

Numerical Techniques for Maximum Likelihood Estimation of Continuous-Time Diffusion Processes

Garland B. Durham and A. Ronald Gallant*

November 9, 2001

Abstract

Stochastic differential equations often provide a convenient way to describe the dynamics of economic and financial data, and a great deal of effort has been expended searching for efficient ways to estimate models based on them. Maximum likelihood is typically the estimator of choice; however, since the transition density is generally unknown, one is forced to approximate it. The simulation-based approach suggested by Pedersen (1995) has great theoretical appeal, but previously available implementations have been computationally costly. We examine a variety of numerical techniques designed to improve the performance of this approach. Synthetic data generated by a CIR model with parameters calibrated to match monthly observations of the U.S. short-term interest rate are used as a test case. Since the likelihood function of this process is known, the quality of the approximations can be easily evaluated. On data sets with 1000 observations, we are able to approximate the maximum likelihood estimator with negligible error in well under one minute. This represents something on the order of a 10,000-fold reduction in computational effort as compared to implementations without these enhancements. With other parameter settings designed to stress the methodology, performance remains strong. These ideas are easily generalized to multivariate settings and (with some additional work) to latent variable models. To illustrate, we estimate a simple stochastic volatility model of the U.S. short-term interest rate.

*We are grateful for the helpful comments and suggestions of Yacine Aït-Sahalia, Evan Anderson, Tim Bollerslev, Michael Brandt, Ai-Ru Cheng, Frank Diebold, Eric Ghysels, Adrian Pagan, Neil Shephard, George Tauchen, and seminar participants at the NBER Summer Institute 2000 and the University of Pennsylvania. + Garland B. Durham, Department of Economics, W376 Papajohn Bldg., University of Iowa, Iowa City, IA 52242-1000, USA; Phone: 319 335 0844; email: garland-durham@uiowa.edu

A. Ronald Gallant, Department of Economics, CB# 3305, University of North Carolina, Chapel Hill NC 27599-3305, USA; Phone: 919 966 5338; email: ron_gallant@unc.edu

First draft: December 1, 1999

1 Introduction

Stochastic differential equations (SDE) often provide a convenient way to model economic and financial data, and their use has become increasingly common in recent years. Although the process specified by a stochastic differential equation is defined in continuous time, the data which is available is typically sampled at discrete time intervals. The resulting estimation problem turns out to be nontrivial, and considerable energy has been expended in developing computationally (and statistically) efficient estimation schemes.

In this paper we focus primarily on scalar, time-homogeneous processes. In particular, we consider the diffusion process generated by a SDE of the form

$$\begin{aligned}dX &= \mu(X; \theta) dt + \sigma(X; \theta) dW \\ X(t_0) &= X_0\end{aligned}\tag{1}$$

with parameter vector θ . Suppose that the sample $\{X_i = X(t_i), i = 0, \dots, n\}$ is available for analysis. The observations need not be equally spaced.

Ideally, one would like to know the transition density, which would allow one to compute the maximum likelihood estimator with its usual optimality properties. Although exact transition densities are known in only a few isolated cases, several approaches toward approximating the transition density have been proposed.

Lo (1988) suggests numerically solving the Fokker-Planck partial differential equation for each observation. Pedersen (1995b) suggests a simulation-based approach which involves integrating out unobserved states of the process at intermediate points between each pair of observations (see also Santa-Clara 1995; Brandt and Santa-Clara 2001). While this approach, commonly known as simulated maximum likelihood estimation (SMLE), is able to come arbitrarily close to the true transition density, previously available implementations have been computationally burdensome.

Other approaches have been proposed which are much less computationally costly. For example, the process described by (1) has a first-order approximation given by the discrete-time process

$$\begin{aligned}\tilde{X}_{i+1} &= \tilde{X}_i + \mu(\tilde{X}_i; \theta)\Delta_i + \sigma(\tilde{X}_i; \theta)\Delta_i^{1/2}\epsilon_i \\ \Delta_i &= t_{i+1} - t_i, \quad \epsilon_i \sim N(0, 1).\end{aligned}\tag{2}$$

Under mild regularity conditions, the maximum likelihood estimator based on this approximation is known to converge to the true maximum likelihood estimator as the sampling interval goes to zero (Florens-Zmirou 1989). While this approach is very appealing from a computational viewpoint, the approximation may not be sufficiently accurate for the sampling frequencies at which reliable data is available.

There are various ways which one might improve upon this idea. Elerian (1998) suggests

replacing the Gaussian density in (2) by a noncentral chi-squared density which is derived from the Milstein scheme, an order 2.0 weak approximation to the true process. Shoji and Ozaki (1998) linearize the SDE, obtaining an approximating Ornstein-Uhlenbeck process (the exact transition density of an Ornstein-Uhlenbeck process is known). Kessler (1997) approximates the transition function by a Gaussian density with first and second moments obtained from higher-order Ito-Taylor expansions. Aït-Sahalia (2001) approximates the transition density using a Hermite function with coefficients obtained using higher-order Ito-Taylor expansions. Except for Aït-Sahalia (2001), these methods still require the sampling interval to go to zero to obtain convergence to the true transition density. While this requirement also holds for Aït-Sahalia's approach with Hermite function and Ito-Taylor expansion of fixed order, Aït-Sahalia's approximation may be made arbitrarily accurate with fixed sampling frequency by using a Hermite function and Ito-Taylor expansion of sufficiently high order (given some regularity conditions).

Various method of moments approaches have also been proposed. Chan, Karolyi, Longstaff and Sanders (1992) use moments based on equation (2). Duffie and Singleton (1993), Gallant and Tauchen (1997), Bibby and Sørensen (1995), and Gouriéroux, Monfort and Renault (1993) compute expectations using simulation-based methods. Hansen and Scheinkman (1995) and Duffie and Glynn (1996) use moment conditions obtained from the infinitesimal generator.

The simulation-based methods can be computationally costly but have the advantage of being easily adapted to diffusions with unobserved state variables. Stochastic volatility models and term structure models are important applications where these techniques have been found useful. The Efficient Method of Moments proposed by Gallant and Tauchen (1996) approaches the efficiency of maximum likelihood asymptotically and provides a convenient set of diagnostic measures for model specification.

Markov chain Monte Carlo (MCMC) methods have been proposed by Eraker (2001), Jones (1999a), and Elerian, Chib and Shephard (2001). There is a close relationship between MCMC methods and SMLE. For example, Elerian et al. point out that their importance sampler can also be used with the simulation-based approach of Pedersen (1995b) to substantially reduce the computational effort required to obtain reasonably accurate likelihood approximations.

In this paper we focus on the SMLE approach. The basic idea is quite simple. Suppose that one wishes to obtain the transition density $p(x_t, t; x_s, s)$. The first-order approximation $p^{(1)}(x_t, t; x_s, s)$ defined by (2) will be accurate if the interval $[s, t]$ is sufficiently short. Otherwise, one may partition the interval $s = \tau_1 < \tau_2 < \dots < \tau_M = t$ such that the first-order approximation is sufficiently accurate on each subinterval. The random variables $X(\tau_1), \dots, X(\tau_{M-1})$ are, of course, unobserved and must be integrated out. Because the

process is Markovian, one obtains

$$p(x_t, t; x_s, s) \approx p^{(M)}(x_t, t; x_s, s) \quad (3)$$

$$\equiv \int \prod_{m=0}^{M-1} p^{(1)}(u_{m+1}, \tau_{m+1}; u_m, \tau_m) d\lambda(u_1, \dots, u_{M-1}) \quad (4)$$

where λ denotes Lebesgue measure and we use the convention $u_0 = x_s$ and $u_M = x_t$ to conserve notation. Monte Carlo integration is generally the only feasible way to evaluate this integral.

The theoretical issues involved with this approach are already reasonably well understood. Sufficient conditions for the approximation in (3) to converge are known. While it is certainly of value to extend these conditions, we do not undertake this task here. The theories of Monte Carlo integration and maximum likelihood estimation have also been extensively studied. Nonetheless, although the simulation-based approach is attractive from a theoretical point of view, the computational burden associated with previous implementations has hindered its widespread use. We have found that it can be quite costly to attain even the degree of accuracy provided by the simple first-order approximation (2). It is this shortcoming which we seek to address.

We attack the problem of computational efficiency from two directions. We first seek to improve the approximation in equation (3). This allows one to attain a given level of accuracy with fewer intermediate points. We consider extrapolation techniques and the use of alternatives to the first-order (Euler) approximation of the subtransition densities. Secondly, we examine techniques to accelerate the convergence of the Monte Carlo integration. We consider several importance samplers and random schemes. Finally, we consider transforming the model in such a way as to make the volatility function constant. Working with the transformed rather than the original model turns out to provide a useful improvement in both the accuracy of the approximation (3) as well as the performance of the Monte Carlo integration used to compute (4).

As a test case, we use the square-root specification proposed by Cox, Ingersoll and Ross (1985) as a model for the short-term interest rate. Parameter settings are calibrated to match monthly observations of U.S. short-term interest rate. This model has the advantage that the transition density is available in closed form, which allows us to easily evaluate the accuracy of our approximations. We also tested our techniques using other parameter settings and models with similar results.

On simulated data sets of 1000 observations we are able to obtain estimates in well under one minute (running Fortran code on a 750 MHz PC) which differ negligibly from those obtained by maximizing the exact log-likelihood function. Achieving comparable accuracy without our acceleration techniques would require something on the order of a 10,000-fold

increase in computational effort.

Much of the discussion in this paper may be readily adapted to the multivariate setting. With some additional work, the ideas can also be extended to latent variable models. We outline an approach to approximating the transition density of a continuous-time stochastic volatility model and illustrate by estimating a simple model over weekly observations of the U.S. Treasury bill rate. We speculate that much carries over to the time-inhomogeneous case as well, however, we have not examined such extensions carefully. Although it should be possible to apply techniques similar to those considered here to jump diffusions, this is also beyond the scope of this paper.

A more extensive application illustrating the techniques discussed in this paper may be found in Durham (2000). Further exploration of these and related techniques in multivariate and latent variable settings is underway.

The structure of this paper is as follows. Section 2 introduces the notation and provides some theoretical results, Section 3 describes the benchmarks which we shall use for evaluation of our techniques, Section 4 examines the performance of the simulation-based method without any of our acceleration techniques, Section 5 considers the issue of bias-reduction, Section 6 considers the issue of variance-reduction, Section 7 discusses the results of our numerical experiments, Section 8 extends these ideas to the stochastic volatility model, Section 9 provides an application, and Section 10 concludes.

2 Background

To begin, we define some notation and provide a brief discussion of the theoretical framework. Let (Ω, \mathcal{F}, P) be a probability space, and let W be a Brownian motion defined on it. Let $\{\mathcal{F}_t, t \geq 0\}$ be the filtration generated by W and augmented by the P -null sets of \mathcal{F} . Let Θ be a compact subset of \mathbb{R}^d . We are interested in the parameterized family of scalar diffusion processes $\{X(t; \theta), \theta \in \Theta\}$ generated by the time-homogeneous SDE

$$\begin{aligned} dX &= \mu(X; \theta) dt + \sigma(X; \theta) dW \\ X(t_0; \theta) &= X_0. \end{aligned}$$

Assumption 1. For each $\theta \in \Theta$, (5) has a nonexploding, unique weak solution.

By nonexploding, we mean that there is zero probability that the process diverges to infinity over any fixed time interval. Sufficient conditions ensuring Assumption 1 are well known (e.g., Karatzas and Shreve 1991, Section 5.5). For example, it suffices that μ and σ satisfy global Lipschitz and linear growth conditions. A variety of extensions are also available. Explosiveness would preclude the existence of a transition density and is thus disallowed. Note that stationarity is not required.

For $s < t$, suppose that $X(t; \theta) | X(s; \theta)$ has a transition density $p(x_t, t; x_s, s, \theta)$, and let

$$p^{(1)}(x_t, t; x_s, s, \theta) = \phi(x_t; x_s + \mu(x_s)(t-s), \sigma^2(x_s)(t-s)), \quad (5)$$

where $\phi(x; \mu, \sigma^2)$ is the Gaussian density, be its first-order approximation. Let $s = \tau_0 < \dots < \tau_M = t$ be a partition of the interval $[s, t]$, and let

$$p^{(M)}(x_t, t; x_s, s, \theta) = \int \prod_{m=1}^M p^{(1)}(u_m, \tau_m; u_{m-1}, \tau_{m-1}, \theta) d\lambda(u_1, \dots, u_{M-1}), \quad (6)$$

where $u_0 = x_s$, $u_M = x_t$, and λ denotes Lebesgue measure. This will serve as our approximating density. For clarity, we will often refer to $p^{(1)}(\cdot)$ as a *subtransition density* (or occasionally simply *subdensity*) when used in this context.

Suppose that one has a set of observations $\{X_i = X(t_i; \theta^o), i = 0, \dots, n\}$ of the process generated by (5) with unknown parameter vector θ^o , and let $P_{\theta^o, n}$ denote the probability measure induced by $\{X_0, \dots, X_n\}$. Let $l_n(\theta) = \sum_{i=1}^n \log p(X_i, t_i; X_{i-1}, t_{i-1}, \theta)$ and $l_n^{(M)}(\theta) = \sum_{i=1}^n \log p^{(M)}(X_i, t_i; X_{i-1}, t_{i-1}, \theta)$ denote the log likelihood functions associated with the exact and approximate densities respectively.

