used as a basis for model averaging is still an open question. And it should also be stressed that no prior model probabilities are necessary for the calculation of Bayes factors.

The outline of the article is as follows: Section 2 gives an introduction to SV models, followed in Section 3 by the definition and properties of DIC. Section 4 reviews Chib's (1995) method for calculating the marginal likelihood based on particle filtering and Newton and Raftery's (1994) harmonic mean estimate of the marginal likelihood. In Section 5 we present results from a simulation study and compare the model ranking implied by the marginal likelihood, harmonic mean, and DIC. Section 6 applies DIC to compare the fit of various SV models to a dataset previously analyzed in the literature. We also assess the performance of DIC using the Bayes factor as a gold standard and examine the prior sensitivity of DIC. In Section 7 we present our conclusions.

2. THE STOCHASTIC VOLATILITY MODEL

In both the theoretic finance literature on option pricing and the empirical finance literature, the SV model (Taylor 1982; Hull and White 1987) has received much attention in recent years. It has become a powerful alternative to the autoregressive conditional heteroscedasticity (ARCH) and generalized autoregressive conditional heteroscedasticity (GARCH) models introduced by Engle (1982) and Bollerslev (1986). Ghysels, Harvey, and Renault (1996) and Shephard (1996) gave excellent reviews of the model.

Given a time series of daily returns $\{y_t\}_{t=1}^n$, a basic SV model consists of an observation equation

$$y_t|h_t = \exp(h_t/2)u_t, \qquad t = 1, \ldots, n, \tag{1}$$

that describes the distribution of the data given unknown states, the log-volatilities h_t, and a state equation

$$h_t|h_{t-1} = \mu + \phi(h_{t-1} - \mu) + v_t, \qquad t = 1, \ldots, n, \tag{2}$$

which models the day-to-day variation of the volatilities as a Markov process. Here y_t is the response variable, h_t is the log-volatility process, and the errors u_t and v_t are uncorrelated Gaussian sequences with $u_t \sim N(0, 1)$ and $v_t \sim N(0, \tau^2)$. We collect the three model parameters in a vector $z = (\phi, \mu, \tau^2)$. In Sections 5 and 6 we will introduce extensions of the basic model to more complex SV models. An example of such an extension is the inclusion of a level effect in the observation equation, namely,

$$y_t = x_t^{\gamma} \exp(h_t/2)u_t, \qquad t = 1, \ldots, n,$$

where x_t denotes a time-varying covariate. The parameter γ plays an important role in analyzing interest rate data (for details refer, e.g., to Chan, Karolyi, Longstaff, and Sanders 1992; Brenner, Harjes, and Kroner 1996). In other applications, for example, stock market data, it is common to set this parameter equal to 0 (see Sec. 5).

Classical parameter estimation for this model is extremely difficult, because of the nonanalytic form of the likelihood function. Harvey, Ruiz, and Shephard (1994) employed a quasi-maximum likelihood technique, whereas Sandmann and Koopman (1998) used the maximum likelihood Monte Carlo method. Several method-of-moment approaches such as the efficient method of moments (Gallant and Tauchen 1996; Andersen, Chung, and Sørensen 1999), the spectral method of moments (Singleton 2001; Chacko and Viceira 2003; Knight, Satchell, and Yu 2002), the simulated method of moments (Duffie and Singleton 1993) and the generalized method of moments (Melino and Turnbull 1990; Andersen and Sørensen 1996) have also been used to estimate the model parameters.

Whereas some of the previously mentioned techniques use ad hoc criteria (see Andersen et al. 1999 for a review and comparison of various estimation techniques for the SV model), a Bayesian approach is based on a sound statistical paradigm. Bayesian posterior computations are performed using MCMC techniques. Several different algorithms have been proposed by Jacquier, Polson, and Rossi (1994) and Kim et al. (1998) and further developed in Chib et al. (2002). Although more efficient updating techniques for SV models exist, we employ the all-purpose Bayesian software package BUGS based on the single-update Gibbs sampler as described in Meyer and Yu (2000) for ease of implementation. The SV model is a typical example of a hierarchical model, in which the number of unknowns, that is, the parameters (z) and the unknown states ($h_1, \dots, h_n$), exceeds the number of observations. The number of free parameters in the model could be the number of model parameters (3) or the number of states plus the number of model parameters ($n + 3$) or something in between. In any case this number is not well defined and, thus, precludes the use of AIC or BIC for model comparison. We will show that DIC provides an efficient and straightforward approach to defining the effective number of parameters and to identifying the most appropriate model.

3. THE DEVIANCE INFORMATION CRITERION

Assume, in general, that the distribution of the data, $y = (y_1, \dots, y_n)$, depends on a p-dimensional parameter vector θ. (In the context of an SV model, θ encompasses the parameter vector z and the vector of log-volatilities $h_1, \dots, h_n$.) From a frequentist point of view, model assessment is based on the *deviance*, the difference in the log-likelihoods between the fitted and the saturated model. The saturated model refers to the model with as many parameters as observations, which yields a perfect fit to the data. By analogy, Dempster (1974) suggested examining the posterior distribution of the classical deviance defined by

$$D(\theta) = -2 \ln f(y|\theta) + 2 \ln g(y) \tag{3}$$

for Bayesian model selection. Here $f(y|\theta)$ is the likelihood function, that is, the conditional joint probability density function of the observations given the unknown parameters, and $\ln g(y)$ denotes a fully specified standardizing term that is a function of the data alone [in our applications in Secs. 5 and 6, $g(y) = 1$]. Dempster (1974) proposed comparing plots and potential summaries such as the posterior mean of $D(\theta)$, and Spiegelhalter et al. (2002) followed these suggestions in the development of DIC as a model choice criterion. Based on the posterior distribution of $D(\theta)$, DIC consists of two components, a term that measures goodness of fit and a penalty term for increasing model complexity:

$$\mathrm{DIC} = \bar{D} + p_D. \tag{4}$$

1. The first term, a Bayesian measure of model fit, is defined as the posterior expectation of the deviance

$$\bar{D} = E_{\theta|y}[D(\theta)] = E_{\theta|y}[-2 \ln f(y|\theta)]. \tag{5}$$

The "better" the model fits the data, the larger are the values for the likelihood. $\bar{D}$, which is defined as -2 times log-likelihood, therefore, attains smaller values for "better" models.

2. The second component measures the complexity of the model by the *effective number of parameters*, p_D, defined as the difference between the posterior mean of the deviance and the deviance evaluated at the posterior mean $\bar{\theta}$ of the parameters:

$$\begin{aligned} p_D &= \bar{D} - D(\bar{\theta}) = E_{\theta|y}[D(\theta)] - D(E_{\theta|y}[\theta]) \\ &= E_{\theta|y}[-2 \ln f(y|\theta)] + 2 \ln f(y|\bar{\theta}). \end{aligned} \tag{6}$$

By defining $-2 \ln f(y|\theta)$ as the residual information in the data y conditional on θ and interpreting it as a measure of uncertainty, (6) shows that p_D can be regarded as the expected excess of the true over the estimated residual information in data y conditional on θ. This means we can interpret p_D as the expected reduction in uncertainty due to estimation.

Rearranging (6), one gets $\bar{D} = D(\bar{\theta}) + p_D$. Thus, the DIC defined in (4) can be reexpressed as

$$\mathrm{DIC} = D(\bar{\theta}) + 2p_D, \tag{7}$$

which can be interpreted as a classical "plug-in" measure of fit plus a measure of complexity. Therefore, the Bayesian measure of fit $\bar{D} = D(\bar{\theta}) + p_d$ already includes a penalty term for model complexity and could thus be better thought of as a measure of "model adequacy" rather than pure goodness of fit.

Spiegelhalter et al. (2002) gave an asymptotic justification of DIC in the case where the number of observations n grows with respect to the number of parameters p and where the prior is nonhierarchical and completely specified (i.e., without hyperparameters). In this situation $\mathrm{AIC} = D(\hat{\theta}) + 2p$, where $\hat{\theta}$ denotes the maximum likelihood (ML) estimate. This is the same formula as (7) but with the posterior mean $\bar{\theta}$ substituted by the ML estimate $\hat{\theta}$. Thus, DIC can be seen as a generalization of AIC, and it also can be compared to the Schwarz information criterion $\mathrm{BIC} = -2 \ln f(y|\hat{\theta}) + p \ln n$. In the special case where the prior is flat, a case that corresponds to a frequentist analysis, AIC equals DIC because the ML estimate coincides with the posterior mean. In the context of normal linear regression with uncertainty in the choice of regressors, George and Foster (2001) developed empirical Bayes alternatives to penalized likelihood criteria such as AIC and BIC, and Fernandez, Ley, and Steel (2001) pointed out links of Bayes factors with classical information criteria and provided a unifying framework.

By applying a logarithmic loss function, Spiegelhalter et al. (2002) gave a decision-theoretic justification for DIC and showed that DIC approximately describes the expected posterior loss when adopting a particular model. For additional asymptotic properties of p_D and $\bar{D}$, the interested reader is referred to Spiegelhalter et al. (2002).

