Estimation of Hyperbolic Diffusion Using MCMC Method

Y.K. Tse
School of Economics and Social Sciences
Singapore Management University
Singapore 259756
yktse@smu.edu.sg

Xibin Zhang
Department of Econometrics and Business Statistics
Monash University
Clayton, Victoria 3800, Australia
bill.zhang@buseco.monash.edu.au

Jun Yu
Department of Economics, University of Auckland
Auckland, New Zealand
j.yu@auckland.ac.nz

First draft: September 2002
Revised: September 2003

Abstract: In this paper we propose a Bayesian method to estimate the hyperbolic diffusion model. The approach is based on the Markov chain Monte Carlo (MCMC) method with the likelihood of the discretized process as the approximate posterior likelihood. We demonstrate that the MCMC method provides a useful tool in analyzing hyperbolic diffusions. In particular, quantities of posterior distributions obtained from the MCMC outputs can be used for statistical inference. The MCMC method based on the Milstein scheme is found to perform well with good mixing properties, while the Euler scheme is unsatisfactory. Our simulation study shows that the hyperbolic diffusion exhibits many of the stylized facts about asset returns documented in the discrete-time financial econometrics literature, such as the Taylor effect, a slowly declining autocorrelation function of the squared returns, and thick tails.

Key Words: ARCH, Euler approximation, Hyperbolic diffusion, Long Memory, Markov chain Monte Carlo, Milstein approximation

Acknowledgment: The third author gratefully acknowledges financial support from the Royal Society of New Zealand Marsden Fund under number 01-UOA-015. Some of the work was done when the third author was visiting the National University of Singapore in 2001, whose hospitality is gratefully acknowledged.

1 Introduction

In the finance literature the geometric Brownian motion has been used as a classical model to describe stock price movements. Though useful as a benchmark model in option pricing and other theories, the geometric Brownian motion is unable to reconcile with many known statistical regularities of stock returns, such as excess kurtosis, clustering of volatility and long memory. To this effect, other processes have been suggested, such as jump diffusions (Kou (2002)), stochastic volatility models (Heston (1993)), stochastic volatility plus jumps (Eraker, Johannes, and Polson (2003)), and time-changed Levy process (Carr and Wu (2003)). As a nonlinear diffusion process, the hyperbolic diffusion model proposed by Bibby and Sorensen (1997) has received some attention (see, e.g., Rydberg (1999)). Bibby and Sorensen (1997) demonstrated some success in fitting the stationary distribution of the hyperbolic diffusion to some stock price data, and provided the theory in applying the hyperbolic diffusion to option pricing.

Although the stationary distribution of the hyperbolic diffusion process follows the hyperbolic distribution and hence has a closed-form expression, the transition density has no closed-form solution. Due to the lack of knowledge of the transition density, econometric estimation of the model using the *exact* likelihood approach is intractable, though an approximate likelihood method based on discretization may be adopted. To circumvent this difficulty Bibby and Sorensen (1997) estimated the hyperbolic diffusion using the martingale estimating function method. However, although the estimator based on martingale estimating functions is consistent and asymptotically normally distributed, it is inefficient in general. Furthermore, computation of the standard errors of the resulting estimates is difficult and requires techniques such as parametric bootstrapping.

In this paper we propose to use the Markov chain Monte Carlo (MCMC) method to estimate the parameters of the hyperbolic diffusion with the discretized likelihood of the diffusion process as an approximate posterior. Like the maximum likelihood (ML) approach in the classical framework, the MCMC method offers a full likelihood-based inference based on Bayesian analysis. In the case of the hyperbolic diffusion, the discretized approximate ML approach is found to encounter difficulties in numerical convergence. The MCMC method, however, provides a general mechanism to sample the parameter vector from its posterior distribution, and hence avoids the need for numerical optimization and enables exact finite-sample inferences via Monte Carlo methods.

In the financial econometrics literature a number of stylized facts have been well documented in describing the statistical properties of equity return series. Several models in the discrete-time domain have been found to be able to generate time series with such stylized properties. In the continuous-time domain, however, the success in this aspect has been much weaker.[1] On the other hand, while most empirical works rely on discrete-time models due to their simplicity in estimation, theories in option pricing are usually based on continuous-time models. The hyperbolic diffusion is a promising continuous-time model that describes empirically equity price data and can be applied to option pricing. Empirical illustrations reported by Rydberg (1999) show that a member of the generalized hyperbolic diffusion can induce long-memory features in the squared return.[2] In this paper we report the ability of the hyperbolic diffusion in reproducing other stylized facts documented in the financial econometrics literature.

This paper is organized as follows. Section 2 reviews the hyperbolic diffusion model and its properties. We discuss how the Euler and Milstein schemes can be used to discretize the model, and thus to provide approximations to the posterior likelihood. Some stylized facts about equity return series are summarized and related to the hyperbolic diffusion. Section 3 describes the MCMC method. In Section 4 we fit the model to three stock market indexes over a decade of daily data using the MCMC method based on both the Milstein and Euler schemes. Statistical inference is then made via the posterior quantities. In Section 5 we examine the statistical properties of sample paths generated

[1] A notable exception is the time-changed Levy process proposed recently by Carr and Wu (2003). These authors, however, did not provide any empirical analysis of their model. Eraker, Johnannes, and Polson (2003) proposed a diffusion model with jumps in both the return and the volatility, but they did not examine the statistical properties of their process.

[2] Rydberg (1999) considered the normal inverse Gaussian diffusion, which differs from the hyperbolic diffusion in the form of the stationary density.

by the hyperbolic diffusion. We find that many of the stylized facts for stock returns in the empirical finance literature documented by Ryden, Teräsverta and Asbrink (1998) are satisfied. Section 6 concludes.

2 Hyperbolic Diffusion and Some Stylized Facts of Stock Returns

Consider the following continuous-time parametric diffusion

$$dX_t = \mu(X_t, \theta)dt + \sigma(X_t, \theta)dW_t, \tag{1}$$

where X_t is a state variable, W_t is a standard Brownian motion defined on the probability space $(\Omega, \Im^B, (\Im_t^B)_{t\geq 0}, P)$, $\mu(\cdot,\cdot)$ and $\sigma(\cdot,\cdot)$ are known functions, and θ is a vector of unknown parameters.

Many empirical studies have shown that asset returns are not normally distributed. Barndorff-Nielsen (1978) suggested using the hyperbolic distribution to describe unconditional asset returns. The density of the hyperbolic distribution is proportional to $1/b^2(x)$, with

$$b(x) = \exp\left\{\frac{1}{2}\left[\alpha\sqrt{\delta^2 + (x-\mu)^2} - \beta(x-\mu)\right]\right\}, \tag{2}$$

where α, β, δ and μ are the parameters of the distribution satisfying $\alpha > |\beta| \geq 0$ and $\delta > 0$. It is noted that δ is the scale parameter, μ is the location parameter, β determines the symmetry (the distribution is symmetrical about μ if $\beta = 0$) and α determines the steepness of the distribution.

We assume that the stock price S_t depends on the state variable X_t as follows

$$S_t = \exp(X_t + \kappa t), \tag{3}$$

where κ is the (constant) drift rate. Following Bibby and Sorensen (1997) we consider the following hyperbolic diffusion process to describe the movement of stock prices[3]

$$dS_t = S_t\left\{\left[\kappa + \frac{1}{2}\sigma^2 b^2(\ln S_t - \kappa t)\right]dt + \sigma b(\ln S_t - \kappa t)dW_t\right\}. \tag{4}$$

[3]Note that μ in equation (2) and σ in equation (4) are parameters of the diffusion. They should not be confused with $\mu(\cdot,\cdot)$ and $\sigma(\cdot,\cdot)$, which are known functions of the drift and diffusion terms, respectively.