Assumption 2. For all $s < t$, $x_s \in \text{supp}[X(s; \theta^o)]$, $\theta \in \Theta$, and $M \geq 1$, the densities $p(\cdot, t; x_s, s, \theta)$ and $p^{(M)}(\cdot, t; x_s, s, \theta)$ exist.

Pedersen (1995a) provides sufficient conditions for Assumption 2 to hold as well as regularity conditions ensuring that

$$\lim_{M \rightarrow \infty} p^{(M)}(\cdot, t; x_s, s, \theta) = p(\cdot, t; x_s, s, \theta) \quad \text{in } L^1(\lambda). \quad (7)$$

We note that Pedersen's results are obtained for multivariate processes. Pedersen's Theorem 2 allows for time-inhomogeneous processes. While this theorem requires a constant diffusion function, we shall see in Section 3 that, for scalar processes at least, this does not impose a material constraint. Pedersen's Theorem 3 allows for a variable diffusion function but imposes other conditions.

Although Pedersen's results assume Lipschitz and linear growth conditions on $\mu(\cdot)$ and $\sigma(\cdot)$ that are not satisfied for many applications of economic interest (including notably the CIR square root process), we speculate that suitable extensions should be possible using localization arguments along the lines of e.g. Karatzas and Shreve (1991), Theorem 5.2.5. Similarly, we will examine subtransition densities other than the simple first-order approximation shown in (5) (see Section 5) and alternative random number schemes (see Section 6) which are not covered by Pedersen's results. Again, these extensions seem plausible but formal justification is left for future work. The goal of this paper is practical rather than theoretical, and Pedersen's results will serve as a convenient starting point. In particular, we assume the following, which Pedersen's Theorem 4 shows to be an immediate consequence of (7).

Assumption 3. For each $\theta \in \Theta$,

$$\lim_{M \rightarrow \infty} l_n^{(M)}(\theta) = l_n(\theta) \quad \text{in probability under } P_{\theta^e, n}.$$

The difficulty is how to efficiently evaluate the integral in equation (6). Monte Carlo integration is generally the only feasible approach. To perform Monte Carlo integration, one requires an importance sampler. Fix $s < t$, x_s , x_t , θ , and M , and let $q(u_1, \dots, u_{M-1})$ denote a probability density on \mathbb{R}^{M-1} . This will be our importance sampler. Some techniques for constructing efficient importance samplers are discussed in Section 6.

Let $\{\mathbf{u}_k = (u_{k,1}, \dots, u_{k,M-1}), k = 1, \dots, K\}$ be independent draws from q , and let

$$p^{(M,K)}(x_t, t; x_s, s, \theta) = \frac{1}{K} \sum_{k=1}^K \frac{\prod_{m=1}^M p^{(1)}(u_{k,m}, \tau_m; u_{k,m-1}, \tau_{m-1}, \theta)}{q(u_{k,1}, \dots, u_{k,M-1})} \quad (8)$$

where $u_{k,0} = x_s$ and $u_{k,M} = x_t$ for all k . Then given Assumption 4 below, the strong law of large numbers implies that

$$\lim_{K \rightarrow \infty} \left| p^{(M,K)}(x_t, t; x_s, s, \theta) - p^{(M)}(x_t, t; x_s, s, \theta) \right| = 0 \quad \text{a.s.} \quad (9)$$

A somewhat stronger condition provides \sqrt{n} -convergence (see Geweke 1989).

Assumption 4. Let $U_0 = x_s$, $U_M = x_t$, $\theta \in \Theta$, and q be fixed, and let (U_1, \dots, U_{M-1}) be a random vector with density q . Then

$$E \left[\frac{\prod_{m=1}^M p^{(1)}(U_m, \tau_m; U_{m-1}, \tau_{m-1}, \theta)}{q(U_1, \dots, U_{M-1})} \right] < \infty.$$

Our goal is to approximate $\log l_n(\theta)$ for a given realization of the process. For this it will suffice to be able to approximate $p(x_t, t; x_s, s, \theta)$ for arbitrary $s < t$, x_s , x_t , and θ . If we can do this with arbitrary precision, and if the log likelihood function is continuous and Θ is compact, then we can evaluate the maximum likelihood estimator at this realization with any desired level of accuracy. We do not treat the estimator obtained by optimizing the approximate log likelihood with a fixed setting of the tuning parameters as an object of independent interest.

3 Benchmarks

The specification

$$dX = \theta_2(\theta_1 - X)dt + \theta_3\sqrt{X}dW \quad (10)$$

with θ_1 , θ_2 , and θ_3 positive was proposed by Cox et al. (1985) to model short-term interest rates. Since this model has a known transition density and is frequently used in applications, it provides a convenient means of evaluating the effectiveness of our numerical methods. If we let $\Delta = t - s > 0$, $c = 2\theta_2/[\theta_3^2(1 - e^{-\theta_2\Delta})]$, and $Y = 2cX$, then $Y_t|Y_s$ is distributed as noncentral chi-squared with $4\theta_2\theta_1/\theta_3^2$ degrees of freedom and noncentrality parameter $Y_s e^{\theta_2\Delta}$. Or equivalently,

$$p(x_t, t; x_s, s) = ce^{-u-v}(v/u)^{q/2} I_q(2\sqrt{uv}), \quad (11)$$

where $u = cx_s e^{-\theta_2\Delta}$, $v = cx_t$, $q = 2\theta_2\theta_1/\theta_3^2 - 1$, and $I_q(\cdot)$ is the modified Bessel function of the first kind of order q .

For any experiments where synthetic data from the CIR model is required, we generate it directly using draws from the noncentral chi-squared.

Our base case uses the parameter settings $\theta^o = (0.06, 0.5, 0.15)$ and $\Delta = 1/12$. These settings are identical to those used in Ait-Sahalia (2001) for ease of comparison, and are said to be calibrated to match monthly observations of the U.S. treasury bill rate. We also test the methods discussed in this paper with other models and parameter settings with similar results.

We have found that better results are often obtained if the SDE is first transformed to make the diffusion term of constant magnitude. With the CIR model, for example, setting $Y = \sqrt{X}$ and applying Ito's Lemma gives

$$dY = \left[\frac{\theta_2}{2Y}(\theta_1 - Y^2) - \frac{\theta_3^2}{8Y} \right] dt + \frac{\theta_3}{2} dW.$$

If $p_Y(y_t, t; y_s, s)$ denotes the transition density of the transformed process, then the density of the original process is obtained in the usual manner by

$$p(x_t, t; x_s, s) = p_Y(y_t, t; y_s, s) \left| \frac{dy}{dx} \right| = \frac{p_Y(\sqrt{x_t}, t; \sqrt{x_s}, s)}{2\sqrt{x_t}}.$$

In general, the appropriate transformation is given by $Y = G(X)$, where G satisfies

$G'(x) = 1/\sigma(x)$. The constant of integration is irrelevant. Ito's Lemma then implies

$$\begin{aligned} dY &= G'(X) dX + \frac{1}{2}G''(X)\sigma^2(X) dt \\ &= \left[\frac{\mu(X)}{\sigma(X)} - \frac{1}{2}\sigma'(X) \right] dt + dW \\ &= \left[\frac{\mu[G^{-1}(Y)]}{\sigma[G^{-1}(Y)]} - \frac{1}{2}\sigma'[G^{-1}(Y)] \right] dt + dW. \end{aligned}$$

In many cases G can be obtained analytically; otherwise, it may require a numerical integration. This does not pose any serious difficulties. If the parameter vector enters into σ nonlinearly, the transformation will have to be recomputed for each candidate parameter, which may be inconvenient.

This transformation goes back to at least Doss (1977) and is also used by Shoji and Ozaki (1998) and Ait-Sahalia (2001). (In contrast to those papers, our methodology does not *require* that the model be transformed.) The reason underlying its effectiveness appears to be that it makes the process closer to Gaussian. This improves the performance of the approximation $p^{(M)}$ as well as that of the importance samplers.

We compare the effectiveness of the various approximation techniques using several different measures. First, we look at some density plots. We fix a value for x_s and consider a range of values for x_t . For each value of x_t , we approximate the density $p(x_t, t; x_s, s)$ a number of times using different seeds for the random number generator. We then compute the difference between the true and approximate log densities and plot the median and interquartile range of the approximation errors for each x_t .

The figures in this paper are obtained using the process defined by (10), with $\Delta = t - s = 1/12$, $x_s = 0.10$, $x_t \in [.05, .15]$, and 1024 repetitions. For reference, the exact transition density with these settings is shown in Figure 1.

Since the object of ultimate interest is the log likelihood, it seems appropriate to examine the error in the log rather than the level of the density. Suppose, for example, that $p(x_t, t; x_s, s) = 10$, and we are able to compute it with an error of ± 0.1 . This term would contribute an error of ± 0.01 to the log likelihood. On the other hand, if $p(x_t, t; x_s, s) = 0.2$, the contribution to the log likelihood of the same approximation error would be in the range of $[-0.7, 0.4]$. If the approximation error is greater than the level of the density, the computed density can be negative. This is clearly catastrophic for the log likelihood. These distinctions are obscured if one examines the approximation error for the level of the density.

The second measure which we examine is the root mean squared error (RMSE) of the log density approximation. We approximate this by generating $n = 100,000$ simulated

observations from the model and computing a sample analog, i.e.

$$\text{RMSE} = \left\{ \int \left(\log \hat{p}(y|x) - \log p(y|x) \right)^2 p(y, x) dy dx \right\}^{1/2} \quad (12)$$

$$\approx \left\{ \frac{1}{n} \sum_{i=1}^n \left(\log \hat{p}(x_{i+1}|x_i) - \log p(x_{i+1}|x_i) \right)^2 \right\}^{1/2} \quad (13)$$

where we have denoted the approximate transition density by \hat{p} . It is convenient to assume that the integral in (12) exists. At any rate, the sum in (13) certainly exists for a fixed realization $\{x_1, \dots, x_n\}$, which is all we really need in order to compare across approximation techniques.

Finally, we are interested in the accuracy of the parameter estimates obtained by maximizing the approximate rather than the exact log likelihood. To measure this, we generate $J = 512$ data sets of length $n = 1000$ and compute $\hat{\theta}$ for each repetition using the exact log likelihood and the various approximations. We compute the RMSE of the exact maximum likelihood estimates with respect to the parameter vector used to actually generate the data, and the RMSE of the simulated maximum likelihood estimates with respect to the exact maximum likelihood estimates, i.e.

$$\begin{aligned} \text{RMSE}_{\text{TRUE-MLE}} &= \left\{ \frac{1}{J} \sum_{j=1}^J (\hat{\theta}_{\text{MLE}} - \theta^o)^2 \right\}^{1/2} \\ \text{RMSE}_{\text{MLE-SMLE}} &= \left\{ \frac{1}{J} \sum_{j=1}^J (\hat{\theta}_{\text{MLE}} - \hat{\theta}_{\text{SMLE}})^2 \right\}^{1/2}. \end{aligned}$$

A reasonable goal might be to obtain an approximation error on the order of 1% of the error inherent in the MLE itself. We are able to easily obtain this goal for our test case.

Virtually any method which one might reasonably consider should be able to approximate the log likelihood function with arbitrary precision given sufficient time. The key issue is how quickly one is able to obtain sufficiently accurate results. Thus we also report computational costs.

As a matter of implementation, variance in the Monte Carlo integral can result in a great deal of jaggedness in the likelihood surface, which will severely degrade the performance of the optimizer. However, this issue is easily addressed if for each evaluation of the likelihood function one uses the same seed for the random number generator used to draw samples for the Monte Carlo integration. This is especially critical if one is computing numerical derivatives. At any rate, for many of the methods which we examine, it is relatively straightforward to obtain analytical derivatives.

4 Simulation method without acceleration techniques

To establish a baseline, we begin by examining the simulation method as implemented by Pedersen (1995b), i.e., without any of our acceleration techniques. The importance sampler used by Pedersen is constructed by simulating paths on each subdivided interval using the Euler scheme. Suppose that $s < t$, $x_s = X(s)$, and $x_t = X(t)$ are given. The importance sampler is defined by the mapping $T^{(M)} : (W_1, \dots, W_{M-1}; \theta) \mapsto (u_1, \dots, u_{M-1})$ given by the recursion

$$u_{m+1} = u_m + \mu(u_m; \theta)\delta + \sigma(u_m; \theta)\delta^{1/2}W_{m+1}, \quad m = 0, \dots, M-2 \quad (14)$$

where $u_0 = x_s$, $\delta = (t-s)/M$, and $W = (W_1, \dots, W_{M-1})$ is a multivariate standard normal.

In this case, equation (8) simplifies considerably. Since the density of the importance sampler q is identical to the first $M-1$ factors of the numerator, they cancel and one is left with

$$p^{(M,K)}(x_t, t; x_s, s, \theta) = \frac{1}{K} \sum_{k=1}^K p^{(1)}(x_t, t; u_{k,M-1}, \tau_{M-1}, \theta), \quad (15)$$

where the $\{u_{k,M-1}, k = 1, \dots, K\}$ are drawn from the $(M-1)^{\text{st}}$ component of q . An alternate interpretation of equation (15) is to consider the right-hand side as the sample analog of $E\left[p^{(1)}(x_t, t; u_{M-1}, \tau_{M-1}, \theta)\right]$, where the expectation is over u and with respect to the distribution induced by $X(\tau_{M-1})|X(s) = x_s$.