In striking contrast to calculating Bayes factors, computing DIC via MCMC is almost trivial. An estimate of $\bar{D}$ is easily calculated from the MCMC output by monitoring $D(\theta)$ and then taking the sample mean of the simulated values of $D(\theta)$. The effective number of parameters p_D can be obtained by evaluating $D(\theta)$ at the sample average of the simulated values of θ and subtracting this plug-in estimate of the deviance from the estimate of $\bar{D}$ (see also Sec. 5.3).

So far, no efficient method has been developed for calculating reasonably accurate MC standard errors of DIC. Zhu and Carlin (2000) explored this problem, but their approach using the multivariate delta method yields poor results. Their final recommendation is the "brute force" approach, which is simply replicating the calculation of DIC some N times and estimating VAR(DIC) by its sample variance

$$\widehat{\mathrm{VAR}}(\mathrm{DIC}) = \frac{1}{N-1}\sum_{k=1}^{N}(\mathrm{DIC}_k - \overline{\mathrm{DIC}})^2.$$

Although this is a painfully time-consuming approach, it at least gives an indication of the inherent variability of DIC.

4. MARGINAL LIKELIHOOD AND HARMONIC MEAN

Because we are going to compare the performance of DIC with that of Chib's marginal likelihood method and the harmonic mean in the next two sections, it is worthwhile to first review Chib's method for calculating the marginal likelihood and Newton and Raftery's (1994) method for estimating the marginal likelihood by the harmonic mean of the sampled likelihood values.

4.1 Chib's Marginal Likelihood

By definition, the marginal likelihood $m(y)$ is the integral of the likelihood function with respect to the prior density $\pi(z)$:

$$m(y) = \int f(y|z)\pi(z)\,dz, \tag{8}$$

with z denoting the vector of parameters in the model. As solving this integral would require high-dimensional integration, Chib (1995) suggested evaluating the marginal likelihood by rearranging Bayes's theorem

$$m(y) = \frac{f(y|z)\pi(z)}{\pi(z|y)},$$

where $\pi(z|y)$ denotes the posterior probability density function of z. Thus, the log-marginal likelihood, which is referred to as $\ln L$ in the following, can be calculated by

$$\ln L = \ln m(y) = \ln f(y|z) + \ln\pi(z) - \ln\pi(z|y), \tag{9}$$

where z is an appropriately selected high-density point (in this article we simply use the posterior mean $\bar{z}$). The first term on the right side of (9) is the log-likelihood evaluated at the posterior mean of the parameter vector z (marginalized over the latent volatilities h_t) and the second term is the log prior density evaluated at $\bar{z}$. The third quantity involves the posterior density, which is only known up to a normality constant. However, an approximation can be obtained by using a multivariate kernel density estimate as suggested in Kim et al. (1998) (see also Silverman 1986, chap. 4), which is based on the posterior MCMC sample of z.

We mentioned in Section 2 that the log-likelihood function $\ln f(y|z)$ has no analytical form for the SV model as it is marginalized over the latent states $h_1, \ldots, h_n$, and this is why the exact maximum likelihood method is extremely difficult to implement. However, it is possible to approximate the likelihood by making use of the so-called *particle filter* method. Important contributions in this area include Gordon, Salmond, and Smith (1993), Kitagawa (1996), and Pitt and Shephard (1999a). By successive conditioning, the log-likelihood $\ln f(y|\bar{z})$ can be decomposed into

$$\ln f(y|\bar{z}) = \ln f(y_1|\bar{z}) + \sum_{t=1}^{n-1}\ln f(y_{t+1}|Y_t, \bar{z}), \tag{10}$$

where $Y_t = (y_1, \ldots, y_t)$ collects the available data up to time t. Taking the latent volatilities into account, each one-step-ahead prediction density has a mixture representation as

$$\begin{aligned} f(y_{t+1}|Y_t, \bar{z}) &= \int f(y_{t+1}|h_{t+1}, Y_t, \bar{z}) f(h_{t+1}|Y_t, \bar{z})\,dh_{t+1} \\ &= \int f(y_{t+1}|h_{t+1}, Y_t, \bar{z}) \\ &\quad \times \left[\int f(h_{t+1}|h_t, \bar{z}) f(h_t|Y_t, \bar{z})\,dh_t\right] dh_{t+1} \end{aligned}$$

and can, thus, be estimated by

$$\frac{1}{M}\sum_{i=1}^{M} f\big(y_{t+1}|h_{t+1}^{(i)}\big),$$

where $h_{t+1}^{(i)}|h_t^{(i)}$ is drawn from the state equation (2) given samples $h_t^{(i)}$, the so-called *filtered particles*, from $f(h_t|Y_t, \bar{z})$.

In this article we utilize Kitagawa's algorithm for particle filtering, which is applicable to a broad class of nonlinear non-Gaussian multidimensional state space models of the form

$$\begin{aligned} y_t &= H(x_t, u_t), \\ x_t &= F(x_{t-1}, v_t), \end{aligned} \tag{11}$$

where x_t is a k-dimensional state vector (here $x_t = h_t$ is the one-dimensional log-volatility), v_t is an l-dimensional white-noise sequence with density $q(v)$, u_t is a one-dimensional white-noise sequence with density $r(u)$ and assumed uncorrelated with $\{v_s\}_{s=1}^n$, and H and F are possibly nonlinear functions. Let $u_t = G(y_t, x_t)$ and let G' be the derivative of G as a function of y_t. The density of the initial state vector is assumed to be $p_0(x)$. We now summarize all the steps involved in Kitagawa's algorithm:

1. Generate M l-dimensional particles from $p_0(x)$, $f_0^{(j)}$ for $j = 1, \ldots, M$.
2. Repeat the following steps for $t = 1, \ldots, n$.
   a. Generate M l-dimensional particles from $q(v)$, $v_t^{(j)}$ for $j = 1, \ldots, M$.
   b. Compute $p_t^{(j)} = F(f_{t-1}^{(j)}, v_t^{(j)})$ for $j = 1, \ldots, M$.
   c. Compute $\alpha_t^{(j)} = r(G(y_t, p_t^{(j)}))$ for $j = 1, \ldots, M$.
   d. Resample $\{p_t^{(j)}\}_{j=1}^M$ to get $\{f_t^{(j)}\}_{j=1}^M$ with probabilities proportional to $\{r(G(y_t, p_t^{(j)})) \times |G'(y_t, p_t^{(j)})|\}_{j=1}^M$.

It can be seen that almost all the SV models presented in the next two sections can be rewritten in the state space form (11); hence, it is straightforward to modify the preceding algorithm to fit our needs. The only exception is Model 5, which violates the assumption of no correlation between u_t and v_{t+1}. When

Model 5 is introduced in Section 5, we will show how a simple rewrite of the model allows for a direct use of Kitagawa's algorithm.

We should point out that more efficient particle filter algorithms are available. An example is the *auxiliary particle filter* introduced by Pitt and Shephard (1999a); see the implementation of this particle filter algorithm in Kim et al. (1998), Pitt and Shephard (1999b), Chib, Nardari, and Shephard (1999) and Chib et al. (2002) in the context of SV models. Our experience suggests that by choosing $M = 50{,}000$ for Kitagawa's algorithm one obtains very similar results to the auxiliary particle filter method with $M = 2{,}500$.

4.2 Harmonic Mean

Newton and Raftery (1994) suggested the calculation of approximate Bayes factors for model comparison using the harmonic mean of the sampled likelihood values as a simulation-consistent estimator of the required marginal likelihood. Let θ denote the parameter vector (augmented by latent volatilities), that is, $\theta = (z, h_1, \ldots, h_n)$, as in Section 3. Similar to (8), the marginal likelihood $m(y)$ can be expressed as

$$m(y) = \int f(y|\theta) f(\theta)\, d\theta,$$

where $f(\theta)$ denotes the joint prior density function of θ. The importance sampling method for evaluating this integral is to generate a sample $\{\theta^{(i)}; i = 1, \ldots, M\}$ from a so-called importance sampling density $f^*(\theta)$. Under quite weak assumptions, a simulation-consistent estimate of $m(y)$ is given by

$$\hat{m}(y) = \frac{\sum_{i=1}^{M} w_i f(y|\theta^{(i)})}{\sum_{i=1}^{M} w_i}, \tag{12}$$

where $w_i = f(\theta^{(i)})/f^*(\theta^{(i)})$. The Gibbs sampler gives us a sample $\theta^{(i)}$ approximately drawn from the posterior density $f^*(\theta) = f(\theta|y) = f(y|\theta)f(\theta)/m(y)$. Using these $\theta^{(i)}$ in (12) yields the harmonic mean estimator of $m(y)$:

$$\hat{m}_{hm}(y) = \left\{ \frac{1}{M} \sum_{i=1}^{M} \frac{1}{f(y|\theta^{(i)})} \right\}^{-1}. \tag{13}$$

Here $\hat{m}_{hm}(y)$ converges almost surely to the correct value $m(y)$ as M goes to ∞, but it does not, in general, satisfy a Gaussian central limit theorem as $1/f(y|\theta)$ is often not square integrable with respect to the posterior distribution. Thus, a few outlying values $\theta^{(i)}$ with small likelihood values can have a large effect on the estimate. For this reason Newton and Raftery (1994) also proposed modified estimators that are much more stable than the straight harmonic mean that we used here.