Bibby and Sorensen (1997) obtained some interesting statistical properties of the process S_t. For instance, they showed that the marginal distribution of $\ln S_t$ is hyperbolic and hence $\ln S_t$ is approximately hyperbolically distributed after a sufficiently long time period. Also, the distribution of increments over short intervals has thick tails while an increment over a long interval follows a distribution that is close to being hyperbolic.

To derive the dynamic properties of stock returns, we apply Ito's lemma to obtain

$$dX_t = \sigma\, b(X_t) dW_t, \tag{5}$$

which represents a diffusion process with no drift. As dW_t are uncorrelated over nonoverlapping intervals, increments of the log-prices (i.e., the continuously compounded rates of return) are serially uncorrelated. Similar to the stochastic volatility (SV) and autoregressive conditional heteroscedasticity (ARCH) models, the squared increments of the log-prices are generally serially correlated. In other words, return series of hyperbolic diffusion are likely to exhibit volatility clustering as demonstrated by SV and ARCH models, which have been found to be successful in describing many stylized facts of equity return series.

To understand why a hyperbolic diffusion generates volatility clustering and long memory properties, we apply the Euler approximation to the diffusion model for the log-price (i.e., equation (1)) and obtain

$$Y_t \approx \sigma \,\exp\left[\frac{1}{2}\left\{\alpha\sqrt{\delta^2 + (\sum_{i=1}^{\infty} Y_{t-i} - \mu)^2} - \beta(\sum_{i=1}^{\infty} Y_{t-i} - \mu)\right\}\right] e_t, \tag{6}$$

where $Y_t = \ln S_{t+\Delta t} - \ln S_t$ denotes the return and $e_t \sim$ iid $N(0, \Delta t)$. Equivalently this equation can be re-written as

$$\begin{aligned} Y_t &\approx \sigma \,\exp\left\{\frac{1}{2}h_t\right\} e_t \\ h_t &= \alpha\sqrt{\delta^2 + (\sum_{i=1}^{\infty} Y_{t-i} - \mu)^2} - \beta(\sum_{i=1}^{\infty} Y_{t-i} - \mu). \end{aligned}$$

Comparing the above specification with the well-known ARCH(∞) model (Engle (1982)),

$$\begin{aligned} Y_t &= \sigma_t e_t \\ \sigma_t^2 &= \alpha_0 + \sum_{i=1}^{\infty} \alpha_i Y_{t-i}^2, \end{aligned}$$

it can be seen that the hyperbolic diffusion model can be regarded as a special case of the following nonlinear ARCH(∞) model:

$$\begin{aligned} Y_t &= \sigma \exp\left\{\frac{1}{2}h_t\right\} e_t \\ h_t &= f(Y_{t-1}, Y_{t-2}, \cdots). \end{aligned}$$

This suggests that the hyperbolic diffusion model may generate return series that exhibit ARCH effects. Furthermore, the nonlinear relationship between h_t and Y_{t-i} may cause long-memory properties in the absolute return as well as the squared return.

Before we discuss the MCMC estimation methods for the hyperbolic diffusion we summarize here some stylized facts for equity return series in the empirical finance literature, which may be used as benchmarks for empirical equity price processes. Let r_t denote the return of a stock. Ryden, Teräsverta and Asbrink (1998) summarized the following dynamic properties of r_t found in many empirical studies:

1. r_t are not autocorrelated, except possibly at lag one.

2. The autocorrelation functions of $|r_t|$ and r_t^2 decay slowly. The decay is much slower than the exponential rate of the autocorrelation function of a stationary ARMA process.

3. $\text{corr}(|r_t|, |r_{t-k}|) > \text{corr}(|r_t|^\phi, |r_{t-k}|^\phi)$, $\phi \neq 1$. The autocorrelations of powers of absolute return are highest at power one. This is called the Taylor effect.

In additional to the above dynamic properties, the following are two well known static properties:

1. Returns often show strong evidence that the marginal distribution has thick tails.

2. Returns often show weak evidence that the marginal distribution is skewed.

3 Discretization of Hyperbolic Diffusion

Although the unconditional distribution of the hyperbolic diffusion process is hyperbolic, the transition density is unknown. Therefore, the exact ML method is difficult to implement. Bibby and Sorensen (1997) suggested using the martingale estimating function approach of Bibby and Sorensen (1995) to estimate the diffusion models. This approach, however, requires knowledge of the conditional expectation and conditional variance of the underlying diffusion which are known only for very simple models, such as those with a linear drift in the state variable. Hence, although the martingale estimating function method provides estimates that are consistent and asymptotically normal, implementation of the method is difficult in practice.

We propose to use the MCMC method to estimate the hyperbolic diffusion after discretizing the model. In the next section we shall discuss the MCMC method. In this section, we outline the Euler and Milstein schemes for the discretization of the hyperbolic diffusion model. The discretized schemes provide approximations to the likelihood function, which is in turn used to approximate the posterior to facilitate Bayesian analysis. It is well known that the Milstein scheme provides an approximation with improved accuracy over the Euler scheme on approximating the underlying diffusion (Milstein (1978), Kloeden and Platen (1992)). As a consequence, it is expected that the likelihood and posterior calculated from the Milstein scheme would provide better approximations to the true counterparts than those from the Euler scheme. Indeed, Elerian (1998) compared the performance of the ML method based on these two schemes in the context of a univariate CIR model (Cox, Ingersoll, and Ross (1985)) and found that the Milstein scheme offers improvements over the Euler scheme.

The Euler scheme approximates a general diffusion process such as equation (1) by the following expansion

$$X_{t+\Delta t} = X_t + \mu(X_t, \theta)\Delta t + \sigma(X_t, \theta)\Delta W_t,$$

where $\Delta W_t = \varepsilon_t \sqrt{\Delta t}$ with $\varepsilon_t \sim$ iid $N(0, 1)$. Assuming constant priors for all the parameters, and given $n + 1$ observations of $\mathbf{x} = \{x_t : t = 0, 1, ..., n\}$, the logarithmic likelihood of θ upon dropping the constant term, is

$$\log p_E(\theta|\mathbf{x}) = -\frac{1}{2}\sum_{t=1}^{n}\log(\sigma(x_t,\theta)^2\Delta t) - \frac{1}{2}\sum_{t=1}^{n}\frac{(x_t - x_{t-1} - \mu(x_t,\theta)\Delta t)^2}{\sigma(x_t,\theta)^2\Delta t}, \tag{7}$$

which can be used directly as an approximate logarithmic likelihood for ML estimation or as an approximate logarithmic posterior for the MCMC algorithm. Hereafter, we refer to $p_E(\theta|\mathbf{x})$ defined in (7) as the Euler likelihood.

Taking a higher-order term in the Taylor expansion, the Milstein approach approximates a general diffusion process by the following equation

$$X_{t+\Delta t} = X_t + \mu(X_t,\theta)\Delta t + \sigma(X_t,\theta)\Delta W_t + \frac{1}{2}\sigma(X_t,\theta)\frac{\partial\sigma(X_t,\theta)}{\partial X_t}\left[(\Delta W_t)^2 - \Delta t\right], \tag{8}$$

which can be rewritten as

$$X_{t+\Delta t} - X_t - \mu(X_t,\theta)\Delta t + g(X_t,\theta)\Delta t = \sigma(X_t,\theta)\sqrt{\Delta t}\,\varepsilon + g(X_t,\theta)\,\Delta t\,\varepsilon^2, \tag{9}$$

where $g(X_t,\theta) = \frac{1}{2}\sigma(X_t,\theta)(\partial\sigma(X_t,\theta)/\partial X_t)$. Let

$$a = \sigma(X_t,\theta)\sqrt{\Delta t}, \qquad b = g(X_t,\theta)\Delta t, \tag{10}$$

then equation (8) can be represented by

$$Y = a\varepsilon + b\varepsilon^2 = b\left[\left(\varepsilon + \frac{a}{2b}\right)^2 - \frac{a^2}{4b^2}\right], \tag{11}$$

where $Y = X_{t+\Delta t} - X_t - \mu(X_t,\theta)\Delta t + g(X_t,\theta)\Delta t$.