Throughout this paper, we use the method of antithetic variates when drawing random numbers. This is a commonly used variance-reduction technique in simulation-based methods. To implement antithetic variates, one draws only $K/2$ samples from the multivariate normal and simulates two paths from each: $T_+^{(M)}(W) = T^{(M)}(W)$ is as described above, and $T_-^{(M)}(W) = T^{(M)}(-W)$ is its ‘mirror image’. While we have found antithetic variates to provide only marginal benefit, the cost is also negligible.

Figure 2 illustrates the approximation error which results from computing the log density using this approach. The settings $K = 256$ and $M = 8$ or $M = 32$ are used (recall that K is the number of sample paths and M is the number of subintervals). Panels (a)-(c) use the untransformed model. Panel (d) uses the transformed model, which appears to provide little benefit in this case. Increasing M reduces bias, but at the cost of greater variance. Reducing the variance is costly since it is $\mathcal{O}(K^{-1/2})$.

Upon comparison with the tables and figures in Section 7, the reason why this approach has not seen wide-spread use is readily apparent. It would take a great deal of effort even to match the accuracy of the simple first-order approximation, at least for our test model.

5 Bias-reduction techniques

There are two sources of approximation error which we wish to address: bias due to the first-order approximation used in the construction of $p^{(M)}$, and variance resulting from the Monte Carlo integration.

We begin with the bias. While it is possible to drive the bias to zero by partitioning the intervals between observations sufficiently finely, this can be computationally costly. We examine two approaches toward reducing the number of subintervals required to obtain a given level of accuracy. The first is to replace the first-order approximation used in equation (5) by a higher-order method. There are several possibilities which one might try.

Elerian (1998) suggests using a transition density derived from a scheme due to Milstein (1978). If the volatility function $\sigma(\cdot)$ is constant, the density is identical to that of the first-order approximation, otherwise it is given by

$$p_{\text{Elerian}}(x_t, t; x_s, s) = \frac{z_t^{-1/2}}{|A|\sqrt{2\pi}} \exp\left(-\frac{C + z_t}{2}\right) \cosh(\sqrt{Cz_t})$$

where

$$\begin{aligned} z_t &= \frac{x_t - B}{A} \\ \Delta &= t - s \\ A &= \frac{\sigma(x_s)\sigma'(x_s)\Delta}{2} \\ B &= -\frac{\sigma(x_s)}{2\sigma'(x_s)} + x_s + \mu(x_s)\Delta - \frac{\sigma(x_s)\sigma'(x_s)\Delta}{2} \\ C &= \frac{1}{(\sigma'(x_s))^2\Delta}. \end{aligned}$$

Note that $\sigma'(\cdot)$ denotes the derivative of σ .

Kessler (1997) suggests using a Gaussian transition density, but rather than using the first-order approximations for the mean and variance, he proposes using higher-order Ito-Taylor approximations. We try a second-order implementation, i.e.

$$p_{\text{Kessler}}(x_t, t; x_s, s) = \phi(x_t; \tilde{\mu}, \tilde{\sigma}^2)$$

where

$$\begin{aligned} \tilde{\mu} &= x_s + \mu(x_s)\Delta + \left[\mu(x_s)\mu'(x_s) + \frac{\sigma^2(x_s)\mu''(x_s)}{2} \right] \frac{\Delta^2}{2} \\ \tilde{\sigma}^2 &= x_s^2 + \left\{ 2\mu(x_s)x_s + \sigma^2(x_s) \right\} \Delta + \left\{ 2\mu(x_s)[\mu'(x_s)x_s + \mu(x_s)] \right. \\ &\quad \left. + \sigma(x_s)\sigma'(x_s) \right\} + \sigma^2(x_s)[\mu''(x_s)x_s + 2\mu'(x_s) + \sigma'(x_s)\sigma'(x_s) \\ &\quad \left. + \sigma(x_s)\sigma''(x_s) \right\} \frac{\Delta^2}{2} - \tilde{\mu}^2. \end{aligned}$$

Notice that for some models and parameter settings, it is possible to obtain $\tilde{\sigma}^2 < 0$. The code should include a check to watch out for this.

Shoji and Ozaki (1998) suggest a method which they refer to as *local linearization*. Their approach requires a model with constant volatility; however, as shown in Section 3, this results in little loss of generality. Given

$$dX = \mu(X) dt + \sigma dW$$

(σ is constant) and fixed x_s , one begins with an application of Ito's Lemma,

$$d\mu(X) = \frac{1}{2}\sigma^2\mu''(X) dt + \mu'(X) dX.$$

Using the first-order Taylor expansion, we define

$$\tilde{\mu}(x_t) = \mu(x_s) + \mu'(x_s)(x_t - x_s) + \frac{1}{2}\sigma^2\mu''(x_s)(t - s).$$

The approximate density will be obtained from

$$d\tilde{X} = \tilde{\mu}(\tilde{X}) dt + \sigma dW,$$

which is an Ornstein-Uhlenbeck process. One obtains

$$p_{\text{Sh\&Oz}}(x_t, t; x_s, s) = \phi(x_t; \tilde{\mu}, \tilde{\sigma}^2)$$

where

$$\begin{aligned}\tilde{\mu} &= x_s + \frac{\mu(x_s)}{\mu'(x_s)}K + \frac{\sigma^2\mu''(x_s)}{2[\mu'(x_s)]^2}[K - \mu'(x_s)\Delta] \\ \tilde{\sigma}^2 &= \frac{\sigma^2}{2\mu'(x_s)}\left(e^{2\mu'(x_s)\Delta} - 1\right) \\ K &= e^{\mu'(x_s)\Delta} - 1.\end{aligned}$$

Nowman (1997) suggests a similar approach, but simply treating the volatility as if it were constant on each sample interval rather than transforming the model so that it actually is constant. While he examines only the special case where the drift function is linear, a plausible extension would be to use a first-order Taylor expansion for the drift function as described above. The resulting approximation is analogous to that of Shoji and Ozaki, but replacing σ with $\sigma(x_s)$.

While Elerian (1998) uses the Milstein density in the context of a simulation-based approach, Kessler (1997), Shoji and Ozaki (1998), and Nowman (1997) approximate the transition density between observations directly (i.e., without using intermediate points).

Another approach to obtaining higher-order methods is extrapolation. Given, for example, a first-order method, one may construct a second-order method as follows:

$$\begin{aligned} p^{(M)} &= p + K\Delta + \mathcal{O}(\Delta^2) \\ p^{(2M)} &= p + K\Delta/2 + \mathcal{O}(\Delta^2) \\ p_E^{(2M)} &= 2p^{(2M)} - p^{(M)} \\ &= p + \mathcal{O}(\Delta^2) \end{aligned}$$

where K is some unknown constant. If the approximate likelihoods are stochastic (i.e., computed by simulation), extrapolation reduces bias but at the cost of greater variance. Since it is possible to obtain a negative value for the extrapolated density, any implementation of this technique should check for positivity and fall back to the non-extrapolated value in case of trouble.

Extrapolation is a well-known bias-reduction technique for computing expectations of diffusion processes (see Kloeden and Platen 1992, Section 15.3). That we are able to apply the technique in the present context is because our approach to approximating the transition density is essentially an expectation. For Pedersen's method, it is easy to see from equation (15) that

$$p^{(M)}(x_t, t; x_s, s) = \int p^{(1)}(x_t, t; u, \tau_{M-1}, \theta) dP_{M-1}^{(M)}(u) \quad (16)$$

where $P_{M-1}^{(M)}$ is the measure induced by the Euler scheme at τ_{M-1} . In general, one obtains

$$p^{(M)}(x_t, t; x_s, s) = \int p^{(1)}(x_t, t; u, \tau_{M-1}, \theta) \rho_{M-1}^{(M)}(u) dQ_{M-1}^{(M)}(u) \quad (17)$$

where $Q_{M-1}^{(M)}$ is the measure induced by the $(M-1)^{\text{st}}$ component of the importance sampler and $\rho_{M-1}^{(M)}$ is the Radon-Nikodym derivative of $P_{M-1}^{(M)}$ with respect to $Q_{M-1}^{(M)}$. These expressions may also be derived directly from (6).

6 Variance-reduction techniques

We examine two approaches to reducing the variance of the Monte Carlo integration shown in equation (8): importance sampling and random number schemes. Some of the techniques are illustrated in Figure 3.

A basic principle of Monte Carlo integration is that one should draw points with higher probability in regions where the integrand is larger. Figure 4 illustrates why Pedersen's method performs so poorly. The paths in the figure are sampled using the Euler scheme with $x_s = 0.08$, $\Delta = t - s = 1/12$, and the SDE given in Section 3. The terminal point of each path represents a draw from $P_{M-1}^{(M)}$. The curve represents the integrand of the right-hand side of (16) as a function of u with $x_t = 0.11$. It is clear that most of the samples are drawn from regions where the integrand has little mass. The importance samplers discussed in this section are designed to address this shortcoming. Elerian et al. (2001) appear to have been the first to consider the idea of using efficient importance sampling in this context.

The first importance sampler we consider is based on the Brownian bridge. A Brownian bridge is a Brownian motion started at x_s at time s and conditioned to terminate at x_t at time t . The sampler is constructed in a manner similar to the Euler scheme. In this case, the mapping $T^{(M)} : (W_1, \dots, W_{M-1}; \theta) \mapsto (u_1, \dots, u_{M-1})$ is defined by the recursion

$$u_{m+1} = u_m + \tilde{\mu}(u_m, \tau_m)\delta + \sigma(u_m; \theta)\delta^{1/2}W_{m+1}$$

where the drift is given by

$$\tilde{\mu}(x, \tau) = \frac{x_t - x}{t - \tau}.$$

This is a true Brownian bridge if and only if σ is constant (which will be the case if we first transform the model as discussed in Section 3).

Although it is certainly possible to compute the approximate density directly from (8), there is an interesting interpretation of this sampler based on Girsanov's Theorem. Consider the processes $dX = \mu(X) dt + \sigma(X) dW$ and $d\tilde{X} = \tilde{\mu}(\tilde{X}) dt + \sigma(\tilde{X}) d\tilde{W}$ with initial condition $X(s) = \tilde{X}(s) = x_s$. Girsanov's Theorem tells us that the Radon-Nikodym derivative of the probability measure generated by X with respect to that generated by \tilde{X} is given by

$$d\rho = \rho k(\tilde{X}) d\tilde{W} \tag{18}$$

with initial condition $\rho(s) = 1$ and

$$k(x) = \frac{\mu(x) - \tilde{\mu}(x)}{\sigma(x)}.$$

We can thus obtain the continuous-time expression

$$p(x_t, t; x_s, s) = \int p(x_t, t; u, \tau_{M-1}) \rho_{M-1}(u) dQ_{M-1}(u)$$

where Q_{M-1} is the probability measure induced by $\tilde{X}(\tau_{M-1})$.

The integral is computed by generating samples $\{(u_{k,M-1}, r_{k,M-1}), k = 1, \dots, K\}$ from the joint process $(\tilde{X}_{M-1}^{(M)}, \rho_{M-1}^{(M)})$ using the Euler scheme and then computing

$$p^{(M,K)}(x_t, t; x_s, s, \theta) = \frac{1}{K} \sum_{k=1}^K p^{(1)}(x_t, t; u_{k,M-1}, \tau_{M-1}) r_{k,M-1}.$$

It is easy to show by direct calculation of $\rho_{M-1}^{(M)}$ that this expression is equivalent to (8). We have found it to be more stable to base the Euler scheme for ρ on

$$d(\log \rho) = -\frac{k^2}{2} dt + k d\tilde{W},$$

with initial condition $\log \rho(s) = 0$, rather than (18).

The second importance sampler which we consider draws u_{m+1} from a Gaussian density based on the first-order approximation conditional on u_m and x_t . That is, treating u_m and $u_M = x_t$ as fixed, one draws u_{m+1} from the density

$$\begin{aligned} p(u_{m+1}|u_m, u_M) &= p(u_{m+1}|u_m)p(u_M|u_{m+1})/p(u_M|u_m) \\ &\approx \phi(u_{m+1}; u_m + \bar{\mu}\delta, \bar{\sigma}^2\delta) \\ &\quad \times \phi(u_M; u_{m+1} + \bar{\mu}\delta^*, \bar{\sigma}^2\delta^*) \\ &\quad \div \phi(u_M; u_m + \bar{\mu}\delta^+, \bar{\sigma}^2\delta^+) \\ &= \phi(u_{m+1}; u_m + \tilde{\mu}_m\delta, \tilde{\sigma}_m^2\delta) \end{aligned}$$

where $\delta = (t - s)/M$, $\delta^* = t - \tau_{m+1}$, $\delta^+ = t - \tau_m$, $\bar{\mu} = \mu(u_m)$, $\bar{\sigma} = \sigma(u_m)$, and

$$\tilde{\mu}_m = \left(\frac{u_M - u_m}{t - \tau_m} \right), \quad \tilde{\sigma}_m^2 = \left(\frac{M - m - 1}{M - m} \right) \bar{\sigma}^2.$$

Notice that this importance sampler turns out to be identical to the Brownian bridge sampler except for the factor $(M - m - 1)/(M - m)$ in the variance. While it is not entirely obvious that this should be the case, we shall see that this modification results in much better performance. We shall refer to this sampler as the modified Brownian bridge.