5. A SIMULATION STUDY

The main objective of this simulation study is to see whether DIC is capable of identifying the true model from which the data are generated. Following suggestions by the referees, we also calculate Chib's marginal likelihood and the harmonic mean estimate for each model within the set of competing models. However, we want to point out an argument by Spiegelhalter et al. (2002, rejoinder) that cautions against using the Bayes factor (or marginal likelihood) as a gold standard against which to assess DIC. The Bayes factor addresses how well the prior has predicted the observed data, whereas DIC addresses how well the posterior might predict future data generated by the same mechanism that gave rise to the observed data. Thus, these criteria cannot, in general, be expected to arrive at the same conclusions as they are designed to answer different questions. Especially for the practical selection of models of financial time series, we consider this posterior predictive outlook of DIC to be potentially more relevant.

We simulate a dataset comprising 2,000 observations from a stochastic volatility model that includes a jump component as described later. The data are plotted in the first panel of Figure 1. This SV + jumps model (Model 6) is very similar to the one proposed in the simulation analysis by Chib et al. (2002). We use the BUGS (Bayesian Inference Using Gibbs Sampling) software package (Spiegelhalter, Thomas, Best and Gilks 1996), available online at *http://www.mrc-bsu.cam.ac.uk/bugs/welcome.shtml*, for posterior computation. BUGS is an easy-to-learn and easy-to-use Bayesian software package that implements the Gibbs sampler for generating samples from a Markov chain whose equilibrium distribution is the posterior distribution. As demonstrated by Meyer and Yu (2000), it can be applied to fit stochastic volatility models. Although more efficient Markov chain Monte Carlo techniques exist for fitting SV models (Kim et al. 1998), the use of BUGS is highly attractive due to the ease of implementation. In the following, we describe the list of competing models under consideration.

5.1 The Models

We fit eight different stochastic volatility models to the simulated data, including the true model from which the data are generated (Model 6). For each of the models we list the observation and state equations (for $t = 1, \ldots, n$) and their distributional assumptions. For all cases we assume u_t and $\{v_s\}_{s=1}^{n}$ are uncorrelated unless we claim otherwise. Prior distributions for the unknown parameters are stated in Section 5.2.

Model 1. This model is identical to the basic SV model in Section 2:

$$y_t|h_t = \exp(h_t/2)u_t, \qquad u_t \overset{\text{i.i.d.}}{\sim} N(0,1),$$
$$h_t|h_{t-1}, \mu, \phi, \tau^2 = \mu + \phi(h_{t-1} - \mu) + v_t, \qquad v_t \overset{\text{i.i.d.}}{\sim} N(0,\tau^2),$$

with $h_0 \sim N(\mu, \tau^2)$.

Model 2. An additional nonzero mean α is added to the observation equation:

$$y_t|h_t, \alpha = \alpha + \exp(h_t/2)u_t, \qquad u_t \overset{\text{i.i.d.}}{\sim} N(0,1),$$
$$h_t|h_{t-1}, \mu, \phi, \tau^2 = \mu + \phi(h_{t-1} - \mu) + v_t, \qquad v_t \overset{\text{i.i.d.}}{\sim} N(0,\tau^2).$$

Model 3. An AR(2) process for the state equation:

$$y_t|h_t = \exp(h_t/2)u_t, \qquad u_t \overset{\text{i.i.d.}}{\sim} N(0,1),$$
$$h_t|h_{t-1}, \mu, \phi, \psi, \tau^2 = \mu + \phi(h_{t-1} - \mu) + \psi(h_{t-2} - \mu) + v_t,$$
$$v_t \overset{\text{i.i.d.}}{\sim} N(0,\tau^2).$$

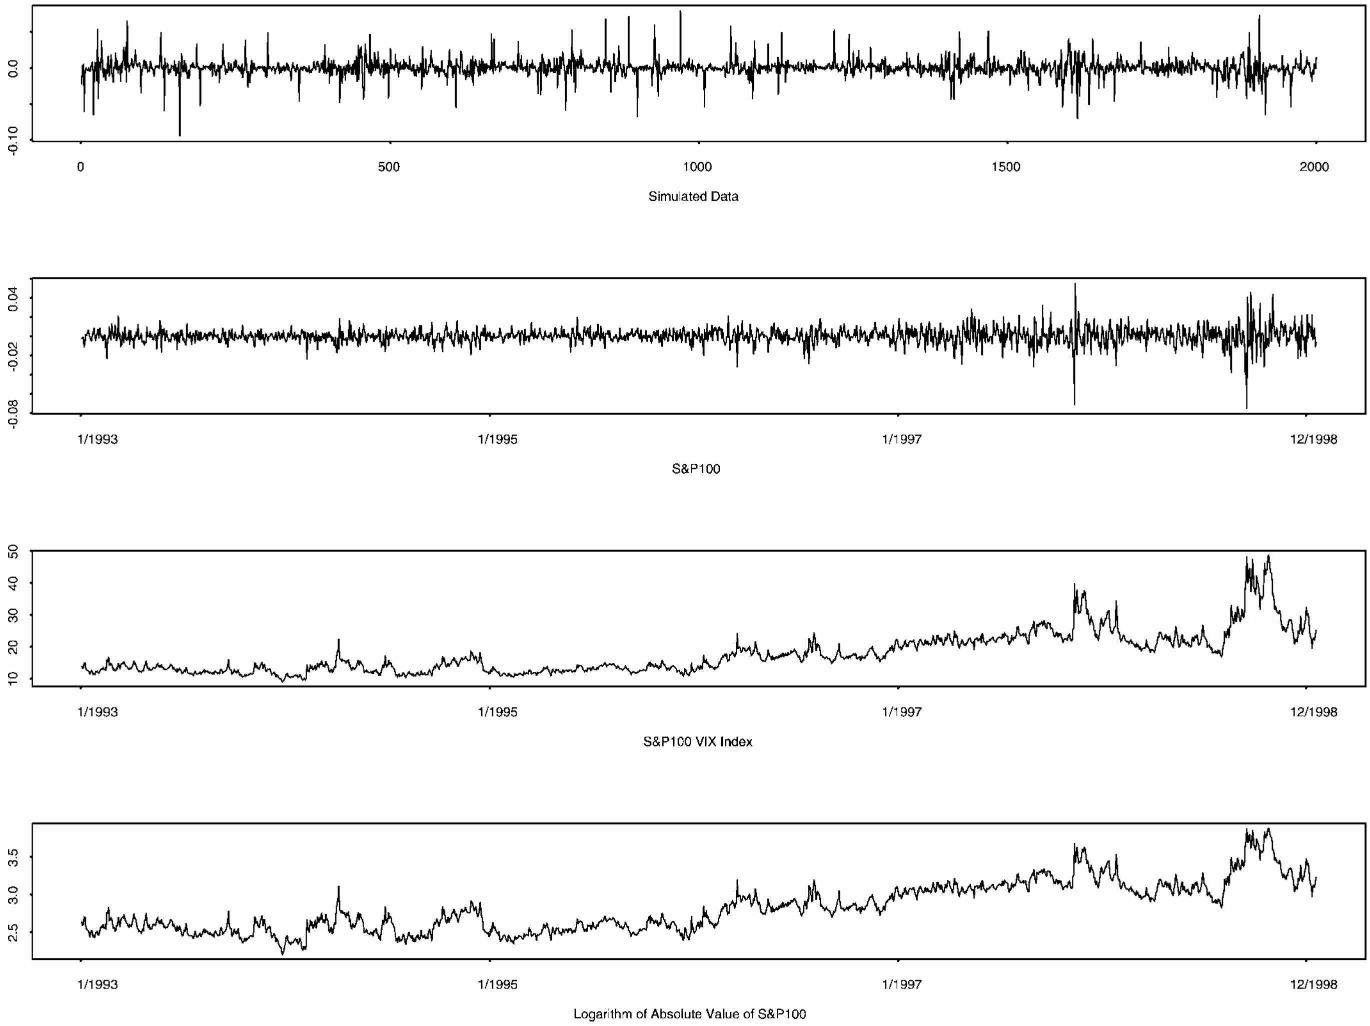

Figure 1. Time Series Plots for Simulated Data, S&P 100, VIX, Logarithm of Absolute Value of S&P 100 Returns.