The normality assumption implies that $\left(\varepsilon + \frac{a}{2b}\right)^2$ (denoted by, say, Z) follows a noncentral χ^2 distribution with 1 degree of freedom and noncentrality parameter $\lambda = a^2/(4b^2)$. Elerian (1998) showed that the density of Z is given by

$$f(z) = \frac{1}{2}\exp\left\{-\frac{\lambda+z}{2}\right\}\left(\frac{z}{\lambda}\right)^{-1/4} I_{-1/2}\left(\sqrt{\lambda z}\right), \tag{12}$$

where

$$I_{-1/2}(w) = \sqrt{\frac{2}{w}}\sum_{j=0}^{\infty}\frac{(w/2)^{2j}}{j!\Gamma(j+1/2)} = \sqrt{\frac{2}{\pi w}}\cosh(w),$$

with $\cosh(w) = (1/2)\{\exp(w) + \exp(-w)\}$ being the hyperbolic cosine function. Hence the density of Y is

$$f^*(y) = \frac{1}{b}f\left(\frac{y}{b} + \frac{a^2}{4b^2}\right). \tag{13}$$

and the logarithmic likelihood upon dropping the constant, is given by

$$\log p_M(\theta|\mathbf{x}) = \sum_{t=1}^{n} \left[\log \left\{ f \left(\frac{y_t}{b} + \frac{a^2}{4b^2} \right) \right\} - \log(b) \right], \tag{14}$$

where $y_t = x_t - x_{t-1} - \mu(x_{t-1}, \theta)\Delta t + g(x_{t-1}, \theta)\Delta t$. Again, the above quantity can be used either directly as an improved approximate logarithmic likelihood for ML estimation or as an improved approximate logarithmic posterior for the purpose of MCMC simulation. Hereafter, we refer to $p_M(\theta|\mathbf{x})$ defined in (14) as the Milstein likelihood.

4 Estimating Hyperbolic Diffusion via MCMC

4.1 MCMC

The MCMC strategy has proved useful in many statistical applications, and has many advantages compared to traditional independent sampling methods. Geweke (1999) provided a survey of the fundamental principles of subjective Bayesian inference in econometrics and the implementation of these principles using posterior simulation methods, emphasizing the importance of simulation methods and describing the implementation of MCMC simulation for Bayesian inference. Gilks, Richardson and Spiegelhatler (1996) presented a collection of papers on the application of MCMC algorithms. In econometrics and finance many successful applications of the MCMC method can be found (e.g., Eraker (2001), and Elerian, Chib and Shephard (2002)). We refer readers to Chib (2001) and Johannes and Polson (2003) for recent surveys on the applications of MCMC in econometrics and finance.

Bayesian inference concerning a parameter vector θ conditional on data $\mathbf{x}$ is made via the posterior density $p(\theta|\mathbf{x})$. By the Bayes theorem, the posterior takes the form

$$\pi(\theta|\mathbf{x}) = c\, p(\mathbf{x}|\theta)\, \pi(\theta), \tag{15}$$

where c is a normalizing constant, $p(\mathbf{x}|\theta)$ is the likelihood of $\mathbf{x}$ conditional upon θ and $\pi(\theta)$ is the prior density of θ. The Bayesian approach requires that statistical inference be based on the posterior. Dealing with the posterior, however, is often analytically intractable. Nonetheless, if we can sample the parameter vector from the posterior, statistical inference about the parameter vector can be made using the usual Monte Carlo approach. The MCMC method aims to provide a general mechanism to sample the parameter vector from its posterior density. While simulating directly from the posterior distribution is typically very difficult, the MCMC method sets up a Markov chain so that its stationary distribution is the same as the posterior density. When the Markov chain converges, the simulated values may be regarded as a sample obtained from the posterior.

There are two broad categories of algorithms for implementing MCMC, which are, respectively, the Gibbs sampler and the Metropolis-Hastings algorithm. Let the current state be denoted as $\theta = (\theta_1, \theta_2, \cdots, \theta_p)$, and assume that the full conditional densities of θ_i's are available. The Gibbs sampler generates the next state θ', in which each component is generated from a sequence of conditional densities. The Metropolis-Hastings algorithm generates a candidate θ' from a proposal density denoted by $q(.|\theta)$. The proposal density should satisfy certain properties, such as the reversibility condition discussed in Chib and Greenberg (1995) and Gilks, Richardson and Spiegelhatler (1996), among many others. The candidate is then accepted with probability $T(\theta, \theta')$, which is defined by

$$T(\theta, \theta') = \min \left\{ 1, \frac{\pi(\theta'|\mathbf{x})q(\theta|\theta')}{\pi(\theta|\mathbf{x})q(\theta'|\theta)} \right\}. \tag{16}$$

If the candidate is accepted, the next state is set to θ'. Otherwise, the chain does not move. Robert and Casella (1999, Chapter 7) showed that the Gibbs sampler is equivalent to a composition of p Metropolis-Hastings algorithms with acceptance probabilities uniformly equal to 1. Robert and Casella (1999) presented detailed discussions on the use of the Metropolis-Hastings algorithm and the Gibbs sampler.

As the full conditional density is often difficult to derive, the Metropolis-Hastings algorithm is generally adopted in complex problems. In this paper we use the Metropolis-Hastings algorithm for its simplicity. In what follows we briefly describe the procedure of the algorithm.

Step 1: Given the current state $\theta^{(i)}$, generate a candidate θ' from the proposal density

$q(\cdot|\theta^{(i)})$.

Step 2: Calculate the acceptance probability $T(\theta^{(i)}, \theta')$ according to (16).

Step 3: Accept the proposal with probability $T(\theta^{(i)}, \theta')$ and set $\theta^{(i+1)} = \theta'$. Otherwise, reject the candidate and set $\theta^{(i+1)} = \theta^{(i)}$.

Step 4: Repeat the previous steps to obtain a chain $\{\theta^{(0)}, \theta^{(1)}, \theta^{(2)}, \cdots\}$, where $\theta^{(0)}$ denotes the initial state of θ. Discard the burn-in values (up to $\theta^{(d)}$, say) obtained whilst the chain converges in distribution to the joint posterior. Then the remaining values, $\{\theta^{(d+1)}, \theta^{(d+2)}, \cdots\}$ are a correlated chain simulated from $\pi(\theta|\mathbf{x})$, and have the same stationary transition density as $\pi(\theta|\mathbf{x})$.

Two important points should be noted. First, the calculation of $T(\theta^{(i)}, \theta')$ does not require knowledge of the normalizing constant in the posterior function. Second, if the proposal density is symmetric, that is, $q(x|y) = q(y|x)$, then the acceptance probability reduces to $\pi(\theta'|\mathbf{x})/\pi(\theta^{(i)}|\mathbf{x})$. Moreover, if $q(y|x)$ is a function of $|y - x|$, the resulting algorithm is called random-walk Metropolis-Hastings algorithm, which has been widely used in practice due to its simplicity.

4.2 Empirical Results

In this section we apply the random-walk Metropolis-Hastings algorithm to the discretized diffusion processes and present empirical results based on some real data sets. The data series considered are the MSCI World Index, the MSCI Europe Index and the NYSE Index. The series consist of weekly observations from January 1, 1990 to December 31, 2000.