The third importance sampler which we consider was proposed by Elerian et al. (2001). The idea is to approximate the target density by a multivariate normal with mean and variance based on a second-order Taylor expansion of the log target density about the mode.

The log target density is given by

$$\log p^{(M)}(u_1, u_2, \dots, u_{M-1}|x_s, x_t) = \sum_{m=0}^{M-1} \log p^{(1)}(u_{m+1}, \tau_{m+1}; u_m, \tau_m).$$

One samples $u = (u_1, u_2, \dots, u_{M-1})$ from $N(\mu^*, \Sigma^*)$, where

$$\mu^* = \underset{u}{\operatorname{argmax}} \log p(u|x_s, x_t)$$

$$\Sigma^* = - \left[\frac{\partial^2}{\partial u' \partial u} \log p(\mu^*|x_s, x_t) \right]^{-1}.$$

In practice, one obtains μ^* by starting with $\hat{u} = (\hat{u}_1, \dots, \hat{u}_{M-1})$, where $\hat{u}_m = x_s + m(x_t - x_s)/M$, and taking a single Newton step towards the maximum. The derivatives of $\log p(u|x_s, x_t)$ are straightforward but tedious to compute.

The key feature of this sampler is that it draws paths in one shot rather than recursively. Implementing this sampler requires solving a system of $M-1$ linear equations and computing a Cholesky decomposition. To obtain reasonable performance, it is essential that one take advantage of the tridiagonal nature of the relevant matrices.

As always with importance sampling, one should ensure that the tails of the sampling density are not too thin; otherwise, it will not be possible to drive down the variance of the Monte Carlo integral despite using a large number of sample paths. One way to address this problem is by using Student t rather than normal increments in the construction of the sample paths. One might also try the approach suggested by Geweke (1989).

The second category of variance-reduction techniques which we examine is random number schemes. The method of antithetic variates, as discussed in Section 5, is one such scheme, although our results suggests that it provides only marginal benefits in this context.

Recall from equations (16) and (17) that the density approximation may be thought of as an expectation. Kloeden and Platen (1992, Section 14.1) suggest that for computing expectations, the Gaussian increments (W_1, \dots, W_{M-1}) in equation (14) (and similar expressions for the other importance samplers) may be replaced with other random variables satisfying appropriate moment conditions.

One possibility is the random variable which takes on the values 1 and -1 each with probability one-half. In addition to reducing variance, this scheme gives a speed increase, since generating normal deviates can be a significant fraction of the computational effort. Furthermore, if M is sufficiently small it is possible to compute the Monte Carlo integral by summing over all possible branches of the binomial tree of length $M - 1$. For example, setting $M = 8$ would require the computation of $2^7 = 128$ sample paths, which is entirely feasible. Using the techniques discussed in Section 5, it is possible to achieve low bias with small M . In particular, since this random number scheme produces a method with no variance, it is ideally suited for use together with extrapolation.

While expectations computed using the binomial tree scheme are known to converge under appropriate conditions, the properties of this scheme in the context of this paper are uncertain. Therefore, we also consider a related scheme which provides much of the same

benefit by less drastic means. The idea is to control the ‘jaggedness’ of the sample paths by forcing each vector of increments (W_1, \dots, W_{M-1}) to have sample variance one. This may be accomplished simply by using the vector

$$(\tilde{W}_1, \dots, \tilde{W}_{M-1}) = \left\{ \frac{1}{M-1} \sum_{m=1}^{M-1} W_m^2 \right\}^{1/2} (W_1, \dots, W_{M-1}),$$

and may be thought of as a weakening of the two-point idea, which forces the sample variance of each individual increment to be equal to one.

7 Numerical experiments

We first test our various techniques by approximating the transition density as described in Section 3. The settings used for the various approximations will be identified by “sampler-subdensity-M-K”; e.g., the Brownian bridge sampler used together with the first-order (Euler) subtransition density, $M = 8$, and $K = 256$ will be identified as “bridge-euler-8-256”. The RMSE is also computed for these approximation schemes (as described in Section 3). The results are summarized in Table 1.

Figure 5 illustrates the performance of the various subdensity methods when used to compute the transition density directly (i.e., $M = 1$, no simulation). While the error associated with the simple first-order approximation is moderately severe, a factor of 10 improvement is obtained if the model is transformed before applying the first-order approximation. The scheme proposed by Elerian (1998) comes close to obtaining this improvement without needing the transformation step. When applied to the transformed models, the techniques proposed by Shoji and Ozaki (1998) and Kessler (1997) provide an additional order of magnitude improvement over the Euler scheme. Although not shown in the figures, we have found these schemes to be of little benefit when used on the untransformed model. The technique proposed by Nowman (1997) provides nearly no benefit whatsoever.

Figure 6 illustrates the Brownian bridge and modified bridge samplers. The first-order approximation is used for the subtransition densities. The transformation step is not used. Using the Brownian bridge largely solves the main problem associated with Pedersen’s method. The modified bridge provides a further dramatic reduction in variance. Notice that panels (c) and (d) of Figure 6 use only $K = 8$ sample paths, as compared to $K = 256$ for panels (a) and (b) and Figure 2(a)–(d) (Pedersen’s method). We see that increasing the number of subintervals brings the expected reduction in bias.

Figure 7(a) and (b) illustrate the use of extrapolation and Elerian’s subtransition density scheme respectively to reduce bias. Panel (c) shows the variance reduction due to normalized variates. Panel (d) demonstrates that one still obtains the expected reductions in bias and variance from increasing M and K respectively. All panels in this figure use the untransformed model.

Figure 8 illustrates the Elerian-Chib-Shephard (ECS) sampler. This sampler seems to work well for M relatively small, but the variance goes up dramatically as the number of intermediate points increases. It was not possible to compute the RMSE of the log density approximation with the ECS sampler and $M = 32$, because the sampler often chose points below zero (i.e., outside the range of the model). We follow Elerian et al. (2001) by using the transformation $Y = \log X$ in this case.

While it is possible to obtain reasonably accurate results using the untransformed model,

Figure 9 shows that first transforming the model provides significant benefits, especially when used with the subdensity scheme of Shoji and Ozaki (1998) and normalized variates. With these settings we are able to obtain $\text{RMSE} \approx 0.0006$ with only $M = 8$ and $K = 8$. Computational cost is around 16 seconds to approximate a likelihood with $n = 100,000$ observations using fortran code on a PC running at 750 MHz. It would be virtually impossible to obtain anywhere near this level of accuracy using Pedersen’s method without our acceleration techniques.

Tables 2 and 3 show the errors which result from estimating parameters by maximizing the approximate rather than exact log likelihood. Results are shown for several different settings of the model parameters. The errors are estimated by Monte Carlo simulation with 512 repetitions over synthetic data sets of $n = 1000$ observations. The SMLE estimates are obtained using the modified bridge sampler, Shoji and Ozaki’s subdensity, $M = 16$ and $K = 16$. For comparison, we also compute parameter estimates using the first-order Euler scheme approximation. The transformed model is used for all the experiments shown in these tables. It should be noted that the Euler scheme approximations thus obtained can be expected to be significantly better than those typically obtained by practitioners without implementing the transformation (see Figure 5).

Table 2(a) uses the baseline model settings, $\theta^o = (.06, 0.5, 0.15)$ and $\Delta = 1/12$. Recall that these are calibrated to monthly observations of the U.S. short-term interest rate. For these model settings we also compute parameter estimates using Pedersen’s method with $M = 8$ and $K = 256$. Pedersen’s method is unable to match even the Euler scheme. On the other hand, the approximation error of the SMLE estimates obtained using our techniques are negligible (compare with the sample distribution of $\hat{\theta}_{\text{MLE}} - \theta^o$).

Panel (b) increases the volatility of the model. Note that computing the exact likelihood requires the evaluation of a Bessel function, which in turn requires $2\theta_2\theta_1/\theta_3^2 \geq 1$. Setting $\theta^o = (.06, 0.5, 0.22)$ comes quite close to this boundary. For some samples, the constraint appears to bind when maximizing the likelihood. These samples are discarded. This model causes our methodology some difficulty, apparently because the data often ventures close to the singularity at zero. The estimates are nonetheless quite good.

Panel (c) reduces the model’s volatility to $\theta_3 = 0.03$. Panel (d) sets the mean reversion parameter to $\theta_2 = 0.4$. Again, the constraint $2\theta_2\theta_1/\theta_3^2 \geq 1$ comes into play. Neither of these tests present any difficulty to our methodology.

Panel (e) increases the mean reversion parameter by a factor of 10 to $\theta_2 = 5.0$. Panel (e) uses the baseline setting for θ^o , but stretches the sampling interval to 2 years. Both of these settings result in large biases for the first-order approximation, but pose little difficulty for the SMLE technique.

8 Stochastic volatility

While the previous sections have focused on techniques designed to efficiently approximate the likelihood function for scalar models, most of these ideas are easily generalized to the multivariate setting. With some work, they may also be applied to latent variable models. The short-term interest rate and many other financial time series are well-known to exhibit properties such as fat-tails and persistent volatility patterns which are inconsistent with time-homogeneous scalar models (e.g., Ghysels, Harvey and Renault 1996). A variety of alternative models have been proposed. To illustrate our methodology, we shall examine stochastic volatility (SV) models of the form

$$\begin{aligned}dX &= \mu_X(X) dt + \sigma_X(X) \exp(H) dW_1 \\dH &= \mu_H(X, H) dt + \sigma_H(X, H) dW_2.\end{aligned}$$

Such models have been examined by Gallant and Tauchen (1998), Andersen and Lund (1997), and Eraker (2001) among others. The second component, H , corresponds to an unobserved volatility factor.

In order to obtain a likelihood, the unobserved factor must be integrated out. Several ways of going about this have been proposed for discrete-time models, e.g., Danielsson and Richard (1993), Jacquier, Polson and Rossi (1994), Richard and Zhang (2000), Sandmann and Koopman (1998), Kim, Shephard and Chib (1998), and Pitt and Shephard (1999). In the continuous-time context, it is less straightforward to integrate out the unobserved factor, and alternative approaches are used. The Efficient Method of Moments approach has been used by Gallant and Tauchen (1998) and others. Methods based on the empirical characteristic function have been proposed by Chacko and Viceira (1998) and Singleton (1997). Markov chain Monte Carlo approaches have been proposed by Eraker (2001), Jones (1999b) and Elerian (1999).

We now outline a technique to approximate the likelihood of continuous time SV models. For simplicity, we shall assume that W_1 and W_2 are independent; this is not an essential part of the methodology. We provide a small Monte Carlo study which demonstrates that the procedure is effective and reasonably fast. The methodology is used to estimate a SV model of the U.S. short-term interest rate in Section 9. Further refinements are undoubtedly possible; a more detailed study is currently underway.

The basic idea is relatively straightforward. We are interested in the process $(X(t), H(t))$, where X is observed at times t_0, t_1, \dots, t_n and H is latent. Let $X_i = X(t_i)$, $H_i = H(t_i)$, and $\mathcal{F}_i = \sigma(H_0, X_0, X_1, \dots, X_i)$ for $i = 1, \dots, n$. The goal is to obtain $p(X_{i+1}|\mathcal{F}_i)$. If we knew the distribution of $H_i|\mathcal{F}_i$, we could use

$$p(X_{i+1}|\mathcal{F}_i) = \int p(X_{i+1}|X_i, h_i) dP_{H_i|\mathcal{F}_i}(h_i).$$

We will approximate $H_i|\mathcal{F}_i$. Given the distribution of $H_{i-1}|\mathcal{F}_{i-1}$, it can be propagated forward using

$$p(H_{i-1}|\mathcal{F}_i) = \frac{p(X_i|\mathcal{F}_{i-1}, H_{i-1})p(H_{i-1}|\mathcal{F}_{i-1})}{p(X_i|\mathcal{F}_{i-1})}$$

$$p(H_i|\mathcal{F}_i) = \int p(H_i|X_i, h_{i-1}) dP_{H_{i-1}|\mathcal{F}_i}(h_{i-1}).$$

It remains only to find H_0 . The approach we take is to estimate H_0 as an unknown parameter, although one could equally well integrate it out using an appropriate importance sampler. This is the basic idea of a particle filter (see e.g., Pitt and Shephard 1999 and the references therein).

To implement this idea, we use the following procedure. Consider the model

$$dZ = \mu(Z) dt + \sigma(Z) dW$$

where

$$Z = \begin{pmatrix} X \\ H \end{pmatrix}, \quad \mu = \begin{pmatrix} \mu_X(Z) \\ \mu_H(Z) \end{pmatrix}, \quad \sigma = \begin{pmatrix} \sigma_X(Z) & 0 \\ 0 & \sigma_H(Z) \end{pmatrix}, \quad W = \begin{pmatrix} W_1 \\ W_2 \end{pmatrix}.$$

The problem is to determine the density $p(X_{i+1}|\mathcal{F}_i, \theta)$, where $\mathcal{F}_i = \sigma(H_0, X_0, X_1, \dots, X_i)$.