Model 4. Two independent AR(1) processes as in Harvey et al. (1994), Shephard (1996), Gallant and Tauchen (2001), and Chernov, Gallant, Ghysels, and Tauchen (2003):

$$y_t|h_t = \exp(\mu/2 + h_t^{(1)}/2 + h_t^{(2)}/2)u_t,$$
$$u_t \overset{\text{i.i.d.}}{\sim} N(0,1),$$
$$h_t^{(1)}|h_{t-1}^{(1)}, \phi, \tau^2 = \phi h_{t-1}^{(1)} + v_t^{(1)}, \qquad v_t^{(1)} \overset{\text{i.i.d.}}{\sim} N(0,\tau^2),$$
$$h_t^{(2)}|h_{t-1}^{(2)}, \phi_2, \tau_2^2 = \phi_2 \cdot h_{t-1}^{(2)} + v_t^{(2)}, \qquad v_t^{(2)} \overset{\text{i.i.d.}}{\sim} N(0,\tau_2^2).$$

Model 5. This is Model 1 including a leverage or asymmetric effect by allowing for correlation between u_t and v_{t+1}, that is,

$$\begin{pmatrix} u_t \\ v_{t+1} \end{pmatrix} \overset{\text{i.i.d.}}{\sim} N\left\{\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho\tau \\ \rho\tau & \tau^2 \end{pmatrix}\right\}.$$

This effect is often observed in financial time series, for example, in time series of exchange rates and, even stronger, in stock market data. It reveals the market behavior, first discovered by Black (1976) and described in Engle and Ng (1993).

Although the correlation between u_t and v_{t+1} makes Kitagawa's algorithm not directly applicable, a simple rewrite of this model gives

$$y_t|h_t = \exp(h_t/2)u_t, \qquad u_t \overset{\text{i.i.d.}}{\sim} N(0,1),$$
$$\begin{aligned} h_t|h_{t-1}, y_{t-1}, \mu, \phi, \tau^2 &= \mu + \phi(h_{t-1} - \mu) \\ &\quad + \rho\tau\exp(-.5h_{t-1})y_{t-1} \\ &\quad + \tau\sqrt{1-\rho^2}\, w_t, \end{aligned}$$

where $w_t \overset{\text{i.i.d.}}{\sim} N(0,1)$ and $\operatorname{cor}(u_t, w_t) = 0$. Based on the new representation, steps 2a and 2b in Kitagawa's algorithm can be modified by:

2a. Generate M particles, called $v_t^{(j)}$, $j = 1, \ldots, M$, from a normal distribution with mean $\rho\tau\exp(-.5h_{t-1})y_{t-1}$ and variance $\tau^2(1-\rho^2)$.

2b. Obtain M particles by setting

$$p_t^{(j)} = F\big(f_{t-1}^{(j)}, v_t^{(j)}\big), \qquad j = 1, \ldots, M,$$

where $p_t^{(j)}$ can be regarded as independent draws from $p(h_t|y_{t-1})$.

Model 6. The SV + jumps model includes a jump component and lagged observations in the observation equation:

$$y_t|h_t, s_t, q_t, \beta = \beta y_{t-1} + s_t q_t + \exp(h_t/2)u_t,$$
$$u_t \overset{\text{i.i.d.}}{\sim} N(0,1),$$
$$h_t|h_{t-1}, \mu, \phi, \tau^2 = \mu + \phi(h_{t-1} - \mu) + v_t, \qquad v_t \overset{\text{i.i.d.}}{\sim} N(0, \tau^2),$$

where q_t follows a Bernoulli distribution that takes a value of 1 with probability κ and 0 with probability $1-\kappa$, and $\ln(1+s_t) \sim N(-\delta^2/2, \delta^2)$.

The underlying data are generated from this model using $\mu = -10$, $\phi = .96$, $\tau = .345$, $\beta = .1$, $\kappa = .08$, and $\delta = .03$.

Model 7. This model includes a jump component in the observation equation but without taking the lagged observations into consideration:

$$y_t|h_t, s_t, q_t = s_t q_t + \exp(h_t/2)u_t, \qquad u_t \overset{\text{i.i.d.}}{\sim} N(0,1),$$
$$h_t|h_{t-1}, \mu, \phi, \tau^2 = \mu + \phi(h_{t-1} - \mu) + v_t, \qquad v_t \overset{\text{i.i.d.}}{\sim} N(0, \tau^2).$$

Model 8. Gaussian observation errors are substituted by independent central Student t distributions with ν degrees of freedom:

$$y_t|h_t = \exp(h_t/2)u_t, \qquad u_t \overset{\text{i.i.d.}}{\sim} t_\nu,$$
$$h_t|h_{t-1}, \mu, \phi, \tau^2 = \mu + \phi(h_{t-1} - \mu) + v_t, \qquad v_t \overset{\text{i.i.d.}}{\sim} N(0, \tau^2).$$

5.2 Prior Distributions

For the parameters ϕ and τ^2 of the basic SV model, we follow exactly the prior specifications of Kim et al. (1998): $\tau^2 \sim$ Inverse-Gamma(2.5,.025), which has a mean of .167 and a standard deviation of .024. Defining $\phi = 2\phi^* - 1$, Kim et al. (1998) specified a beta distribution with parameters 20 and 1.5 for ϕ^*, which implies a mean of .86 and a standard deviation of .11. Following Kim et al. (1998), we choose an informative but reasonably flat prior distribution for the parameter μ, a normal distribution with mean -10 and variance 25.

For α in Model 2 a normal distribution with mean parameter $\mu_\alpha = 0$ and variance $\sigma_\alpha^2 = 10$ is specified.

For Model 3 we use the same prior for ϕ as for the basic SV model and center the prior for ψ around 0 using a uniform distribution on $[-1, 1]$.

In Model 4 we again use the same prior for ϕ as for the basic SV model and center a vague prior for ϕ_2 around 0 using a beta distribution with parameters 2 and 2.

The correlation parameter ρ in Model 5 is assumed to be uniformly distributed with support between -1 and 1.

As the parameter β in Model 6 is assumed to be small a priori, we use an informative normal distribution with hyperparameters $\mu_\beta = 0$ and $\sigma_\beta^2 = .2$. The parameter q_t represents the frequency of a jump occurrence with a Bernoulli distribution with parameter κ. Following Chib et al. (2002), we specify a Beta(2, 100) prior distribution, which implies a mean of .02 and suggests that a priori on average one jump in approximately every 50th observation. Finally, as in Chib et al. (2002), we assume that $\ln(\delta)$ follows a normal prior distribution with mean -3.07 and variance .149.

A well-known alternative to the direct fitting of many symmetric but nonnormal error distributions is through scale mixtures of normals (Andrews and Mallows 1974). Thus, in Model 8 we use the alternative mixture representation of a t_d distribution by

$$y_t \sim N\big(0, \exp(h_t)/w_t\big),$$
$$w_t \sim \frac{1}{\nu}\chi_\nu^2 = \text{Gamma}\left(\frac{\nu}{2}, \frac{\nu}{2}\right).$$

We choose a uniform distribution for ν on [2, 128] as in Chib et al. (2002).

5.3 Implementation in WinBUGS

WinBUGS is the BUGS version operating under Windows. A DIC module that automatically calculates values for DIC and related parameters is implemented in the latest WinBUGS version. Even without the DIC module, DIC is easily obtained from any MCMC output.

The first part of DIC, $\bar{D}$, is easily calculated using the MCMC output $\theta^{(i)}, i = 1, \ldots, N$. We simply calculate $D(\theta^{(i)})$ for $i = 1, \ldots, N$ and estimate $\bar{D}$ by the sample mean $(1/N) \times \sum_{i=1}^N D(\theta^{(i)})$. In practice, using BUGS, this is accomplished by adding the variable $D(\theta)$. For the second part, the effective number of parameters p_D, we only need to evaluate $D(\theta)$ at the sample posterior mean $\bar{\theta} = (1/N)\sum_{i=1}^N \theta^{(i)}$. WinBUGS offers several useful convergence-checking criteria available in an attached CODA (Convergence Diagnosis and Output Analysis Software for Gibbs sampling output; Best, Cowles, and Vines 1995) module running, for example, under S-Plus. It is necessary to check whether convergence has been achieved because it is crucial that the sample is taken from the stationary distribution. The CODA package consists of a selection of model-checking criteria, one of which is the Heidelberger–Welch test (Heidelberger and Welch 1983). All the results we report in this article are based on samples that have passed the Heidelberger–Welch convergence test for all parameters.

5.4 Results

In Table 1 we report means and standard errors (numbers in parentheses) of both prior and posterior distributions, as well as the computing time to generate 100 iterations for each of the eight models. The numbers in square brackets are the true values of the parameters. In Table 2 we report Chib's marginal likelihood, harmonic mean, and DIC together with $\bar{D}$ and p_D for each of the eight models as well as their associated rankings by each criterion. For SV Models 1–5, after a burn-in period of 50,000 iterations and a follow-up period of 250,000, we stored every 20th iteration. Due to higher posterior correlations among the parameters and thus slower convergence of the Gibbs sampler in the remaining models, we chose a burn-in period of 100,000 iterations, a follow-up period of 900,000, and stored every 40th iteration. All calculations were performed on a Pentium-III PC, 550 MHz, running the WinBUGS 131 version updated with the DIC tool.