As argued before, the transition density of hyperbolic diffusions does not have closed-form, making the direct ML approach difficult. However, the transition density of discretized models under both schemes has an analytic expression, which, in theory, can be used to obtain the approximate ML estimates. Before we carried out the Bayesian MCMC analysis, we implemented the approximate ML estimation but found that numerical optimizations rarely converged. This experience indicates that the likelihood function of the discretized models is not well behaved. As a conditional simulation method, MCMC avoids any numerical difficulties associated with numerical optimizations. We now describe the details of the implementation of the MCMC method for the estimation of the hyperbolic diffusion.

4.2.1 Empirical Results under the Milstein Likelihood

Assume that the priors of the parameters are given by: $\kappa \sim N(0, 10)$, $\alpha \sim \Gamma(1, 20)$, $\delta^2 \sim \Gamma(0.05, 20)$, $\mu \sim N(5, 10)$, $\beta \sim U(-\alpha, \alpha)$ and $\sigma^2 \sim IG(5, 0.05)$, where U, Γ, IG refer to the uniform, gamma, and inverted gamma densities, respectively. These priors are very flat and nearly noninformative. The joint prior of all the parameters, denoted as $\pi(\theta)$, is the product of these marginal priors. Based on the Milstein likelihood $p_M(\theta|\mathbf{x})$, we can obtain the joint posterior

$$\pi(\theta|x) \propto \pi(\theta) p_M(\theta|x).$$

In the implementation of the random-walk Metropolis-Hastings algorithm, the proposal density is uniform on $[-0.5, 0.5]$, and the parameter vector θ is updated as follows:

$$\theta' = \theta + \tau\varepsilon,$$

where θ' is the proposal for θ, ε is a vector of random numbers drawn from the uniform density on $[-0.5, 0.5]$, and τ is a tuning parameter which is chosen so that the acceptance rate is between 20% and 30%. In addition, τ may be either a scalar- or vector-constant. Generally speaking, if the parameters are of weak correlation and their values are of the same scale, τ can be a scalar constant. Otherwise, τ should be a constant vector, so that each parameter is assigned a specific tuning parameter.

4.2.2 Convergence Checking

In the implementation of the MCMC algorithm, the sampled path, denoted by $\{\theta^{(i)} : i = 1, 2, \cdots, N\}$, forms a Markov chain whose stationary density is the posterior $\pi(\theta|\mathbf{x})$, and the output is summarized in terms of the ergodic averages in the form of:

$$\overline{f}_N = \frac{1}{N} \sum_{i=1}^{N} f(\theta^{[i]}), \tag{17}$$

where $f(\cdot)$ is a real-valued function to be estimated. Roberts (1996) pointed out that most of the Markov chains produced in MCMC converge geometrically to the stationary distribution $\pi(\theta|\mathbf{x})$, and one of the most important consequences of the geometric convergence is that the central limit theorem of ergodic averages is invoked, i.e.,

$$\sqrt{N}\left(\overline{f}_N - E_\pi[f(\theta)]\right) \xrightarrow{D} N(0, \sigma_f^2), \tag{18}$$

where $E_\pi[\cdot]$ denotes the expectation operator under $\pi(\theta|\mathbf{x})$, and the convergence is in distribution. To assess the accuracy of the ergodic average as an estimate of $E_\pi[f(\theta)]$, it is essential to estimate σ_f^2. One of the most commonly used methods for estimating σ_f^2 is the batch mean, which is discussed extensively in Roberts (1996).

To estimate σ_f^2 using the batch mean, the MCMC algorithm is run for $N = m \times n$ iterations, where n is sufficiently large so that

$$y_k = \frac{1}{n} \sum_{i=(k-1)n+1}^{kn} f(\theta^{[i]}), \tag{19}$$

for $k = 1, 2, \cdots, m$, are approximately independently distributed as $N(E_\pi[f(\theta)], \sigma_f^2/n)$. Therefore σ_f^2 can be estimated by

$$\hat{\sigma}_f^2 = \frac{n}{m-1} \sum_{k=1}^{m} (y_k - \overline{f}_N)^2, \tag{20}$$

where $\overline{f}_N$ is defined in equation (17). Thus, the standard error of $\overline{f}_N$ can be estimated by $\sqrt{\hat{\sigma}_f^2/N}$, which is called the batch-mean standard error or the Monte Carlo standard error, and is commonly used for checking the mixing performance.

In addition to the batch-mean standard error, one may also compute the standard deviation $\tilde{\sigma}_f$ directly based on the sampled paths using the formula

$$\tilde{\sigma}_f = \left\{ \frac{1}{N-1} \sum_{i=1}^{N} \left[f(\theta^{[i]}) - \overline{f}_N \right]^2 \right\}^{1/2}. \tag{21}$$

Kim, Shephard and Chib (1998) indicated that the mixing performance of the sampled paths can be measured using the simulation inefficiency factor (SIF), also called the integrated autocorrelation time by Sokal (1996), which is estimated as the variance of the sample mean divided by the variance of the sample mean from a hypothetical sampler that draws independent random observations from the posterior distribution. Meyer and Yu (2000) showed that SIF is given by

$$\text{SIF} = \frac{\hat{\sigma}_f^2}{\tilde{\sigma}_f^2}. \tag{22}$$

In the empirical applications, the burn-in period is taken as 10,000 iterations and the number of total recorded iterations after the burn-in period is 50,000. Based on the sampled path for each data set, we calculate the ergodic average (or mean) and standard deviations. The MC standard errors are obtained using the batch-mean approach described in equations (17) through (19) with $f(x) = x$. The number of batches is $m = 50$, and there are n =1,000 draws in each batch.

Table 1. MCMC Results of the Milstein Scheme under Specific Priors

Data	Para.	Mean	CI	SD	MCSE	SIF	AC
MSCI World	κ	-0.08130	(-0.12732, -0.03913)	0.02225	0.00091	82.85	0.20
	α	1.49535	(1.25644, 1.74913)	0.12335	0.00356	41.70	0.24
	δ^2	0.02058	(0.00034, 0.09382)	0.02625	0.00075	40.37	0.27
	μ	6.48077	(6.30844, 6.66141)	0.09040	0.00358	78.44	0.21
	β	0.37230	(0.08179, 0.58669)	0.12743	0.00314	30.28	0.22
	σ^2	0.00765	(0.00579, 0.00958)	0.00095	0.00004	67.41	0.26
MSCI Europe	κ	-0.01871	(-0.06965, 0.02870)	0.02526	0.00094	69.56	0.25
	α	1.56359	(1.11674, 1.81534)	0.17638	0.00505	40.97	0.25
	δ^2	0.04023	(0.00057, 0.19857)	0.05389	0.00239	98.58	0.24
	μ	6.31699	(6.10846, 6.54861)	0.11330	0.00374	54.62	0.24
	β	0.27217	(-0.21448, 0.65201)	0.22147	0.00674	46.37	0.25
	σ^2	0.01043	(0.00704, 0.01347)	0.00159	0.00006	71.74	0.23
NYSE	κ	-0.03571	(-0.08039, 0.00643)	0.02218	0.00079	63.95	0.23
	α	1.65771	(1.29826, 2.83869)	0.13635	0.00308	25.46	0.23
	δ^2	0.01400	(0.00018, 0.06829)	0.02008	0.00054	36.65	0.27
	μ	5.73335	(5.54733, 5.92901)	0.09667	0.00359	68.83	0.24
	β	0.27616	(-0.03691, 0.51257)	0.13933	0.00326	27.38	0.22
	σ^2	0.00805	(0.00616, 0.01016)	0.00102	0.00004	64.42	0.24

Note: CI refers to the 95% confidence interval. SD refers to the standard deviation computed through (20). MCSE refers to the Monte Carlo standard error computed through the batch-mean approach. SIF refers to the simulation inefficiency factor given by (22). AC refers to the acceptance probability.