The first-order approximation is given by

$$p^{(1)}(z_t, t; z_s, s, \theta) = \phi(z_t; z_s + \mu(z_s)(t - s), \Sigma(z_s)(t - s)),$$

where ϕ is the multivariate Gaussian density and $\Sigma = \sigma\sigma^T$. Also, let $s = \tau_0 < \tau_1 < \dots < \tau_M = t$ be a partition of the interval $[s, t]$, and let

$$p^{(M)}(z_t, t; z_s, s, \theta) = \int \prod_{m=1}^M p^{(1)}(w_m, \tau_m; w_{m-1}, \tau_{m-1}, \theta) d\lambda(w_1, \dots, w_{M-1})$$

where $w_m \in \mathbb{R}^2$ for $m = 0, \dots, M$, λ is Lebesgue measure in $\mathbb{R}^{2(M-1)}$, and we use the convention $w_0 = z_s$ and $w_M = z_t$. Since H is unobserved, it must be integrated out as well. One obtains

$$p^{(M)}(X_{i+1}|\mathcal{F}_i, \theta) = \int p^{(M)}\left(\begin{pmatrix} X_{i+1} \\ h_{i+1} \end{pmatrix}, t_{i+1}; \begin{pmatrix} X_i \\ h_i \end{pmatrix}, t_i, \theta\right) dP_{H_i, H_{i+1}|\mathcal{F}_i}(h_i, h_{i+1}).$$

As in the scalar case, the integrals are evaluated using Monte Carlo integration. Let $q\left(v_0, \begin{pmatrix} u_1 \\ v_1 \end{pmatrix}, \begin{pmatrix} u_2 \\ v_2 \end{pmatrix}, \dots, \begin{pmatrix} u_{M-1} \\ v_{M-1} \end{pmatrix}, v_M\right)$ be an importance sampler on \mathbb{R}^{2M} , and let

$$\mathbf{v}_k = \left(v_{k,0}, \begin{pmatrix} u_{k,1} \\ v_{k,1} \end{pmatrix}, \begin{pmatrix} u_{k,2} \\ v_{k,2} \end{pmatrix}, \dots, \begin{pmatrix} u_{k,M-1} \\ v_{k,M-1} \end{pmatrix}, v_{k,M}\right), \quad k = 1, 2, \dots, K$$

be draws from q . Then one defines

$$p^{(M,K)}(X_{i+1}|\mathcal{F}_i, \theta) = \frac{1}{K} \sum_{k=1}^K \frac{\prod_{m=1}^M p^{(1)}\left(\begin{pmatrix} u_{k,m} \\ v_{k,m} \end{pmatrix}, \tau_m; \begin{pmatrix} u_{k,m-1} \\ v_{k,m-1} \end{pmatrix}, \tau_{m-1}, \theta\right)}{q\left(v_{k,0}, \begin{pmatrix} u_{k,1} \\ v_{k,1} \end{pmatrix}, \begin{pmatrix} u_{k,2} \\ v_{k,2} \end{pmatrix}, \dots, \begin{pmatrix} u_{k,M-1} \\ v_{k,M-1} \end{pmatrix}, v_{k,M}\right)}$$

with the convention $u_{k,0} = X_i$ and $u_{k,M} = X_{i+1}$ for all k .

The theoretical framework is essentially unchanged from the scalar case. In particular, sufficient conditions for

$$\lim_{M \rightarrow \infty} p^{(M)}(\cdot, t; z_s, s, \theta) = p(\cdot, t; z_s, s, \theta) \quad \text{in } L^1(\lambda)$$

may be found in Pedersen (1995b). For a fixed realization $(x_0, x_1, \dots, x_{i+1})$, one obtains (with standard regularity conditions)

$$\lim_{K \rightarrow \infty} p^{(M,K)}(x_{i+1}; \mathcal{F}_i, \theta) \rightarrow p^{(M)}(x_{i+1}; \mathcal{F}_i, \theta)$$

from the strong law of large numbers in the usual way.

The issue of how to construct the importance sampler remains to be addressed. The approach taken in this paper is as follows. Suppose that the distribution of $H_i | \mathcal{F}_i$ is known. We need a way to sample

$$\mathbf{w} = \left(v_0, \begin{pmatrix} u_1 \\ v_1 \end{pmatrix}, \begin{pmatrix} u_2 \\ v_2 \end{pmatrix}, \dots, \begin{pmatrix} u_{M-1} \\ v_{M-1} \end{pmatrix}, v_M \right).$$

First, draw v_0 from $H_i | \mathcal{F}_i$ and let $u_0 = X_i$. Now proceed recursively for $m = 0, \dots, M-2$: given (u_m, v_m) , draw u_{m+1} using the modified Brownian bridge sampler described in Section 6,

$$\begin{aligned} u_{m+1} &\sim N(u_m + \tilde{\mu}_{X,m} \delta, \tilde{\sigma}_{X,m}^2 \delta) \\ \tilde{\mu}_{X,m} &= \left(\frac{X_{i+1} - u_m}{t_{i+1} - \tau_m} \right) \\ \tilde{\sigma}_{X,m}^2 &= \sigma_X^2(u_m, v_m) \left(\frac{M - m - 1}{M - m} \right), \end{aligned}$$

and draw v_{m+1} ‘blindly’,

$$\begin{aligned} v_{m+1} &\sim N(v_m + \tilde{\mu}_{H,m} \delta, \tilde{\sigma}_{H,m}^2 \delta) \\ \tilde{\mu}_{H,m} &= \mu_H(u_m, v_m) \\ \tilde{\sigma}_{H,m}^2 &= \sigma_H^2(u_m, v_m), \end{aligned}$$

where $\delta = (t_{i+1} - t_i)/M$. And finally, draw v_M using the same procedure as for $m = 1, \dots, M-1$.

Now we need to propagate the distribution of $H_i | \mathcal{F}_i$ forward. Consider the K draws $\{v_{k,M}, k = 1, \dots, K\}$ from the sampler, together with the corresponding weights

$$\lambda_k = \frac{\prod_{m=1}^M p^{(1)} \left(\begin{pmatrix} u_{k,m} \\ v_{k,m} \end{pmatrix}, \tau_m; \begin{pmatrix} u_{k,m-1} \\ v_{k,m-1} \end{pmatrix}, \tau_{m-1}, \theta \right)}{q \left(v_{k,0}, \begin{pmatrix} u_{k,1} \\ v_{k,1} \end{pmatrix}, \begin{pmatrix} u_{k,2} \\ v_{k,2} \end{pmatrix}, \dots, \begin{pmatrix} u_{k,M-1} \\ v_{k,M-1} \end{pmatrix}, v_{k,M} \right)}.$$

Let $\{\bar{\lambda}_k, k = 1, \dots, K\}$ denote the weights after normalizing so that they sum to one. The collection $\{(v_{k,M}, \bar{\lambda}_k), k = 1, \dots, K\}$ of points and weights may be thought of as representing a discrete approximation to the density $H_{i+1}|\mathcal{F}_{i+1}$.

Various techniques are available to form a distribution function from this approximation and draw points from it. While one could simply use the approximation

$$P_{H_{i+1}|\mathcal{F}_{i+1}}(h|\mathcal{F}_{i+1}) \approx \sum_{k: v_{k,M} \leq h} \bar{\lambda}_k,$$

the resulting likelihood function is discontinuous in θ and will cause difficulties for the optimizer. Therefore, we use a Hermite function to approximate the density and draw points from it instead.

We repeat the Monte Carlo experiment of Eraker (2001) using the methodology described above. Synthetic data sets of length $n = 500$ are generated from the model

$$\begin{aligned} dX &= (\theta_1 + \theta_2 X) dt + \exp(H/2) X^{1/2} dW_1 \\ dH &= (\theta_3 + \theta_4 H) dt + \theta_5 dW_2 \end{aligned} \tag{19}$$

with parameter vector $\theta^\circ = (0.0002, -0.002, -0.3, -0.03, 0.3)$. This model is said to be calibrated to the U.S. short-term interest rate with time measured in days. Data is generated using the Euler scheme with 50 subintervals. Estimates are obtained using $K = 128$ sample paths and $M = 8$ subintervals. Little is lost by reducing the number of sample paths to $K = 64$. We do not recommend using less than this. Increasing the number of subintervals does not improve the precision of the estimates significantly. The mean, standard error, and RMSE of parameter estimates over 512 Monte Carlo repetitions are shown in Table 4. For comparison, we reproduce Eraker's results as well. Our results are very similar. Computational cost is around 1.2 seconds per evaluation of the likelihood function on a 750 MHz PC.

Note that the estimates of θ_1 and θ_3 appear to be quite imprecise. This is due to the parametric form of the model. In contrast, the estimates of θ_1/θ_2 and θ_3/θ_4 are quite good. While we have maintained the model in the form used by Eraker, it might be better to instead use

$$\begin{aligned} dX &= \theta_2(X - \mu) dt + \sigma \exp(H) X^{1/2} dW_1 \\ dH &= \theta_4 H dt + \theta_5 dW_2, \end{aligned}$$

where $\mu = -\theta_1/\theta_2$ and $\sigma = \exp(-\theta_3/\theta_4)$.

Note that our methodology can be easily used to estimate discrete time SV models by setting $M = 1$. We have tested the methodology using some of the models examined by Jacquier et al. (1994) and others with similar results.

9 Application

To illustrate the methodologies proposed in this paper, we estimate some simple models of the short-term interest rate. We use the data set previously examined by Gallant and Tauchen (1998), which consists of 1809 weekly observations of the three-month Treasury bill rate (January 5, 1962 – August 30, 1996). Rates are annualized and quoted on a discount basis. The data are plotted in Figure 10.

We first fit a simple scalar model, $dX = (\theta_1 + \theta_2 X) dt + \theta_3 X^{\theta_4} dW$, commonly referred to as the constant elasticity of volatility (CEV) model. It has been studied previously (using other estimators and data) by Chan et al. (1992), Tauchen (1995), Aït-Sahalia (1996), Conley, Hansen, Luttmer and Scheinkman (1997), and others.

We also fit the stochastic volatility (SV) model given by

$$dX = (\theta_1 + \theta_2 X) dt + \theta_3 X^{\theta_4} e^H dW_1 \quad (20)$$

$$dH = \theta_5 H dt + \theta_6 dW_2, \quad (21)$$

with W_1 and W_2 independent. Similar models have been examined by Gallant and Tauchen (1998), Andersen and Lund (1997), and Eraker (2001). Time is measured in years.

Maximum likelihood estimates and log likelihoods are given in Table 5. For the scalar model, we have used $M=16$ and $K=16$ with the modified bridge sampler, transformed model, normalized variance random scheme, and Shoji and Ozaki's subdensity. For the SV model, we use $M=8$ and $K=256$ with the techniques described in the preceding section. Since the scalar model is nested within the SV model, the restriction can be tested using, for example, the likelihood ratio statistic. Although setting $\theta_6 = 0$ causes θ_5 to become unidentified, this issue may be addressed along the lines of e.g. Gallant (1977) or Andrews and Ploberger (1994). In any event, the SV model results in a huge improvement in the log likelihood. The scalar model does not appear to be plausible.

Notice that the estimates for θ_1 and θ_2 are insignificantly different from zero in both models. In a more exhaustive study, Durham (2000) finds that the constant term in the drift is needed to avoid having an attracting boundary at zero, but that the benefits of including additional terms are negligible. This is true for both the scalar and stochastic volatility models.

Notice also that the estimates for θ_3 and θ_4 are similar for both models (the unconditional mean of H is zero, and thus $\exp(H)$ vacillates around one). The estimates are slightly lower in the SV model since this model generates conditional densities with thicker tails, which helps to catch the extreme events. These parameters are estimated quite precisely in the scalar model. Including the unobserved component reduces the precision of these estimates, but not to an unacceptable degree. There is a great deal of covariance between $\hat{\theta}_4$ and $\hat{\theta}_5$

in the SV model. Volatility tends to be high when interest rates are high; the estimator has difficulty distinguishing whether this is due to persistence or a level effect. Much of the lack of precision in these estimates appears to be due to this.

The volatility component is an Ornstein-Uhlenbeck process. Its unconditional distribution is Gaussian with mean zero and variance $\theta_6^2/(-2\theta_5) \approx 0.34$. The mean reversion parameter is around -4, which corresponds to a half-life of about 2 months.

10 Conclusions

Despite the theoretical advantages of maximum likelihood estimation, it has been seldom used in estimating continuous-time diffusion models. The transition densities are not known for most models of interest, and previously available approximation techniques have either been of questionable accuracy or computationally intensive.

The simulation-based approach suggested by Pedersen (1995b) and Santa-Clara (1995) is appealing from a theoretical and intuitive viewpoint; however, we find that it can be prohibitively costly to attain even the accuracy of the simple first-order approximation. Eraker (2001), Jones (1999a), and Elerian et al. (2001) propose interesting MCMC approaches to estimating diffusion models. Elerian et al. suggest the idea of improved importance sampling to accelerate the convergence of the Monte Carlo integration at the heart of the simulated likelihood approach. However, computational cost remains high.

We build upon this work, examining other importance samplers, alternate random schemes, higher-order subtransition densities, extrapolation, and a variance-stabilizing transformation of the model. Combining these ideas results in highly efficient approximations. When applied to synthetic data ($n = 1000$) generated by a CIR model with parameters calibrated to match monthly observations of the U.S. short-term interest rate, we are able to obtain maximum likelihood estimates with negligible approximation error in well under one minute. The log likelihood function itself can be approximated with great accuracy in about 0.1 second.