From the examination of these two tables, we first note that the correct model (Model 6) is estimated by MCMC with reasonably accurate results for all six parameters. Moreover, the

Table 1. Parameter Estimates for Simulated Data

Parameter	*Distribution*	Model *1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*
μ	Prior	−10.0 (5.00)	−10.0 (5.00)	−10.0 (5.00)	−10.0 (5.00)	−10.0 (5.00)	−10.0 (5.00)	−10.0 (5.00)	−10.0 (5.00)
[−10]	Posterior	−9.836 (.0844)	−9.837 (.0847)	−9.845 (.0933)	−9.859 (.1615)	−9.837 (.0848)	−10.13 (.1499)	−10.08 (.1474)	−10.31 (.1512)
ϕ	Prior	.86 (.11)	.86 (.11)	.86 (.11)	.86 (.11)	.86 (.11)	.86 (.11)	.86 (.11)	.86 (.11)
[.96]	Posterior	.6502 (.0398)	.6496 (.0401)	.4146 (.0703)	.9494 (.0182)	.6518 (.0399)	.9353 (.0139)	.9330 (.0147)	.9457 (.0143)
τ	Prior	.12 (.05)	.12 (.05)	.12 (.05)	.12 (.05)	.12 (.05)	.12 (.05)	.12 (.05)	.12 (.05)
[.345]	Posterior	1.128 (.0617)	1.130 (.0622)	1.205 (.0567)	.2803 (.0646)	1.130 (.0624)	.3959 (.0438)	.3988 (.0450)	.3244 (.0462)
α	Prior	—	.00 (3.16)	—	—	—	—	—	—
	Posterior	—	3.7e^{-4} (.0001)	—	—	—	—	—	—
ψ	Prior	—	—	.00 (.58)	—	—	—	—	—
	Posterior	—	—	.2569 (.0754)	—	—	—	—	—
ϕ_2	Prior	—	—	—	.00 (.45)	—	—	—	—
	Posterior	—	—	—	.2135 (.0963)	—	—	—	—
τ_2	Prior	—	—	—	.38 (.16)	—	—	—	—
	Posterior	—	—	—	1.187 (.0606)	—	—	—	—
ρ	Prior	—	—	—	—	.00 (.58)	—	—	—
	Posterior	—	—	—	—	−.079 (.0423)	—	—	—
β	Prior	—	—	—	—	—	.00 (.45)	—	—
[.1]	Posterior	—	—	—	—	—	.0807 (.0121)	—	—
κ	Prior	—	—	—	—	—	.02 (.01)	.02 (.01)	—
[.08]	Posterior	—	—	—	—	—	.0727 (.0099)	.0691 (.0100)	—
δ	Prior	—	—	—	—	—	.05 (.02)	.05 (.02)	—
[.03]	Posterior	—	—	—	—	—	.0342 (.0026)	.0348 (.0028)	—
ν	Prior	—	—	—	—	—	—	—	65.0 (36.4)
	Posterior	—	—	—	—	—	—	—	2.563 (.1866)
Time (seconds)		4.30	4.84	7.54	7.38	31.31	36.45	31.29	42.59

correct model provides the smallest value for DIC as well as for the posterior mean of the deviance despite the fact that the effective number of parameters is not the smallest. We get only a slightly larger value of DIC for the SV + jumps model without lagged observations (Model 7). This is because differences between this model and the correct model are very small. Not surprisingly, this model is ranked second by DIC. All the other models clearly perform worse. For example, compared with

Table 2. Chib's Marginal Likelihood, Harmonic Mean, and DIC for Simulated Data

	ln L		*Harmonic mean*		*DIC*			
Model	*Value*	*Ranking*	*Value*	*Ranking*	*Value*	*Ranking*	$\bar{D}$	p_D
1	6,472.67	7	6,888.43	7	−13,442.5	7	−14,002.5	560.0
2	6,467.51	8	6,882.43	8	−13,441.9	8	−14,003.4	561.5
3	6,474.38	5	6,890.43	6	−13,463.3	4	−14,040.7	577.4
4	6,495.51	4	6,948.43	4	−13,496.2	3	−14,102.4	606.2
5	6,472.82	6	6,901.42	5	−13,453.9	5	−14,018.9	565.0
6	6,569.16	1	7,172.00	1	−14,450.0	1	−14,582.7	132.7
7	6,548.27	2	7,102.92	2	−14,362.0	2	−14,485.3	123.3
8	6,517.44	3	6,949.62	3	−13,448.0	6	−14,096.1	648.1

Table 3. Deviance and Harmonic Mean (HM) Summaries for Simulated Data

Model	$\bar{D}_{\min}$	$\bar{D}_{\max}$	$p_{D\,min}$	$p_{D\,max}$	$DIC_{\min}$	$DIC_{\max}$	$HM_{\min}$	$HM_{\max}$
1	−14,003.5	−14,000.1	557.3	560.5	−13,443.8	−13,441.5	6,871.41	6,890.55
2	−14,004.4	−14,001.1	561.5	563.4	−13,443.9	−13,439.7	6,860.43	6,886.79
3	−14,040.7	−14,036.8	574.4	577.8	−13,465.4	−13,461.2	6,870.30	6,893.43
4	−14,103.9	−14,100.2	603.3	606.4	−13,499.7	−13,495.6	6,939.66	6,957.15
5	−14,022.9	−14,018.9	565.0	568.1	−13,456.5	−13,453.9	6,890.98	6,912.18
6	−14,585.7	−14,580.0	131.7	133.4	−14,452.0	−14,448.6	7,139.71	7,172.00
7	−14,485.3	−14,479.5	123.3	124.7	−14,362.0	−14,354.8	7,091.90	7,111.97
8	−14,099.3	−14,096.0	648.1	650.2	−13,453.1	−13,448.1	6,933.15	6,949.62

DIC values of −14,450 and −14,362 for the two models with jumps, the basic SV model provides a DIC value of −13,442.5. In fact, the DIC margins among all the models excluding the jump models are reasonably small. For example, DIC of the third best model differs from that of the worst model by 54.3, whereas the difference between the second best and the third best is 865.8. Moreover, the effective number of parameters is much larger for all the models except the jump models and none of these models fits the data as well as the jump models, indicated by the highest value for the posterior mean of the deviance. Not surprisingly, the higher values of $\bar{D}$ and p_D add up to the higher DIC values.

Model 4, with two independent AR(1) components, gives a relatively good fit, being ranked the best fitting after the jump models by DIC and the best fitting after the jump and SV-t models by Chib's marginal likelihood. It can thus be considered as a good alternative to using SV models with jumps.

Another interesting result emerging from these two tables is the performance of DIC relative to Chib's marginal likelihood and the harmonic mean. Neither DIC nor the harmonic mean provides the same model ranking as Chib's marginal likelihood but the differences are not substantial. Differences between the two marginal likelihood methods and DIC are not surprising as the focus of DIC is different to that of the marginal likelihood methods, as explained in detail in the previous sections.

The computing time to generate 100 iterations suggests that the MCMC program runs substantially slower for the SV Models 5–8 than for the SV Models 1–4. This is because most of the full conditional distributions for SV Models 5–8 are no longer log-concave and a Metropolis–Hastings updating step is needed. To conserve space, the correlograms are not plotted in the article, but they are available from the authors upon request.

Comparison between DIC and Chib's marginal likelihood reveals that the mixture normal-Gamma t SV model (Model 8) is the only cause of the discrepancy. Here it is helpful to divide DIC into a pure measure of fit $D(\bar{\theta})$ and a measure of complexity $2p_D$ as in (7) to see that the t SV model is heavily penalized by its high effective number of parameters. Considering $D(\bar{\theta})$ gives a value of −14,715.4 for the true (Model 6) and a value of −14,744.1 for the t SV model (Model 8). Thus, the t SV model provides a better fit but its high complexity tips the scales. Although not reported, we have also estimated the nonscale mixture t SV model and found that the performance of these two representations are quite different. The nonscale mixture t SV model performs even worse than the scale mixture t SV model according to DIC. It has been recognized that different mixture distributions can generate different DIC values, due to the fact that different mixture distributions correspond to different prediction problems, and more research and experience is needed as to the performance of DIC in the area of mixture models (Spiegelhalter et al. 2002).

Table 3 shows the smallest and largest values for DIC and the harmonic mean, the number of effective parameters p_D, and the goodness of fit $\bar{D}$, respectively, obtained for six runs for each of the seven models. It serves to demonstrate that DIC varies from one run to another but is reasonably stable across runs. This is in contrast to the well-known instability problem of the harmonic mean, which is apparent from the large discrepancies between the smallest and largest values for the harmonic mean. However, the reader should note that slightly modified estimators of the harmonic mean as proposed by Newton and Raftery (1994) are much more stable and do not suffer from the lack of a central limit theorem.

6. AN EMPIRICAL STUDY

6.1 The Data

The dataset consists of 1,512 mean-corrected daily continuously compounded returns, y_t, in decimals, on the Standard & Poors (S&P) 100 index, covering the period of time between January 1993 and December 1998. The S&P 100 index returns have been used often in the literature. For instance, Blair, Poon, and Taylor (2001a) estimated the GJR-GARCH model proposed by Glosten, Jagannathan, and Runkle (1993) based on the S&P 100 index returns for four different sample periods from March 1984 to December 1998, one of which is identical to that in this article. We also use data from the Chicago Board Options Exchange Market Volatility Index (VIX) for the same period of time as a covariate, measuring the so-called implied volatility. For a detailed explanation of the Chicago Board Options Exchange Market Volatility Index, the reader is referred to Hol and Koopman (2000) and Fleming, Ostdiek, and Whaley (1995). Both data series are plotted in the second and third panels of Figure 1.