We plot the MCMC sample paths of the parameters the World Index in Figures 1.1 to 1.6, and the autocorrelation functions (ACFs) of these sample paths in Figures 2.1 to 2.6.[4] These plots show that the sample paths are reasonably well mixed. Table 1 summarizes the ergodic averages, standard deviations, 95% Bayes confidence intervals, Monte Carlo standard errors, and the simulation inefficiency factors for each data set. The Bayes confidence interval can be used to test the significance of each parameter. For example, for the World Index all parameter estimates are significantly different from zero. Note that the posterior means of the steepness parameter α are quite similar across the three indexes. While the Europe Index and the NYSE Index are symmetrical (the sampled posterior β is not significantly different from zero), the World index is asymmetric. For the scale and volatility parameters (i.e., δ and σ), the World and NYSE indexes (but not the Europe Index) have similar posterior means.

4.2.3 Robustness to the Choice of the Priors

To examine the robustness of the results with respect to the choice of the priors, we alter the priors in two ways. First, we keep the prior distributions in the same family as before but change some hyperparameters. The results are very similar. Second, we use a different set of prior distributions, which are now the uniform density. As any constant in the posterior will be eliminated from both the nominator and denominator when computing the acceptance probability, we effectively assume that the support of each uniform prior is wide enough for any update of the associated parameter. Then we use the same MCMC procedure as before and summarize the results in Table 2.

A comparison with the results in Table 1 reveals the following conclusions. Firstly, SIFs are either comparable (in most cases) to or marginally higher (in a few cases) than those under the priors adopted in Section 4.2.1, suggesting that the mixing performance is not affected much by the change of priors. Secondly, there is no obvious difference in the ergodic averages and CIs under both sets of priors, suggesting that the posterior distribution is robust to the change of priors.

[4]To save space, the plots for the other two indexes are not presented.

Table 2. MCMC Results of the Milstein Scheme under Uniform Priors

Data	Para.	Mean	CI	SD	MCSE	SIF	AC
MSCI World	κ	-0.07950	(-0.12554, -0.03552)	0.02263	0.00083	67.41	0.29
	α	1.53933	(1.26136, 1.85925)	0.14764	0.00442	44.30	0.26
	δ^2	0.02636	(0.00037, 0.12875)	0.03635	0.00153	87.95	0.27
	μ	6.48224	(6.30377, 6.66300)	0.09125	0.00292	52.16	0.28
	β	0.35768	(0.06210, 0.58873)	0.13177	0.00382	41.24	0.25
	σ^2	0.00741	(0.00513, 0.00953)	0.00109	0.00004	62.65	0.26
MSCI Europe	κ	-0.01422	(-0.06831, 0.03407)	0.02553	0.00100	76.17	0.27
	α	1.54634	(0.99561, 2.12376)	0.27217	0.01247	102.84	0.29
	δ^2	0.05139	(0.00068, 0.24200)	0.06502	0.00350	146.10	0.29
	μ	6.29925	(6.06792, 6.54124)	0.12231	0.00391	51.06	0.27
	β	0.27821	(-0.23627, 0.73844)	0.24608	0.01044	89.79	0.27
	σ^2	0.01006	(0.00584, 0.01371)	0.00199	0.00010	134.66	0.26
NYSE	κ	-0.03117	(-0.07709, 0.01361)	0.02332	0.00094	81.94	0.25
	α	1.65771	(1.33921, 2.04863)	0.17835	0.00744	86.90	0.25
	δ^2	0.01999	(0.00024, 0.11207)	0.03067	0.00152	119.81	0.26
	μ	5.73078	(5.54249, 5.92304)	0.09845	0.00370	72.31	0.25
	β	0.23127	(-0.11584, 0.48954)	0.15436	0.00511	54.95	0.27
	σ^2	0.00766	(0.00522, 0.01003)	0.00118	0.00005	87.56	0.24

Note: See Table 1.

4.2.4 Empirical Results under the Euler Likelihood

To compare the MCMC performance based on the Milstein and Euler likelihoods, we apply the sampling algorithm to the posterior

$$\pi(\theta|\mathbf{x}) \propto \pi(\theta)p_E(\theta|\mathbf{x}), \tag{23}$$

where $\pi(\theta)$ is the same as that in Section 4.2.3, and $p_E(\theta|\mathbf{x})$ is the Euler likelihood defined in (7). Applying the sampling algorithm to all three data sets using the same priors as adopted in Section 4.2.1, we obtain the empirical results tabulated in Table 3. The sample paths of parameters and the ACFs of sample paths are plotted, respectively, in Figures 3.1-3.6 and Figures 4.1-4.6.

A few more results emerge from Table 3 and Figures 3-4. First, the mixing performance under the Euler scheme is worse than that under the Milstein scheme, as can be seen in the fact that the sampled paths under the Euler scheme have larger variances than those under the Milstein scheme. This relative performance is also obvious from the plots. We apply the Heidelberger and Welch convergence test (Heidelberger and Welch (1983)) to all the sampled paths, and find that the samples from the Milstein scheme under both sets of priors pass the test for all parameters, whereas the samples from the Euler scheme fail the test for δ^2. Second, the ergodic averages for some parameters are different from those obtained under the Milstein likelihood. For example, the estimated κ and β are, respectively, significantly different from the corresponding estimates reported in Table 1. Better mixing and convergence under the Milstein scheme leads to the conclusion that the empirical results based on the Milstein likelihood are more reliable.

Table 3. MCMC Results of the Euler Scheme under Specific Priors

Data	Para.	Mean	CI	SD	MCSE	SIF	AC
MSCI World	κ	0.18743	(0.15499, 0.24044)	0.02214	0.00073	54.28	0.21
	α	1.16844	(0.80189, 1.60388)	0.20522	0.00669	53.16	0.27
	δ^2	0.02092	(0.00020, 0.12777)	0.03768	0.00173	105.95	0.22
	μ	5.30367	(5.01787, 5.53962)	0.13311	0.00765	165.21	0.24
	β	-0.60743	(-1.27971, 0.19868)	0.38850	0.01894	118.83	0.32
	σ^2	0.00177	(0.00150, 0.00204)	0.00013	0.00001	45.80	0.22
MSCI Europe	κ	0.22699	(0.19302, 0.27736)	0.02113	0.00113	141.94	0.21
	α	1.13359	(0.73979, 1.60613)	0.22233	0.01026	106.52	0.27
	δ^2	0.01659	(0.00015, 0.10510)	0.03186	0.00207	210.23	0.26
	μ	5.00702	(4.39287, 5.48137)	0.25509	0.01898	276.91	0.22
	β	-0.03255	(-1.08518, 0.75978)	0.44467	0.03283	272.52	0.24
	σ^2	0.00159	(0.00134, 0.00185)	0.00013	0.00001	92.74	0.22
NYSE	κ	0.25975	(0.21057, 0.32013)	0.02803	0.00086	46.62	0.29
	α	1.28902	(0.91002, 1.72375)	0.20830	0.00935	100.73	0.29
	δ^2	0.05159	(0.00084, 0.21860)	0.06303	0.00302	114.97	0.20
	μ	4.58943	(4.33953, 4.76834)	0.11107	0.00499	100.91	0.28
	β	-1.11554	(-1.59922, -0.53160)	0.26377	0.01220	107.02	0.29
	σ^2	0.00283	(0.00227, 0.00330)	0.00027	0.00001	77.00	0.23

Note: See Table 1.

5 Empirical Properties of Hyperbolic Diffusions

Rydberg (1999) reported simulation results of the normal inverse Gaussian diffusion in which the autocorrelation function of r_t^2 declines very slowly, thus satisfying partly dynamic property 2 listed in Section 2. In this section we examine in more detail whether the hyperbolic diffusion would give rise to the statistical properties described in Section 2.