Our results suggest that the best performance is obtained using the modified bridge sampler with the subtransition density of Shoji and Ozaki applied to the transformed model. The number of subintervals, M , and sample paths, K , must be determined by experimentation.

Future work will undoubtedly uncover further refinements in these techniques, as well as point to situations where one or the other particular variant may be preferred. While the main focus of this paper is on scalar models, an approach to applying some of the ideas to estimate a two-factor, latent variable model is also proposed. Extending these techniques to models with jump components would also be of interest.

References

- Aït-Sahalia, Y. (1996). Testing continuous-time models of the spot interest rate, *Review of Financial Studies* **9**(2): 385–426.
- Aït-Sahalia, Y. (2001). Maximum likelihood estimation of discretely sampled diffusions: a closed-form approach, *Econometrica* (forthcoming).
- Andersen, T. G. and Lund, J. (1997). Estimating continuous-time stochastic volatility models of the short-term interest rate, *Journal of Econometrics* **77**: 343–377.
- Andrews, D. W. K. and Ploberger, W. (1994). Optimal tests when a nuisance parameter is present only under the alternative, *Econometrica* **62**(6): 1383–1414.
- Bibby, B. M. and Sørensen, M. (1995). Martingale estimating functions for discretely observed diffusion processes, *Bernoulli* **1**: 17–39.
- Brandt, M. W. and Santa-Clara, P. (2001). Simulated likelihood estimation of diffusions with an application to exchange rate dynamics in incomplete markets, *Journal of Financial Economics* (forthcoming).
- Chacko, G. and Viceira, L. M. (1998). Spectral GMM estimation of continuous-time processes, Working Paper, Harvard Business School.
- 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, *Journal of Finance* **47**: 1209–1228.
- Conley, T. G., Hansen, L. P., Luttmer, E. G. J. and Scheinkman, J. A. (1997). Short-term interest rates as subordinated diffusions, *Review of Financial Studies* **10**(3): 525–577.
- Cox, J. C., Ingersoll, J. E. and Ross, S. A. (1985). A theory of the term structure of interest rates, *Econometrica* **53**(2): 385–407.
- Danielsson, J. and Richard, J.-F. (1993). Accelerated Gaussian importance sampler with application to dynamic latent variable models, *Journal of Applied Econometrics* **8**: 153–173.
- Doss, H. (1977). Liens entre équations différentielles stochastiques et ordinaires, *Ann. Inst. H. Poincaré* **13**: 99–125.
- Duffie, D. and Glynn, P. (1996). Estimation of continuous-time Markov processes sampled at random time intervals, Working Paper, Graduate School of Business, Stanford University.

- Duffie, D. and Singleton, K. J. (1993). Simulated moments estimation of Markov models of asset prices, *Econometrica* **61**(4): 929–952.
- Durham, G. B. (2000). Specification analysis of the short-term interest rate using likelihood-based techniques, Working Paper, Department of Economics, University of North Carolina.
- Elerian, O. (1998). A note on the existence of a closed form conditional transition density for the Milstein scheme, Working Paper, Nuffield College, Oxford University.
- Elerian, O. (1999). *Simulation Estimation of Continuous-Time Models with Applications to Finance*, Ph.d. dissertation, Nuffield College, University of Oxford.
- Elerian, O., Chib, S. and Shephard, N. (2001). Likelihood inference for discretely observed non-linear diffusions, *Econometrica* (forthcoming).
- Eraker, B. (2001). MCMC analysis of diffusion models with application to finance, *Journal of Business and Economic Statistics* **19**(2): 177–191.
- Florens-Zmirou, D. (1989). Approximate discrete-time schemes for statistics of diffusion processes, *Statistics* **20**: 547–557.
- Gallant, A. R. (1977). Testing a nonlinear regression specification: A nonregular case, *Journal of the American Statistical Association* **72**(359): 523–530.
- Gallant, A. R. and Tauchen, G. (1997). Estimation of continuous-time models for stock returns and interest rates, *Macroeconomic Dynamics* **1**: 135–168.
- Gallant, A. R. and Tauchen, G. (1998). Reprojecting partially observed systems with application to interest rate diffusions, *Journal of the American Statistical Association* **93**(441): 10–24.
- Gallant, A. R. and Tauchen, G. E. (1996). Which moments to match?, *Econometric Theory* **12**: 657–681.
- Geweke, J. (1989). Bayesian inference in econometric models using Monte Carlo integration, *Econometrica* **57**(6): 1317–1339.
- Ghysels, E., Harvey, A. and Renault, E. (1996). Stochastic volatility, in *Handbook of Statistics 14, Statistical Methods in Finance*, G.S. Maddala and C.R. Rao (eds.), North Holland, Amsterdam.
- Gouriéroux, C., Monfort, A. and Renault, E. (1993). Indirect inference, *Journal of Applied Econometrics* **8**: S85–S118.

- Hansen, L. P. and Scheinkman, J. A. (1995). Back to the future: Generating moment implications for continuous-time markov processes, *Econometrica* **63**: 767–804.
- Jacquier, E., Polson, N. G. and Rossi, P. E. (1994). Bayesian analysis of stochastic volatility models, *Journal of Business and Economic Statistics* **12**(4): 371–389.
- Jones, C. S. (1999a). Bayesian estimation of continuous-time finance models, Working Paper, Simon School of Business, University of Rochester.
- Jones, C. S. (1999b). The dynamics of stochastic volatility, Working Paper, (Nov 1999 version), Simon School of Business, University of Rochester.
- Karatzas, I. and Shreve, S. E. (1991). *Brownian Motion and Stochastic Calculus*, 2nd edn, Springer, New York.
- Kessler, M. (1997). Estimation of an ergodic diffusion from discrete observations, *Scandinavian Journal of Statistics* **24**: 211–229.
- Kim, S., Shephard, N. and Chib, S. (1998). Stochastic volatility: Likelihood inference and comparison with ARCH models, *Review of Economic Studies* **65**: 361–393.
- Kloeden and Platen (1992). *Numerical Solution of Stochastic Differential Equations*, Springer Verlag, Berlin.
- Lo, A. W. (1988). Maximum likelihood estimation of generalized ito processes with discretely sampled data, *Econometric Theory* **4**: 231–247.
- Milstein, G. (1978). A method of second order accuracy integration of stochastic differential equations, *Theory Probab. Appl.* **23**: 396–401.
- Nowman, K. (1997). Gaussian estimation of single-factor continuous time models of the term structure of interest rates, *Journal of Finance* **52**(4): 1695–1706.
- Pedersen, A. R. (1995a). Consistency and asymptotic normality of an approximate maximum likelihood estimator for discretely observed diffusion processes, *Bernoulli* **1**(3): 257–279.
- Pedersen, A. R. (1995b). A new approach to maximum likelihood estimation for stochastic differential equations based on discrete observations, *Scandinavian Journal of Statistics* **22**: 55–71.
- Pitt, M. K. and Shephard, N. (1999). Filtering via simulation: Auxiliary particle filter, *Journal of the American Statistical Association* **446**(94): 590–599.

- Richard, J.-F. and Zhang, W. (2000). Accelerated Monte Carlo integration: an application to dynamic latent variables models, *in* R. S. Mariano, T. Schuermann and M. Weeks (eds), *Simulation-Based Inference in Economics*, Cambridge University Press, Cambridge, UK, chapter 2, pp. 47–70.
- Sandmann, G. and Koopman, S. J. (1998). Estimation of stochastic volatility models via Monte Carlo maximum likelihood, *Journal of Econometrics* **87**: 271–301.
- Santa-Clara, P. (1995). Simulated likelihood estimation of diffusions with an application to the short-term interest rate, Working Paper, Anderson Graduate School of Management, UCLA.
- Shoji, I. and Ozaki, T. (1998). Estimation for nonlinear stochastic differential equations by a local linearization method, *Stochastic Analysis and Applications* **16**(4): 733–752.
- Singleton, K. J. (1997). Estimation of affine asset pricing models using the empirical characteristic function, Working paper (Aug 2000 version), Stanford University.
- Tauchen, G. E. (1995). New minimum chi-square methods in empirical finance, *in* K. Wallace and D. Kreps (eds), *Advances in Econometrics*, Cambridge University Press, Cambridge, UK, chapter 9, pp. 279–317.

Table 1: RMSE of log density approximations.

Untransformed model							
Sampler	subdensity	M	K	NV ^a	extrap ^b	RMSE	time ^c
none	euler	1	1			0.13678	0.2
none	elerian	1	1			0.03550	0.2
none	nowman	1	1			0.14467	0.2
pedersen	euler	8	256			0.19353	539.5
pedersen	euler	8	256			0.19353	538.3
pedersen	euler	32	256			0.66310	2169.8
bridge	euler	8	256			0.05355	550.4
bridge	euler	32	256			0.05570	2211.2
mod bridge	euler	8	8			0.03068	17.5
mod bridge	euler	8	8		x	0.06103	26.3
mod bridge	euler	8	8	x		0.02404	19.1
mod bridge	euler	32	32	x		0.01180	299.3
mod bridge	euler	32	128	x		0.00713	1227.1
mod bridge	elerian	8	8			0.03562	24.6
mod bridge	elerian	8	8	x		0.01856	25.0
mod bridge	elerian	32	32	x		0.01206	401.2
mod bridge	elerian	32	128	x		0.00656	1603.8
ecs	elerian	8	8			0.07207	30.4
ecs	elerian	8	8	x		0.05164	30.3
ecs ^d	elerian	32	32	x		0.46920	403.6
ecs ^d	elerian	32	128	x		0.25973	1582.2
Transformed model							
Sampler	subdensity	M	K	NV ^a	extrap ^b	RMSE	time ^c
none	euler	1	1			0.03790	0.2
none	shoji	1	1			0.01412	0.3
none	kessler	1	1			0.00864	0.2
pedersen	euler	8	256			0.16500	338.2
pedersen	euler	32	256			0.57507	1254.5
mod bridge	euler	8	8			0.00812	10.8
mod bridge	euler	64	8			0.00255	80.7
mod bridge	euler	8	8		x	0.01795	16.6
mod bridge	shoji	8	8			0.00070	16.0
mod bridge	shoji	8	8	x		0.00057	16.4
mod bridge	shoji	32	32	x		0.00030	248.7
ecs	shoji	8	8	x		0.00072	25.3
ecs	shoji	32	32	x		0.00020	315.5

^aAn 'x' in this column indicates normalized variates were used; otherwise, antithetic variates.

^bAn 'x' in this column indicates extrapolation was used.

^cComputing time (in seconds) required to obtain likelihood for $n = 100,000$ observations using fortran code on a 750 MHz PC.

^dWith these settings, the sampler frequently chose points $x_\tau < 0$, thus, following Elerian et al., we used the transformation $Y = \log X$.

Table 2: Approximation errors for parameter estimates.

Results of Monte Carlo study assessing the quality of parameter estimates obtained using various techniques. Each experiment uses 512 replications, each over synthetic data sets of $n = 1000$ observations. The goal is for the distance from the approximations to the exact MLE to be a small fraction of the distance from the exact MLE to the data generating parameter.

(a)		$\theta^o = (.06, 0.5, 0.15), \Delta = 1/12, \text{df}^a=5.33$				
		θ_1	θ_2	θ_3	Log L	
mle – true:	mean	0.00061	0.04890	0.00020	1.54137	
	std err	0.00796	0.12532	0.00340	1.24511	
	rmse	0.00797	0.13442	0.00341	1.98074	
mle – euler:	mean	0.00026	0.01781	0.00374	0.19031	
	std err	0.00003	0.01103	0.00076	0.54499	
	rmse	0.00026	0.02095	0.00382	0.57700	
mle – smle ^b : (pedersen)	mean	-0.00062	-0.00801	-0.00223	8.07971	
	std err	0.00256	0.03681	0.00175	4.76355	
	rmse	0.00264	0.03765	0.00284	9.37802	
mle – smle ^c :	mean	0.00000	-0.00005	-0.00000	-0.00469	
	std err	0.00000	0.00073	0.00001	0.01340	
	rmse	0.00000	0.00073	0.00001	0.01419	
(b)		$\theta^o = (.06, 0.5, 0.22), \Delta = 1/12, \text{df}^a=2.48$				
		θ_1	θ_2	θ_3	Log L	
mle – true:	mean	-0.00031	0.05973	0.00012	1.55119	
	std err	0.01136	0.12808	0.00535	1.25344	
	rmse	0.01136	0.14126	0.00535	1.99386	
mle – euler:	mean	0.00069	0.04941	0.00701	3.01730	
	std err	0.00014	0.03111	0.00136	1.90071	
	rmse	0.00070	0.05838	0.00714	3.56548	
mle – smle ^c :	mean	0.00001	-0.00357	0.00018	0.03479	
	std err	0.00014	0.01678	0.00119	0.93214	
	rmse	0.00014	0.01714	0.00120	0.93187	
(c)		$\theta^o = (.06, 0.5, 0.03), \Delta = 1/12, \text{df}^a=133.33$				
		θ_1	θ_2	θ_3	Log L	
mle – true:	mean	0.00005	0.04378	0.00001	1.57035	
	std err	0.00162	0.12239	0.00070	1.29601	
	rmse	0.00162	0.12992	0.00070	2.03567	
mle – euler:	mean	0.00001	0.01287	0.00067	0.00428	
	std err	0.00000	0.00595	0.00015	0.05930	
	rmse	0.00001	0.01418	0.00069	0.05942	
mle – smle ^c :	mean	0.00000	-0.00011	-0.00000	-0.00227	
	std err	0.00000	0.00005	0.00000	0.00107	
	rmse	0.00000	0.00013	0.00000	0.00251	

^aDegrees of freedom of the exact noncentral chi-square transition density

^bSampler=Pedersen, subdensity=Euler, M=8, K=256, transformed model.