6.2 The Models and Prior Distributions

In this section we fit the models introduced in Section 5 to the preceding dataset. We drop Model 4 from the list due to a great deal of convergence problems that we have encountered (it may be possible to achieve convergence by using different parameterizations or using different MCMC algorithms, however). Instead we consider as an additional extension a model that includes implied volatility:

Model 9. This model is very similar to the SVX model introduced in Hol and Koopman (2000), which includes implied volatility as expressed by an additional covariate x_t:

$$y_t|h_t = \exp(h_t/2)u_t, \qquad u_t \overset{\text{i.i.d.}}{\sim} N(0,1),$$

$$h_t|h_{t-1},\mu,\phi,\tau^2,\lambda = \mu + \phi(h_{t-1}-\mu) + \lambda(x_t - \bar{x}) + v_t,$$

$$v_t \overset{\text{i.i.d.}}{\sim} N(0,\tau^2).$$

The implied volatility is used in this model as an alternative source for predicting volatility and is based on calculations of option price models. The specification of the variance equation is motivated by the empirical result that implied volatilities contain useful information in forecasting future volatilities (see, e.g., Blair, Poon, and Taylor 2001b). In the last panel of Figure 1, we plot the logarithm of the absolute value of S&P 100 returns, which is regarded as an approximation of unobserved log-volatility. It can be seen that the VIX and the logarithm of the absolute value of S&P 100 returns are highly correlated. Note that we demean the observations in vector x_t for convergence purposes.

A priori, λ is assumed to be uniformly distributed in the interval $[-1, 1]$. Due to the inclusion of the implied volatility, it is not clear a priori whether the log-volatility h_t is still highly persistent. Instead of using a rather informative prior of a beta distribution with parameters 20 and 1.5 for ϕ^*, we choose a less informative prior for ϕ^*, namely, a uniform distribution with support between 0 and 1.

6.3 Results

In Table 4 we report means and standard errors (numbers in parentheses) of both prior and posterior distributions, as well as the computing time to generate 100 iterations for each of the eight models. For Models 1–5, after a burn-in period of 50,000 iterations and a follow-up period of 250,000, we stored every 20th iteration. In the remaining models, we chose a burn-in period of 100,000 iterations, a follow-up period of 900,000, and stored every 40th iteration.

Table 4. Parameter Estimates for S&P 100 Data

		Model							
Parameter	*Distribution*	*1*	*2*	*3*	*5*	*6*	*7*	*8*	*9*
μ	Prior	−10.0 (5.00)	−10.0 (5.00)	−10.0 (5.00)	−10.0 (5.00)	−10.0 (5.00)	−10.0 (5.00)	−10.0 (5.00)	−10.0 (5.00)
	Posterior	−9.971 (.2573)	−9.956 (.2408)	−9.986 (.2663)	−9.951 (.2019)	−10.06 (.2889)	−10.07 (.2897)	−10.34 (.3213)	−9.942 (.0460)
ϕ	Prior	.86 (.11)	.86 (.11)	.86 (.11)	.86 (.11)	.86 (.11)	.86 (.11)	.86 (.11)	.00 (.58)
	Posterior	.9803 (.0088)	.9789 (.0094)	.8375 (.1525)	.9743 (.0100)	.9868 (.0070)	.9873 (.0069)	.9923 (.0044)	−.2745 (.1155)
τ	Prior	.12 (.05)	.12 (.05)	.12 (.05)	.12 (.05)	.12 (.05)	.12 (.05)	.12 (.05)	.12 (.05)
	Posterior	.1674 (.0297)	.1729 (.0317)	.1886 (.0442)	.1947 (.0319)	.1331 (.0284)	.1302 (.0286)	.1005 (.0189)	.4375 (.0763)
α	Prior	—	.00 (3.16)	—	—	—	—	—	—
	Posterior	—	1.03e^{-4} (1.7e^{-4})	—	—	—	—	—	—
ψ	Prior	—	—	.00 (.58)	—	—	—	—	—
	Posterior	—	—	.1413 (.1495)	—	—	—	—	—
ρ	Prior	—	—	—	.00 (.58)	—	—	—	—
	Posterior	—	—	—	−.4139 (.0860)	—	—	—	—
β	Prior	—	—	—	—	.00 (.45)	—	—	—
	Posterior	—	—	—	—	.0050 (.0263)	—	—	—
κ	Prior	—	—	—	—	.02 (.01)	.02 (.01)	—	—
	Posterior	—	—	—	—	.0114 (.0076)	.0115 (.0075)	—	—
δ	Prior	—	—	—	—	.05 (.02)	.05 (.02)	—	—
	Posterior	—	—	—	—	.0315 (.0119)	.0324 (.0121)	—	—
ν	Prior	—	—	—	—	—	—	65.0 (36.4)	—
	Posterior	—	—	—	—	—	—	7.306 (1.532)	—
λ	Prior	—	—	—	—	—	—	—	.0 (.58)
	Posterior	—	—	—	—	—	—	—	.1527 (.0159)
Time (seconds)		3.25	3.66	5.70	23.67	27.56	23.65	32.20	4.27

Table 5. Chib's Marginal Likelihood, Harmonic Mean, and DIC for S&P 100 Data

	ln L		Harmonic mean		DIC			
Model	Value	Ranking	Value	Ranking	Value	Ranking	$\bar{D}$	p_D
1	5,227.45	6	5,270.40	8	−10,529.6	8	−10,613.1	83.5
2	5,223.17	8	5,276.64	7	−10,530.0	7	−10,616.9	86.9
3	5,226.17	7	5,277.29	6	−10,531.5	6	−10,616.7	85.2
5	5,229.87	5	5,300.43	3	−10,610.9	4	−10,936.3	325.4
6	5,241.54	3	5,279.58	5	−10,649.7	2	−10,677.2	27.5
7	5,243.76	1	5,288.27	4	−10,653.5	1	−10,679.3	25.8
8	5,242.38	2	5,331.00	1	−10,566.1	5	−10,788.9	222.8
9	5,234.48	4	5,301.30	2	−10,618.8	3	−10,741.3	122.5

From Table 4 it can be seen that the estimated means and standard deviations for the parameters appear quite reasonable and comparable with previous estimates in the literature. For instance, in Model 1, the volatility process is estimated to be highly persistent. In Model 5 the posterior mean of ρ is $-.4139$ with the upper limit of the 95% posterior credibility interval less than 0. This suggests that the leverage effect is an important feature for the S&P 100 index. The parameter estimates for the two SV + jumps models provide similar results for those parameters already covered by the SV models without jumps. As already observed in Chib et al. (2002), we note that the jump parameters κ and δ are not as precisely estimated as other parameters. However, they are well identified as their posterior distributions are substantially different from their prior distributions. The posterior mean of the jump intensity κ is .011, which means an average daily probability of 1.1% of a jump occurring. This implies that a jump can be expected to occur roughly every 90th day. The standard deviation of the jump size is about .03; that is, daily jumps are usually around 6%.

In Model 8 the posterior mean of ν is 7.306 and similar to the values of 7.7 and 8.9 for the S&P 500 index in Sandmann and Koopman (1998) and Chib et al. (2002), respectively. The posterior mean of λ in Model 9 indicates that the implied volatility contains important information about the volatility process. Interestingly, allowing for the implied volatility as a covariate induces a negative posterior mean of the autoregressive coefficient in the model. This finding is similar to what was obtained in Hol and Koopman (2000) based on an S&P 100 index for a different period.

In Table 5 we report Chib's marginal likelihood, harmonic mean, and DIC together with $\bar{D}$ and p_D for each of the eight models as well as their associated rankings by each criterion. The most adequate models to describe the S&P 100 according to DIC are the jump model without lagged observations (Model 7) and the jump model with lagged observations (Model 6), followed by the implied volatility model (Model 9) and the model including the leverage effect (Model 5). Although the posterior means of the deviance for the jump models are higher than those of most of the other models, the effective number of parameters is much lower. The effective number of parameters is around 26 for the jump models, which is less than one-third of the effective number of parameters for the closest competitor. Model 5 has the lowest posterior means of the deviance, which suggests the best fit to the data. However, its effective number of parameters is much higher than that of the other models. In particular, it is more than 10 times larger than that of the jump models. This renders a larger value of DIC.

As for the simulated data, neither DIC nor the harmonic mean provides the same model ranking as Chib's marginal likelihood. According to Chib's marginal likelihood, the strength of evidence to distinguish between the models is much weaker for the S&P 100 data than for the simulated data. For example, the marginal likelihood values from the second best model and the third best model only differ by .84, which is not worth more than a bare mention according to Jeffrey's Bayes factor scale [exp(.84) = 2.316]. Nonetheless, both DIC and Chib's marginal likelihood select Model 7 (i.e., the jump model without lagged observations) as the best performing model, whereas the harmonic mean picks Model 8 (i.e., the t SV model).