To examine the posterior properties of the hyperbolic diffusion, we record 1,000 sampled parameter vectors (i.e., 1 draw for every 50 draws of the MCMC iterations). Using the sampled parameter vector at each recorded draw, we generate a path of daily-price series with 2,000 observations based on Milstein approximation, where a time interval of 15 minutes is used (i.e., we use $\Delta t = 1/7{,}000$ year, assuming 7 hours of trading per day and 250 trading days per year). For each sampled path, we calculate the sample kurtosis, sample skewness, Box-Pierce statistic of the squared returns with 20 lags included, and ACF up to lag order 300 for $|r_t|^\phi$ with $\phi = 1, 1.5, 2$. Table 4 reports the means of the kurtosis, skewness, Box-Pierce statistic and the ACF of these 1,000 sample paths. The last row of the table reports the proportions that the kurtosis is larger than 3 (the kurtosis implied by the normal distribution), the skewness is less than 0, the Box-Pierce statistic is larger than 31.41 (the critical value at the 5% level), and the ACF at $\phi = 1$ is the largest among $\phi = 1, 1.5, 2$ at lags of order 100, 200 and 300 among the 1,000 sample paths.

Several points can be observed from the table. First, there is overwhelming evidence of excess kurtosis (thick tails) in the unconditional distribution of the hyperbolic diffusion, as manifest in the mean of the kurtosis and the proportion of excess kurtosis across the simulated sample paths. This is consistent with static property 1. Second, there is weak evidence of asymmetry in the unconditional distribution: 44.3% of the time we get negative skewness whereas 55.7% of the time we have positive skewness. This is consistent with static property 2. Third, the average value of the Box-Pierce statistic is much larger than the 5% critical value and the proportion of this statistic being significant is

67%, indicating reasonably strong evidence of an ARCH effect. Moreover, like Rydberg (1999), we also find that the ACF of $|r_t|$ and r_t^2 decay very slowly, a pattern inconsistent with the exponential decay. All these results are consistent with the dynamic property 2. Finally, at all three lags considered, most of the time the ACF is highest at power one. This result is consistent with the Taylor effect.

Table 4. Analysis of Statistical Properties of the Hyperbolic Diffusion

	Kurt	Skew	B-P	ACF of $\{|r_t| \ / \ |r_t|^{1.5} \ / \ |r_t|^2\}$		
				Lag 100	Lag 200	Lag 300
Mean	20.52	.098	72.96	.134/.107/.081	.091/.072/.055	.062/.050/.039
Prop	1.00	.443	.668	.736	.732	.624

6 Conclusions

In this paper we propose a Bayesian MCMC method to estimate the hyperbolic diffusion based on the discretized density via the Milstein scheme. Relative to some alternative estimation methods, such as the ML estimation based on the discretized densities and the MCMC method based on the discretized density via the Euler scheme, we find that the MCMC method using the Milstein scheme provides the best empirical results. Apart from showing evidence that the MCMC method is a useful tool for estimating hyperbolic diffusions and making statistical inferences, we have also demonstrated that the hyperbolic diffusion is able to exhibit many of the stylized facts about asset returns documented in the financial econometrics literature.

References

1. Barndorff-Nielsen, O.E. (1978), Hyperbolic Distributions and Distributions on Hyperbolae, Scandinavian Journal of Statistics, 5, 151-157.

2. Bibby, B.M., and M. Sorensen (1995), Martingale Estimation Functions for Discretely Observed Diffusion Processes, Bernoulli, 1, 17-39.

3. Bibby, B.M., and M. Sorensen (1997), A Hyperbolic Diffusion Model for Stock Prices, Finance and Stochastics, 1, 25-41.

4. Carr, P., and L. Wu (2003), Time-Changed Levy Processes and Option Pricing, Journal of Financial Economics, forthcoming.

5. Chib, S. (2001), Markov chain Monte Carlo Methods: Computation and Inference, in Handbook of Econometrics, eds by J.J. Heckman, and E. Leamer, North-Holland, Amsterdam, 5, 3569-3649.

6. Chibs, S., and E. Greenberg (1995), Understanding the Metropolis-Hastings Algorithm, American Statistician, 49, 327-335.

7. Cox, J., J. Ingersoll, and S. Ross (1985), A Theory of the Term Structure of Interest Rates, Econometrica, 53, 385–407.

8. Elerian, O. (1998), A Note on the Existence of a Closed-form Conditional Transition Density for the Milstein Scheme, Economics discussion paper 1998-W18, Nuffield College, Oxford.

9. Elerian, O., S. Chib, and N. Shephard (2001), Likelihood Inference for Discretely Observed Non-Linear Diffusions, Econometrica, 69, 959-993.

10. Engle, R.F. (1982), Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation, Econometrica, 50, 987-1007.

11. Eraker, B. (2001), MCMC Analysis of Diffusion Models with Application to Finance, Journal of Business and Economic Statistics, 19, 177-191.

12. Eraker, B., M. Johannes, and N. Polson (2003), The Impact of Jumps in Volatility and Returns, Journal of Finance, 58, 1269-1300.

13. Geweke, J. (1999), Using Simulation Methods for Bayesian Econometric Models: Inference, Development, and Communication, Econometric Reviews , 18, 1-73.

14. Gilks, W.R., S. Richardson, and D.J. Spiegelhatler (1996), Introducing Makov Chain Monte Carlo, in Markov chain Monte Carlo in Practice (eds. by W.R. Gilks, S. Richardson and D.J. Spiegelhalter), 45-57, London: Chapman & Hall.

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

16. Heston, S.L. (1993), A Closed-Form Solution for Options with Stochastic Volatility, with Application to Bond and Currency Options. Review of Financial Studies 6, 327–343.

17. Johannes, M., and N. Polson (2003), MCMC Method for Financial Econometrics. in Handbook of Financial Econometrics, eds. by Y. Aït-Sahalia and L. Hansen, forthcoming

18. Kim, S., N. Shephard, and S. Chib (1998), Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models, Review of Economic Studies, 65, 361-393.

19. Kloedan, P. E., and E. Platten (1995), Numerical Solution of Stochastic Differential Equations, Springer: New York

20. Kou, S.G. (2002), A Jump-Diffusion Model for Option Pricing, Management Science, 48, 1086-1101.

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

22. Milstein, G.N. (1978), A Method of Second-Order Accuracy Integration of Stochastic Differential Equations, Theory of Probability and its Applications, 23, 396-401.

23. Robert, C.P., and G. Casella (1999), Monte Carlo Statistical Methods, Springer: New York

24. Roberts, G.O. (1996), Markov chain Concepts Related to Sampling Algorithms, in Markov chain Monte Carlo in Practice, eds. by W.R. Gilks, S. Richardson and D.J. Spiegelhalter, 45-57, London: Chapman & Hall.

25. Rydberg, T.H. (1999), Generalized Hyperbolic Diffusion Processes with Applications in Finance, Mathematical Finance, 9, 183-201.

26. Ryden, T., T. Teräsverta, and S. Asbrink (1998), Stylized Facts of Daily Return Series and the Hidden Markov Model, Journal of Applied Econometrics, 13, 217-244.

27. Sokal, A.D. (1996), Monte Carlo Methods in Statistical Mechanics: Foundations and New Algorithms, Lectures at the Cargese Summer School on "Functional Integration: Basics and Applications".