^cSampler=modified bridge, subdensity=Shoji & Ozaki, M=16, K=16, normalized variates, transformed model.

Table 3: Approximation errors for parameter estimates, continued.

		$\theta^o = (.06, 0.4, 0.15), \Delta = 1/12, df^a=4.27$				
		θ_1	θ_2	θ_3	Log L	
(d)	mle – true:	mean	0.00077	0.04589	0.00008	1.55436
		std err	0.01019	0.11201	0.00347	1.28351
		rmse	0.01022	0.12099	0.00347	2.01539
(d)	mle – euler:	mean	0.00027	0.01459	0.00320	0.27005
		std err	0.00005	0.01066	0.00072	0.60164
		rmse	0.00027	0.01807	0.00328	0.65920
(d)	mle – smle ^c :	mean	-0.00000	-0.00011	-0.00001	-0.01275
		std err	0.00002	0.00109	0.00004	0.03611
		rmse	0.00002	0.00109	0.00004	0.03827
		$\theta^o = (.06, 5.0, 0.15), \Delta = 1/12, df^a=53.33$				
		θ_1	θ_2	θ_3	Log L	
(e)	mle – true:	mean	0.00007	0.01506	0.00000	1.59353
		std err	0.00082	0.46302	0.00427	1.39283
		rmse	0.00082	0.46300	0.00427	2.11592
(e)	mle – euler:	mean	0.00019	0.93515	0.02666	0.31509
		std err	0.00001	0.16274	0.00257	0.73929
		rmse	0.00019	0.94919	0.02679	0.80325
(e)	mle – smle ^c :	mean	0.00000	-0.00734	-0.00005	-0.14917
		std err	0.00000	0.00423	0.00002	0.04223
		rmse	0.00000	0.00847	0.00005	0.15502
		$\theta^o = (.06, 0.5, 0.15), \Delta = 2, df^a=5.33$				
		θ_1	θ_2	θ_3	Log L	
(f)	mle – true:	mean	0.00036	0.00126	0.00038	1.52724
		std err	0.00173	0.04305	0.00564	1.24582
		rmse	0.00176	0.04302	0.00564	1.97013
(f)	mle – euler:	mean	0.00367	0.21218	0.05275	21.19312
		std err	0.00020	0.02967	0.00458	5.66304
		rmse	0.00367	0.21424	0.05295	21.93523
(f)	mle – smle ^c :	mean	0.00006	0.00712	0.00041	-2.63215
		std err	0.00004	0.00778	0.00091	0.95334
		rmse	0.00007	0.01054	0.00100	2.79916

^aDegrees of freedom of the exact noncentral chi-square transition density

^bSampler=Pedersen, subdensity=Euler, M=8, K=256, transformed model.

^cSampler=modified bridge, subdensity=Shoji & Ozaki, M=16, K=16, normalized variates, transformed model.

Table 4: Monte Carlo study for SV model

Mean and root mean squared error of parameters estimated on synthetic data ($n = 500$) generated from (19). The sampling frequency is $\Delta t = 1$. The parameters are calibrated to match U.S. short-term interest rates with time measured in days. The Monte Carlo experiment was run for 512 repetitions. Monte Carlo results from Eraker (2001) are included for comparison.

	θ_1	θ_2	θ_1/θ_2	θ_3	θ_4	θ_3/θ_4	θ_5
true	0.00020	-0.00200	-0.10000	-0.30000	-0.03000	0.10000	0.30000
SMLE							
	θ_1	θ_2	θ_1/θ_2	θ_3	θ_4	θ_3/θ_4	θ_5
mean	0.00089	-0.00886	-0.10319	-0.43425	-0.04350	0.10040	0.31734
RMSE	0.00102	0.00986	0.05523	0.27931	0.02802	0.00493	0.06888
Eraker (2001)							
	θ_1	θ_2		θ_3	θ_4		θ_5
mean	0.00127	-0.01271		-0.38174	-0.03873		0.24297
RMSE	0.00154	0.01468		0.21844	0.02209		0.07173

Table 5: CEV and SV models fitted to weekly observations of the 3-month Treasury bill rate, Jan 5, 1962–Aug 30, 1996.

Parameter estimates for CEV model and the SV model in equation (20). Standard errors are in parentheses below the parameter estimates. Time is measured in years.

Model	$\log L$	θ_1	θ_2	θ_3	θ_4	θ_5	θ_6
CEV	512.32	0.8277 (0.5004)	-0.1049 (0.1167)	0.1032 (0.0038)	1.4502 (0.0210)		
SV	913.95	0.176 (0.291)	0.019 (0.066)	0.105 (0.028)	1.247 (0.154)	-4.104 (0.747)	1.682 (0.144)

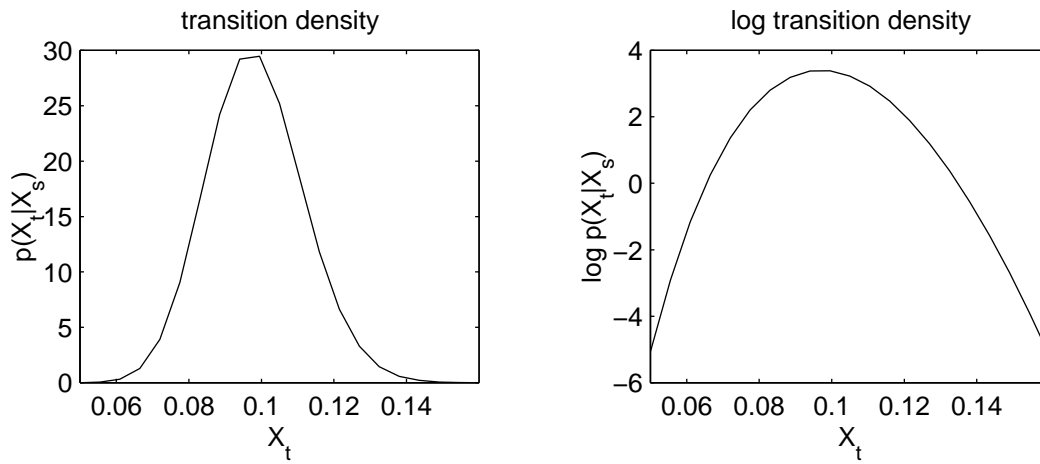


Figure 1: True transition density and log transition density for the CIR model given $\Delta = t - s = 1/12$, $X_s = 0.10$ and $\theta = (0.5, 0.06, 0.15)$.

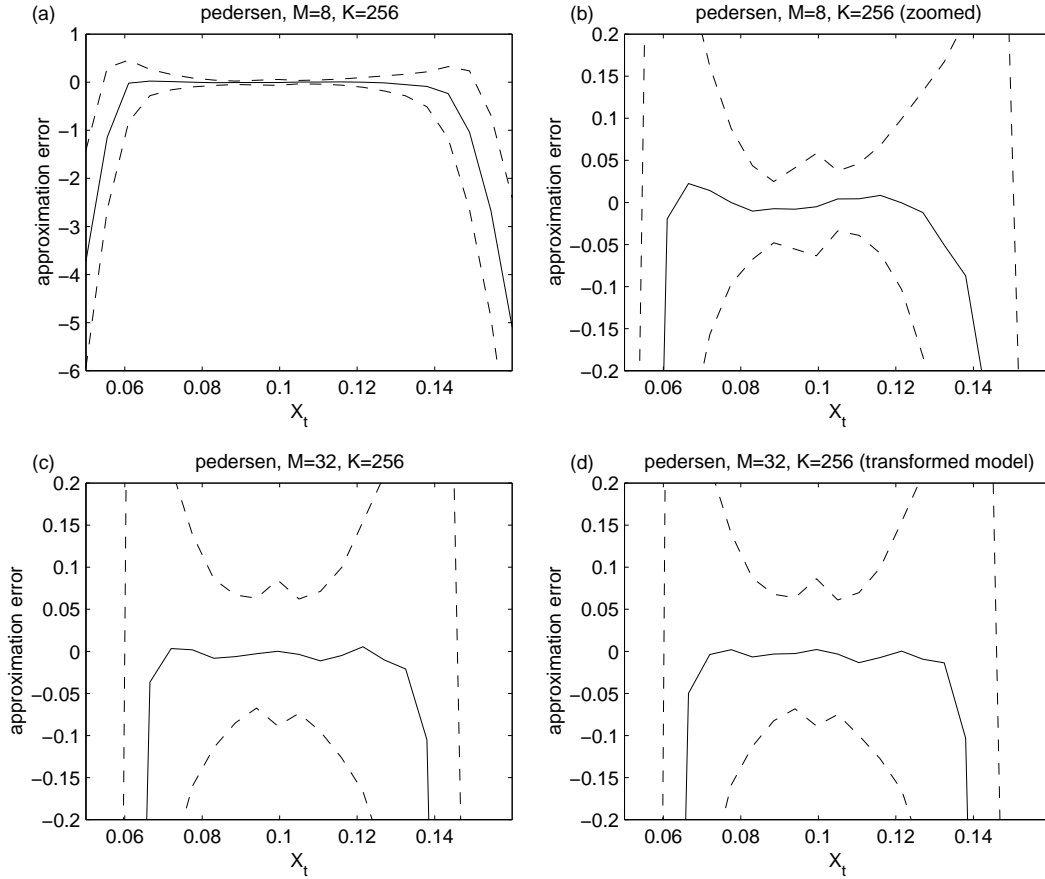


Figure 2: Approximation error, $\log \hat{p}(X_t, t; X_s, s) - \log p(X_t, t; X_s, s)$, using Pedersen's method given $\Delta = t - s = 1/12$, $X_s = 0.10$ and $\theta = (0.5, 0.06, 0.15)$. The median and interquartile range over 1024 repetitions are plotted. The untransformed model is used in panels (a)-(c), and the transformed model is used in panel (d).

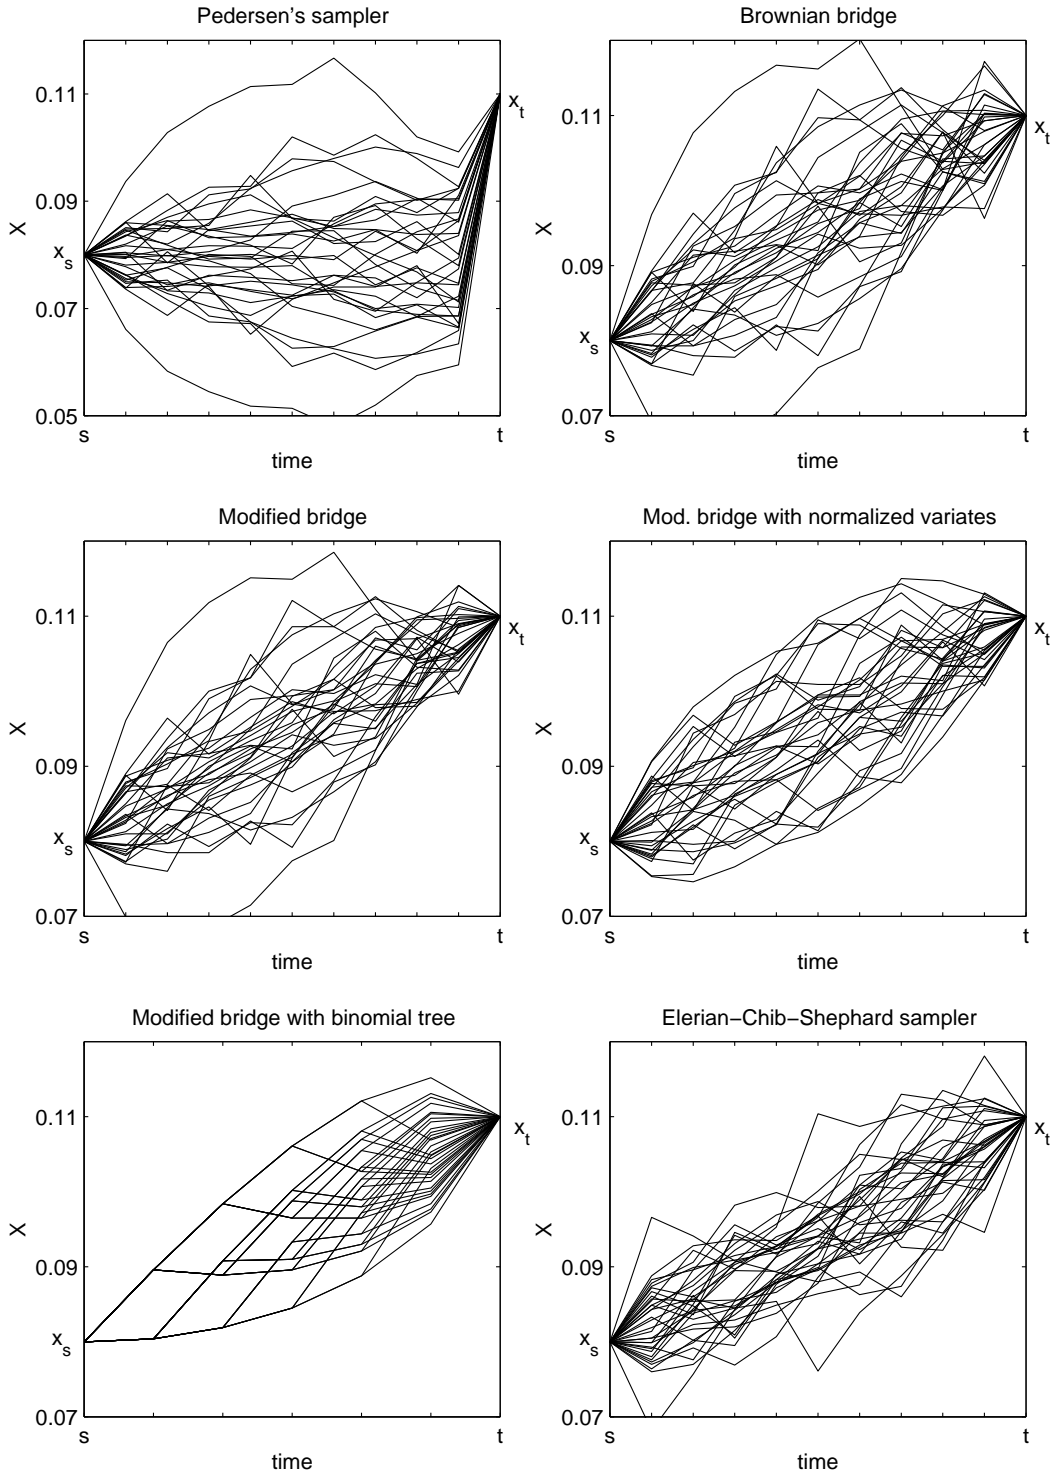


Figure 3: Simulated paths drawn using various importance samplers and random schemes.