A close look at Table 5 reveals that the mixture normal-Gamma t SV model (i.e., Model 8) is the major cause of the discrepancy between the DIC ranking and Chib's marginal likelihood ranking. This is a similar finding to the simulated data. Another minor discrepancy arises from the first three models. Chib's marginal likelihood ranks Model 2 the worst model, whereas DIC ranks Model 1 the worst.

Table 6 shows the smallest and largest values for DIC and the harmonic mean, the number of effective parameters p_D and the goodness of fit $\bar{D}$, respectively, obtained for six runs for each of the seven models. Again it demonstrates that DIC varies from one run to another but is reasonably stable across runs and DIC is more stable than the harmonic mean. Also, it can be seen

Table 6. Deviance and Harmonic Mean (HM) Summaries for S&P 100 Data

Model	$\bar{D}_{\min}$	$\bar{D}_{\max}$	$p_{D\,min}$	$p_{D\,max}$	$DIC_{\min}$	$DIC_{\max}$	$HM_{\min}$	$HM_{\max}$
1	−10,617.1	−10,611.6	81.9	85.8	−10,531.3	−10,527.4	5,266.96	5,275.31
2	−10,618.4	−10,615.9	85.1	88.2	−10,531.4	−10,529.6	5,271.53	5,279.98
3	−10,621.3	−10,613.1	83.1	88.8	−10,532.8	−10,530.0	5,270.33	5,277.29
5	−10,941.4	−10,934.9	323.0	328.6	−10,613.3	−10,610.9	5,297.65	5,308.88
6	−10,681.5	−10,674.9	24.9	29.8	−10,656.7	−10,645.1	5,279.58	5,287.72
7	−10,680.7	−10,675.9	24.1	30.3	−10,655.1	−10,645.7	5,278.56	5,288.27
8	−10,791.5	−10,787.3	222.2	226.1	−10,565.5	−10,566.9	5,324.47	5,331.00
9	−10,741.5	−10,738.8	120.8	124.0	−10,618.8	−10,617.1	5,298.71	5,302.28

Table 7. Sensitivity of DIC and Chib's Marginal Likelihood to the Prior

	Model 1, Prior 2	Model 5, Prior 2	Model 7 Prior 2	Model 7 Prior 3
μ	−9.970 (.2543)	−9.963 (.2240)	−10.10 (.3250)	−10.20 (.3232)
ϕ	.9806 (.0092)	.9768 (.0098)	.9886 (.00678)	.9887 (.00694)
τ	.1680 (.0327)	.1865 (.0317)	.1271 (.0271)	.1266 (.0290)
ρ	—	−.4145 (.0883)	—	—
κ	—	—	.0107 (.0064)	.0121 (.0080)
δ	—	—	.0337 (.0113)	.0298 (.0120)
DIC	−10,530.7	−10,618.1	−10,646.5	−10,659.0
ln L	5,225.39	5,228.03	5,242.27	5,240.88

that the ranges of DIC overlap with each other for the first three models. This explains why the first three models are difficult to distinguish.

6.4 Robustness Check

In this section we examine the implications of alternative prior distributions on DIC and Chib's marginal likelihood. In particular, we focus on a subset of hyperparameters, namely, ϕ and κ. Also, for brevity we only consider a subset of the models, namely, the basic SV model (Model 1), the SV model with a leverage effect (Model 5), and the SV + jumps model without lagged observations (Model 7). Following Chib et al. (2002), we consider the following two alternative priors:

- Prior 2: $\phi^* \sim \mathrm{U}(0, 1)$.
- Prior 3: $\phi^* \sim \mathrm{U}(0, 1)$, $\kappa \sim$ Beta with mean .0385 and standard error .0264.

We reestimate all three models with Prior 2 and Model 7 with Prior 3 and calculate DIC and Chib's marginal likelihood accordingly. The posterior means, standard errors, DIC, and the marginal likelihood are reported in Table 7. A comparison with the results in Table 4 shows that Prior 2 yields a posterior distribution that is almost identical to that with the original prior and that Prior 3 yields a posterior distribution that is reasonably close to that with the original prior. More important, DIC seems quite robust to the change of prior. Moreover, it preserves the ranking of the models considered and the ranking is consistent with that based on the marginal likelihood.

7. CONCLUSION

In this article we have explored the practical performance of DIC as a model selection criterion for comparing various stochastic volatility models. DIC is a Bayesian version of the classical deviance for model assessment. It is particularly suited to compare Bayesian models whose posterior distributions have been obtained using MCMC simulation. Similar to AIC and BIC, DIC comprises two parts, a goodness-of-fit measure, the posterior distribution of the deviance, and a penalty term, the effective number of parameters, measuring complexity. Using the concept of *effective number of parameters*, DIC can be used in complex hierarchical models where the number of unknowns often exceeds the number of observations and the number of free parameters is not well defined. This is in contrast to AIC and BIC, where the number of free parameters needs to be specified. DIC has been implemented as a tool in the BUGS software package.

We carry out a simulation study using an SV + jumps model as the true model. Our estimation results with respect to the simulated data are quite accurate for the true model, and DIC clearly identifies the correct model out of eight different alternatives. If one were to omit the mixture t SV model, DIC would give the same model ranking as Chib's marginal likelihood. By comparing eight different SV models for the S&P 100 index, comprising 1,512 observations from 1993 to 1998, the jump volatility model without lagged observations turns out to be the most adequate as indicated by both DIC and Chib's marginal likelihood. The Monte Carlo error of DIC is fairly low for all the models, indicating a stable performance for model comparison purposes. Finally, DIC appears robust to the change of prior distributions.

ACKNOWLEDGMENTS

The authors thank an Associate Editor and two anonymous referees for constructive comments that have substantially improved the article. They also acknowledge comments from Nicky Best, Genshiro Kitagawa, Angelika van der Linde, Adrian Pagan, Peter Phillips, Werner Ploberger, and seminar participants at the 2001 Australasian Econometric Society Meeting, the New Zealand Econometric Study Group Meeting in Auckland, and the National University of Singapore (both Department of Economics and Center for Financial Engineering). Andreas Berg gratefully acknowledges the support of his research by the DAAD, personal number D/99/19367 within a "DAAD Doktorandenstipendium im Rahmen des gemeinsamen Hochschulsonderprogramms III von Bund und Laendern" and by the Royal Society of New Zealand Marsden Fund for a Ph.D. scholarship under grant 01-UOA-015. The work of Renate Meyer was supported by the Royal Society of New Zealand

Marsden Fund and the University of Auckland Research Committee. Jun Yu acknowledges financial support from the University of Auckland Research Committee and the Royal Society of New Zealand Marsden Fund under grant 01-UOA-015.

[Received October 2001. Revised May 2003.]

REFERENCES

Akaike, H. (1973), "Information Theory and an Extension of the Maximum Likelihood Principle," in *Proceedings of the Second International Symposium on Information Theory*, eds. B. N. Petrov and F. Csaki, Budapest: Akademiai Kiado, pp. 267–281.

Andersen, T. G., Chung, H.-J., and Sørensen, B. E. (1999), "Efficient Method of Moments Estimation of a Stochastic Volatility Model: A Monte Carlo Study," *Journal of Econometrics*, 91, 61–87.

Andersen, T. G., and Sørensen, B. E. (1996), "GMM Estimation of a Stochastic Volatility Model: A Monte Carlo Study," *Journal of Business & Economic Statistics*, 14, 329–352.

Andrews, D. F., and Mallows, C. L. (1974), "Scale Mixtures of Normal Distributions," *Journal of the Royal Statistical Society*, Ser. B, 36, 99–102.

Berger, J. O., and Pericchi, L. R. (1996), "The Intrinsic Bayes Factor for Model Selection and Prediction," *Journal of the American Statistical Association*, 91, 109–122.

Best, N., Cowles, M. K., and Vines, K. (1995), "CODA Convergence Diagnosis and Output Analysis Software for Gibbs Sampling Output Version 0.30," Cambridge, U.K.: MRC Biostatistics Unit.

Black, F. (1976), "Studies of Stock Market Volatility Changes," *Proceedings of the American Statistical Association, Business and Economic Statistics Section*, 177–181.

Blair, B. J., Poon, S.-H., and Taylor, S. J. (2001a), "Modelling S&P 100 Volatility: The Information Content of Stock Returns," *Journal of Banking and Finance*, 25, 1665–1679.

——— (2001b), "Forecasting S&P 100 Volatility: The Incremental Information Content of Implied Volatilities and High Frequency Returns," *Journal of Econometrics*, 105, 5–26.

Bollerslev, T. (1986), "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 31, 307–327.

Brenner, R. J., Harjes, R. H., and Kroner, K. F. (1996), "Another Look at Models of the Short-Term Interest Rate," *Journal of Financial and Quantitative Analysis*, 31, 85–107.

Carlin, B. P., and Louis, T. A. (1996), *Bayes and Empirical Bayes Methods for Data Analysis*, *Monographs on Statistics and Applied Probability*, 69.