Figure 1. MCMC Results for the Milstein Scheme

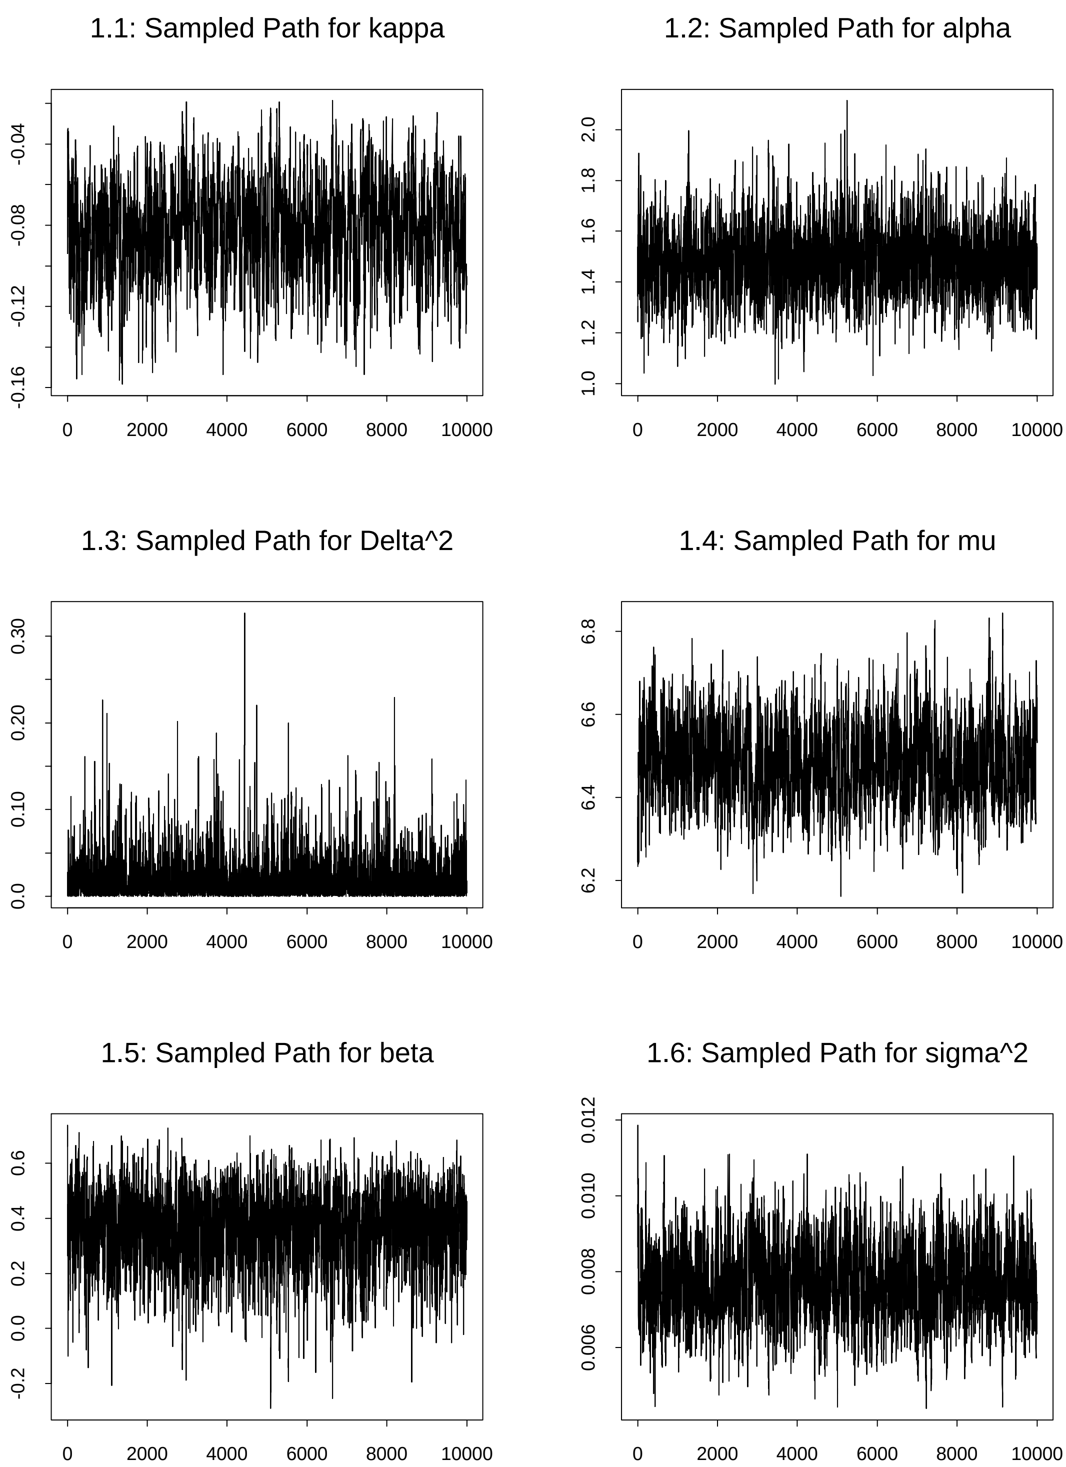

Figure 2. ACF for the Milstein Scheme

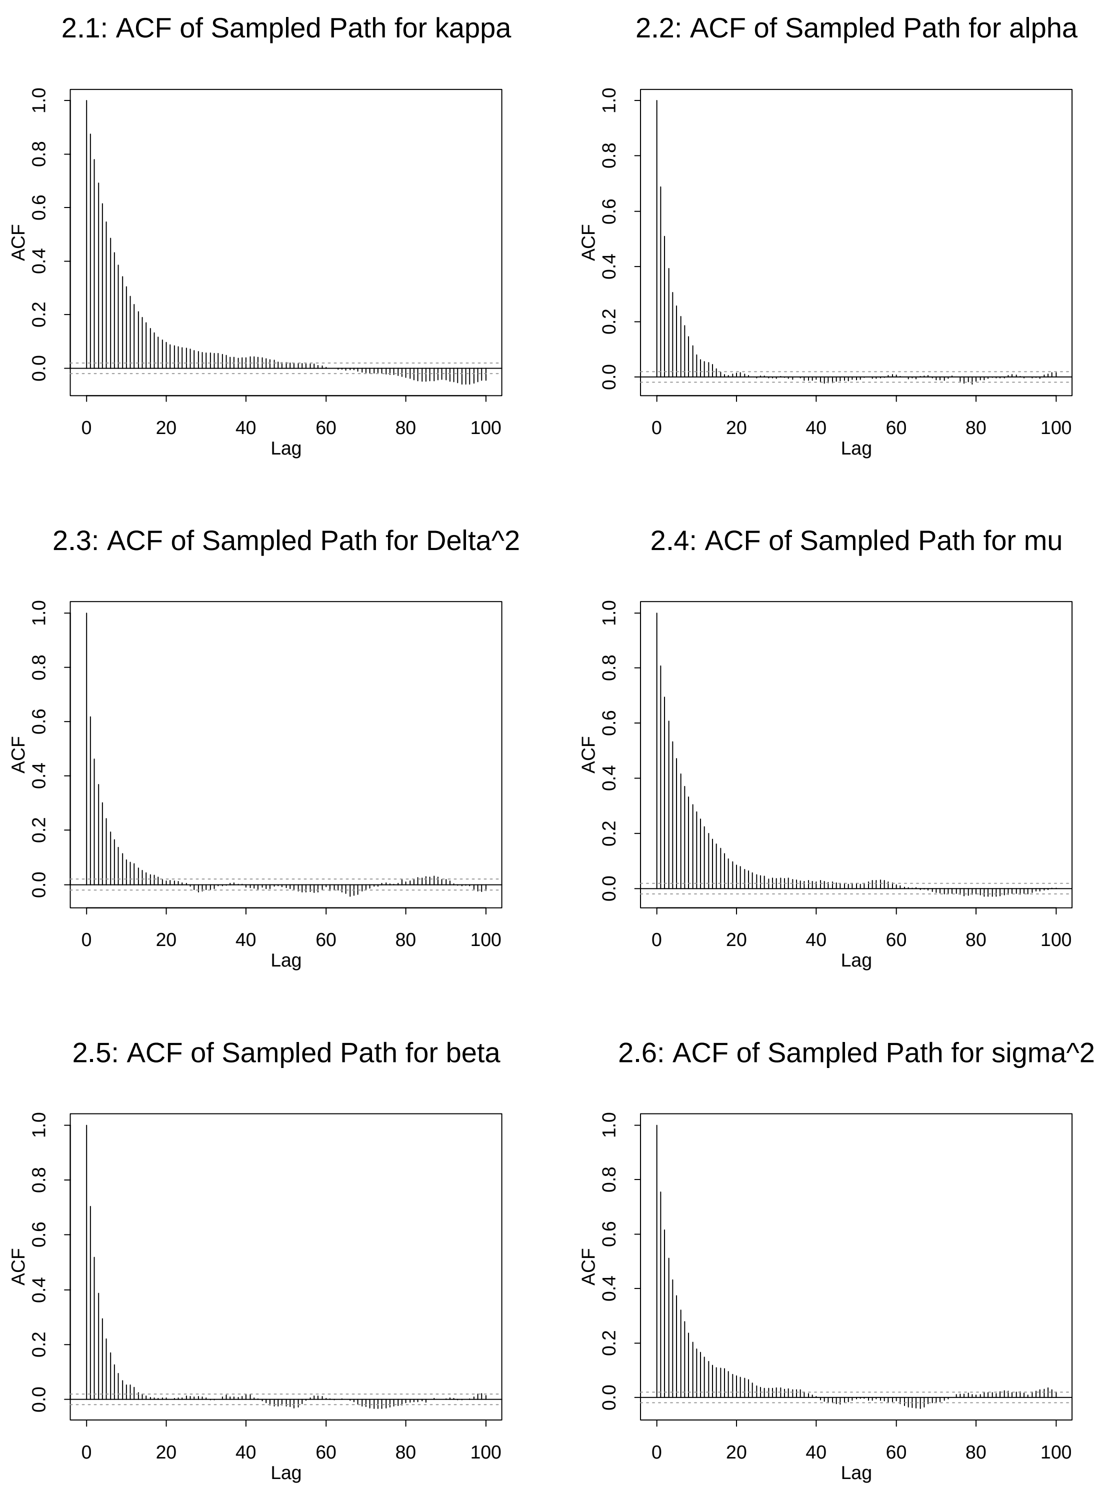

Figure 3. MCMC Results for the Euler Scheme

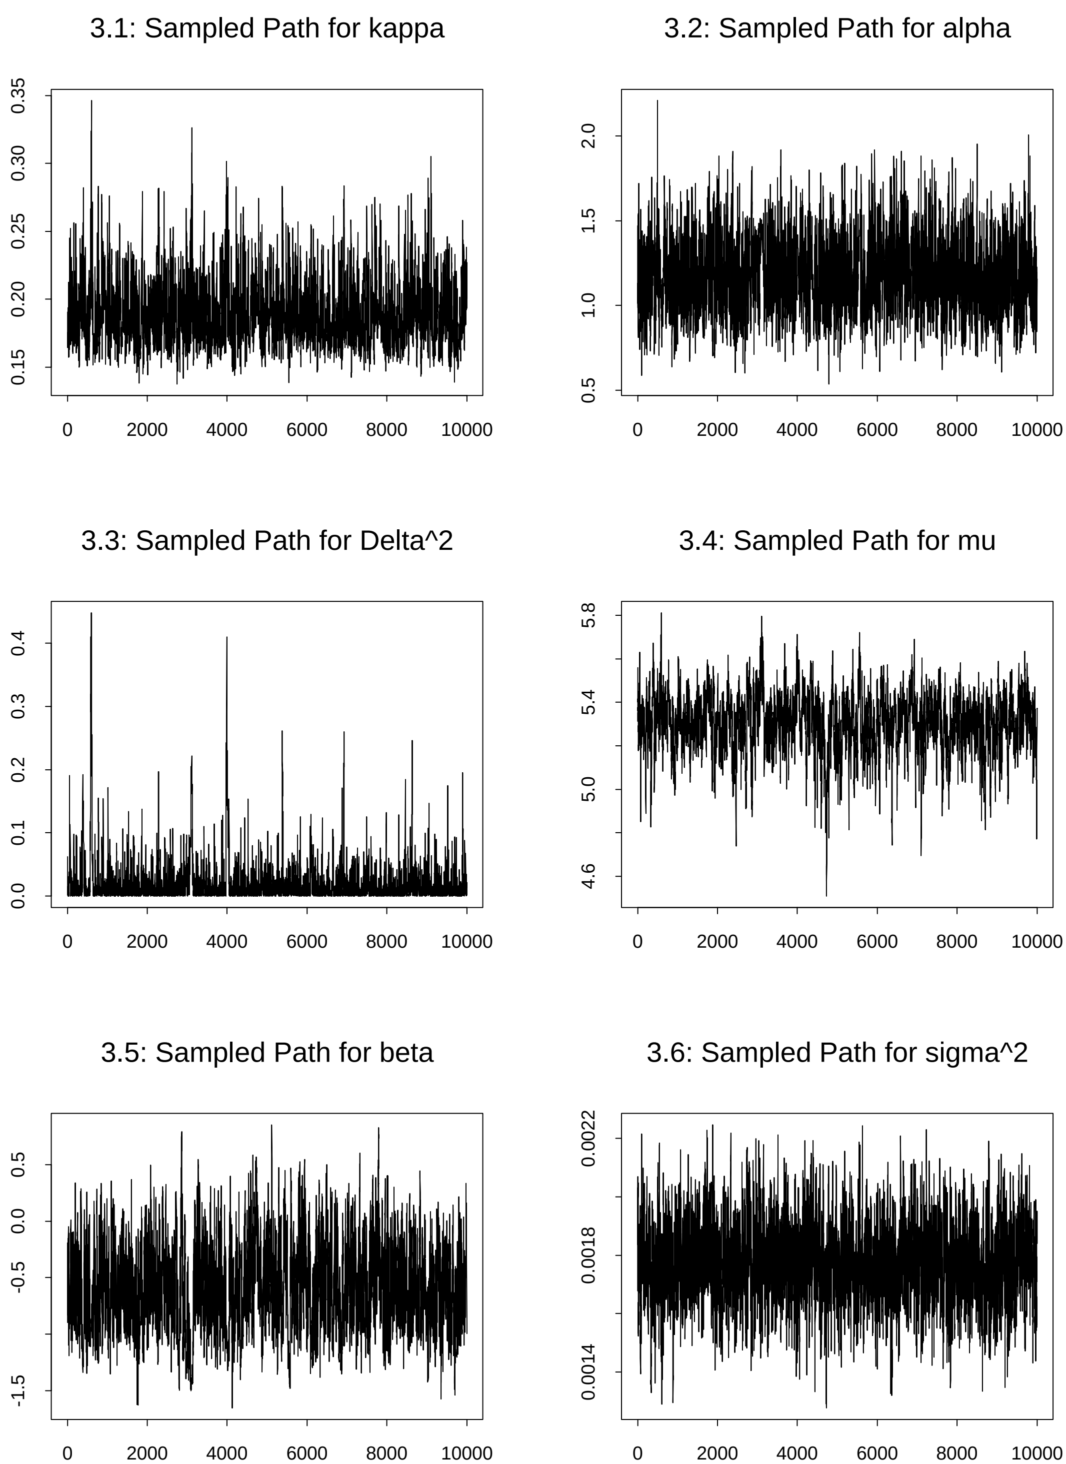

Figure 4. ACF for the Euler Scheme