Chacko, G., and Viceira, L. M. (2003), "Spectral GMM Estimation of Continuous-Time Processes," *Journal of Econometrics*, 116, 259–292.

Chan, K. C., Karolyi, G. A., Longstaff, F. A., and Sanders, A. B. (1992), "An Empirical Comparison of Alternative Models of the Short-Term Interest Rate," *The Journal of Finance*, 47, 1209–1227.

Chernov, M., Gallant, A. R., Ghysels, E., and Tauchen, G. (2003), "Alternative Models for Stock Price Dynamics," *Journal of Econometrics*, 116, 225–257.

Chib, S. (1995), "Marginal Likelihood from the Gibbs Output," *Journal of the American Statistical Association*, 90, 1313–1321.

Chib, S., Nardari, F., and Shephard, N. (1999), "Analysis of High Dimensional Multivariate Stochastic Volatility Models," unpublished paper.

——— (2002), "Markov Chain Monte Carlo Methods Stochastic Volatility Models," *Journal of Econometrics*, 108, 281–316.

Dempster, A. P. (1974), "The Direct Use of Likelihood for Significance Testing," in *Proceedings of Conference on Foundational Questions in Statistical Inference*, University of Aarhus, pp. 335–352.

Doucet, A., de Freitas, N., and Gordon, N. (2001), *Sequential Monte Carlo Methods in Practice*, New York: Springer-Verlag.

Duffie, D., and Singleton, K. J. (1993), "Simulated Moments Estimation of Markov Models of Asset Prices," *Econometrica*, 61, 929–952.

Engle, R. F. (1982), "Autoregressive Conditional Heteroscedasticity With Estimates of the Variance of United Kingdom Inflation," *Econometrica*, 50, 987–1007.

Engle, R. F., and Ng, V. K. (1993), "Measuring and Testing the Impact of News on Volatility," *The Journal of Finance*, 48, 1749–1778.

Fernandez, C., Ley, E., and Steel, M. F. J. (2001), "Benchmark Priors for Bayesian Model Averaging," *Journal of Econometrics*, 100, 381–427.

Fleming, J., Ostdiek, B., and Whaley, R. E. (1995), "Predicting Stock Market Volatility: A New Measure," *Journal of Futures Markets*, 15, 265–302.

Gallant, A. R., and Tauchen, G. (1996), "Which Moments to Match," *Econometric Theory*, 12, 657–681.

——— (2001), "Efficient Method of Moments," working paper, University of North Carolina, Dept. of Economics. Available at www.unc.edu/arg.

Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (1996), *Bayesian Data Analysis*, London: Chapman & Hall.

George, E. I., and Foster, D. P. (2001), "Calibration and Empirical Bayes Variable Selection," *Biometrika*, 87, 731–747.

Ghysels, E., Harvey, A. C., and Renault, E. (1996), "Stochastic Volatility," in *Statistical Models in Finance*, eds. C. R. Rao and G. S. Maddala, Amsterdam: North-Holland, pp. 119–191.

Gilks, W. R., Richardson, S., and Spiegelhalter, D. J. (1996), *Markov Chain Monte Carlo in Practice*, London: Chapman & Hall.

Glosten, L. R., Jagannathan, R., and Runkle, D. E. (1993), "On the Relation Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks," *The Journal of Finance*, 48, 1779–1801.

Gordon, N. J., Salmond, D. J., and Smith, A. E. M. (1993), "A Novel Approach to Nonlinear and Non-Gaussian Bayesian State Estimation," *IEEE Proceedings* F, 140, 107–133.

Han, C., and Carlin, B. P. (2001), "Markov Chain Monte Carlo Methods for Computing Bayes Factors: A Comparative Review," *Journal of the American Statistical Association*, 96, 1122–1132.

Harvey, A. C., Ruiz, E., and Shephard, N. (1994), "Multivariate Stochastic Variance Models," *Review of Economic Studies*, 61, 247–264.

Heidelberger, P., and Welch, P. (1983), "Simulation Run Length Control in the Presence of an Initial Transient," *Operations Research*, 31, 1109–1144.

Hoeting, J. A., Madigan, A., Raftery, A. E., and Volinsky, C. T. (1999), "Bayesian Model Averaging: A Tutorial," *Statistical Science*, 14, 383–417.

Hol, E., and Koopman, S. J. (2000), "Forecasting the Variability of Stock Index Returns With Stochastic Volatility Models and Implied Volatility," Discussion paper, TI 2000-104/4, Tinbergen Institute.

Hull, J., and White, A. (1987), "The Pricing of Options on Asset With Stochastic Volatilities," *The Journal of Finance*, 42, 281–300.

Jacquier, E., Polson, N. G., and Rossi, P. E. (1994), "Bayesian Analysis of Stochastic Volatility Models" (with discussion), *Journal of Business & Economic Statistics*, 12, 371–417.

Kass, R. E., and Raftery, A. E. (1995), "Bayes Factors," *The Journal of the American Statistical Association*, 90, 773–795.

Key, J. T., Pericchi, L. R., and Smith, A. F. M. (1999), "Bayesian Model Choice: What and Why?" in *Bayesian Statistics* (Vol. 6), eds. J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith, Oxford, U.K.: Oxford University Press, pp. 343–370.

Kim, S., Shephard, N., and Chib, S. (1998), "Stochastic Volatility: Likelihood Inference Comparison With ARCH Models," *Review of Economic Studies*, 65, 361–393.

Kitagawa, G. (1996), "Monte Carlo Filter and Smoother for Gaussian Nonlinear State Space Models," *Journal of Computational and Graphical Statistics*, 5, 1–25.

Knight, J. L., Satchell, S. E., and Yu, J. (2002), "Estimation of Stochastic Volatility Model by the Empirical Characteristic Function Method," *Australian and New Zealand Journal of Statistics*, 44, 319–335.

Lavine, M., and Schervish, M. J. (1999), "Bayes Factors: What They Are and What They Are Not," *The American Statistician*, 53, 119–122.

Melino, A., and Turnbull, S. M. (1990), "Pricing Foreign Currency Options With Stochastic Volatility," *Journal of Econometrics*, 45, 239–265.

Meyer, R., and Yu, J. (2000), "BUGS for a Bayesian Analysis of Stochastic Volatility Models," *The Econometrics Journal*, 3, 198–215.

Newton, M. A., and Raftery, A. E. (1994), "Approximate Bayesian Inference With the Weighted Likelihood Bootstrap" (with discussion), *Journal of the Royal Statistical Society*, Ser. B, 56, 3–48.

O'Hagan, A. (1991), "Contribution to the Discussion of 'Posterior Bayes Factors,'" *Journal of the Royal Statistical Society*, Ser. B, 53, 136.

——— (1994), *Kendall's Advanced Theory of Statistics* (Vol. 2), London: Edward Arnold.

Pitt, M., and Shephard, N. (1999a), "Filtering via Simulation: Auxiliary Particle Filter," *Journal of the American Statistical Association*, 94, 590–599.

——— (1999b), "Time Varying Covariances: A Factor Stochastic Volatility Approach" (with discussion), in *Bayesian Statistics* (Vol. 6), eds. J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith, Oxford, U.K.: Oxford University Press, pp. 547–570.

Sandmann, G., and Koopman, S. J. (1998), "Estimation of Stochastic Volatility Models via Monte Carlo Maximum Likelihood," *Journal of Econometrics*, 87, 271–301.

Schwarz, G. (1978), "Estimating the Dimension of a Model," *Annals of Statistics*, 6, 461–464.

Shephard, N. (1996), "Statistical Aspects of ARCH and Stochastic Volatility," in *Time Series Models in Econometrics*, *Finance and Other Fields*, eds. D. R. Cox, D. V. Hinkley, and O. E. Barndoff-Nielson, London: Chapman and Hall, pp. 1–67.

Silverman, B. W. (1986), *Density Estimation for Statistics and Data Analysis*, London: Chapman & Hall.

Singleton, K. J. (2001), "Estimation of Affine Asset Pricing Models Using the Empirical Characteristic Function," *Journal of Econometrics*, 102, 111–141.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and van der Linde, A. (2002), "Bayesian Measures of Model Complexity and Fit" (with discussion), *Journal of the Royal Statistical Society*, Ser. B, 64, 583–639.

Spiegelhalter, D. J., Thomas, A., Best, N. G., and Gilks, W. R. (1996), *BUGS 0.5, Bayesian Inference Using Gibbs Sampling, Manual (Version II)*, Cambridge, U.K.: MRC Biostatistics Unit.

Taylor, S. J. (1982), "Financial Returns Modelled by the Product of Two Stochastic Processes—A Study of the Daily Sugar Prices 1961–75," in *Time Series Analysis: Theory and Practice* (Vol. 1), ed. O. D. Anderson, Amsterdam: North-Holland, pp. 203–226.

Zhu, L., and Carlin, B. P. (2000), "Comparing Hierarchical Models for Spatio-Temporally Misaligned Data Using the Deviance Information Criterion," *Statistics in Medicine*, 19, 2265–2278.