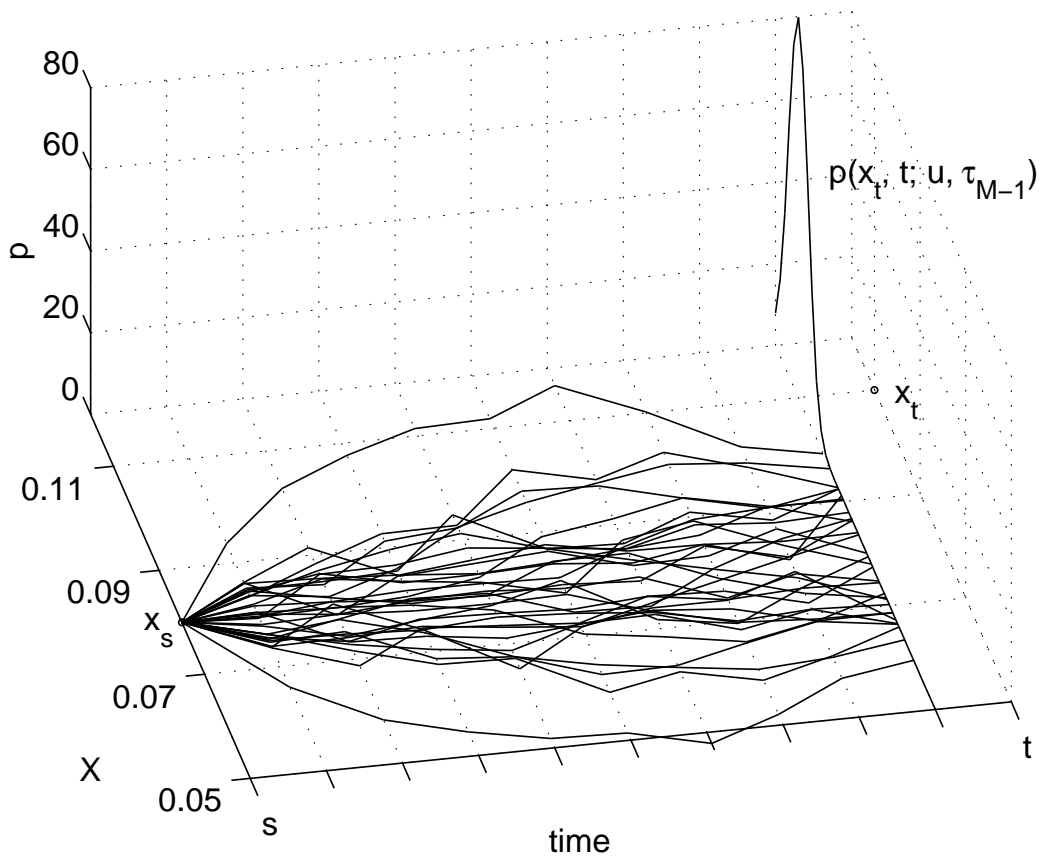


Figure 4: Illustration of equation (16). The terminal points of the sample paths represent draws from $P_{M-1}^{(M)}$, the curve represents the integrand.

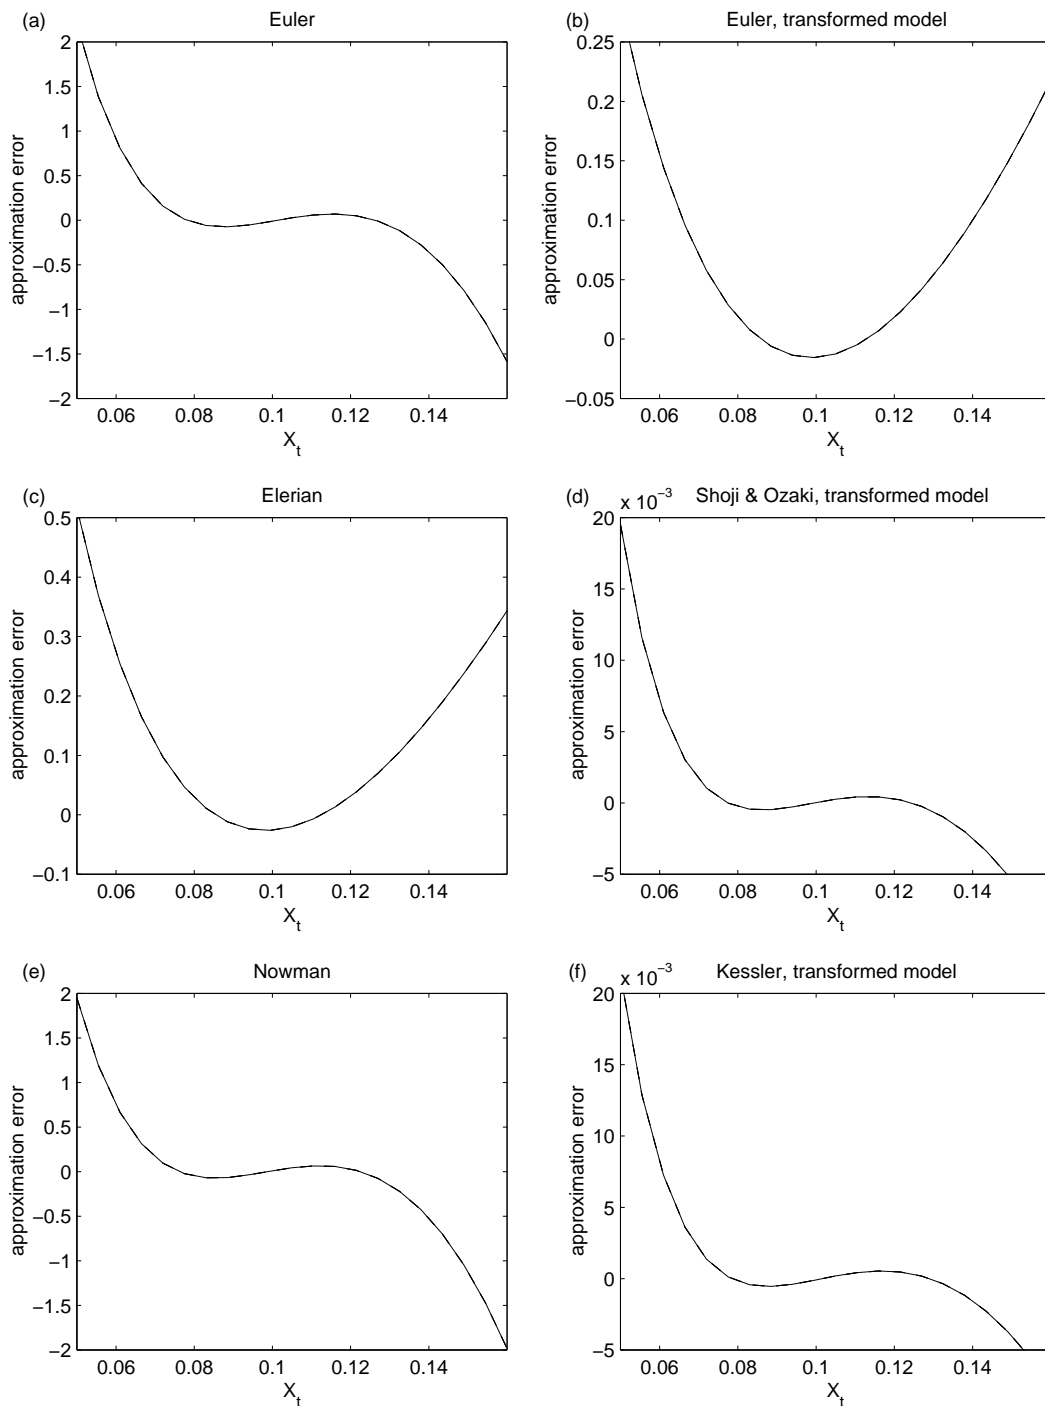


Figure 5: Approximation error, $\log \hat{p}(X_t, t; X_s, s) - \log p(X_t, t; X_s, s)$, for various schemes *without* using simulation (i.e., $M = 1$, $K = 1$) given $\Delta = t - s = 1/12$, $X_s = 0.10$ and $\theta = (0.5, 0.06, 0.15)$.

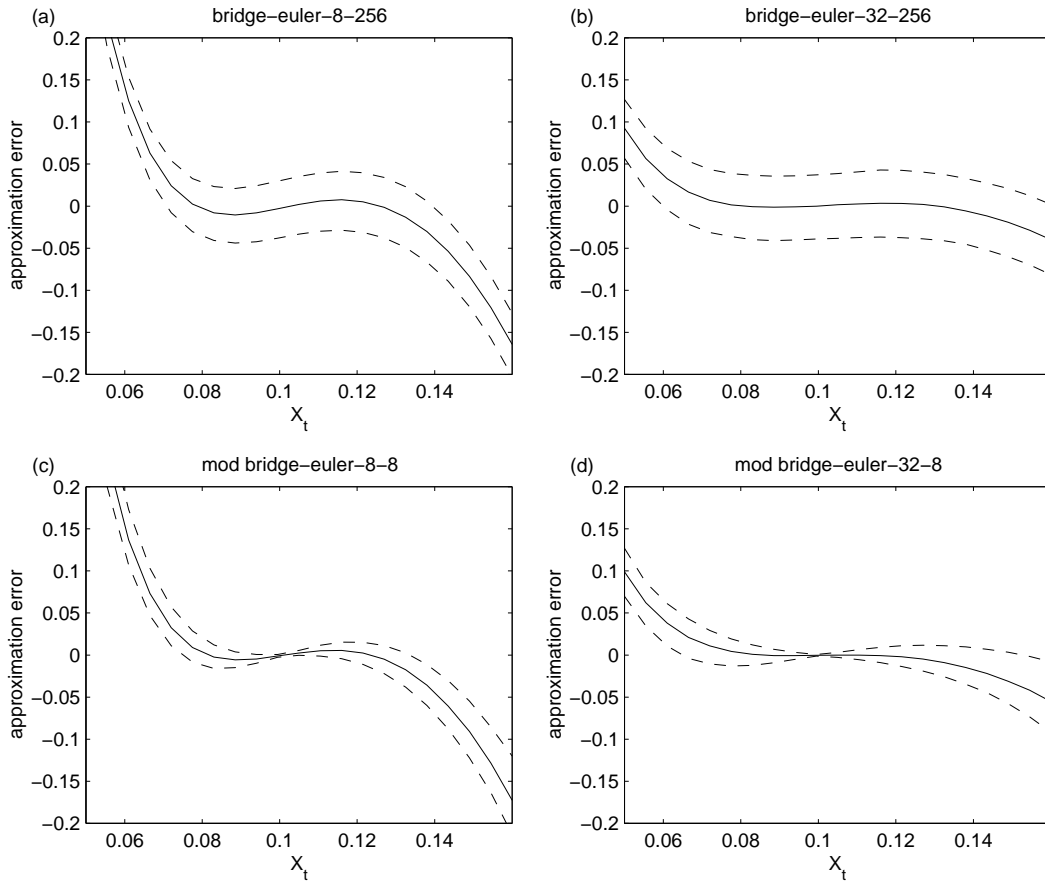


Figure 6: Approximation error, $\log \hat{p}(X_t, t; X_s, s) - \log p(X_t, t; X_s, s)$, for various simulated likelihood schemes given $\Delta = t - s = 1/12$, $X_s = 0.10$ and $\theta = (0.5, 0.06, 0.15)$. The median and interquartile range over 128 repetitions are plotted. The untransformed model is used.

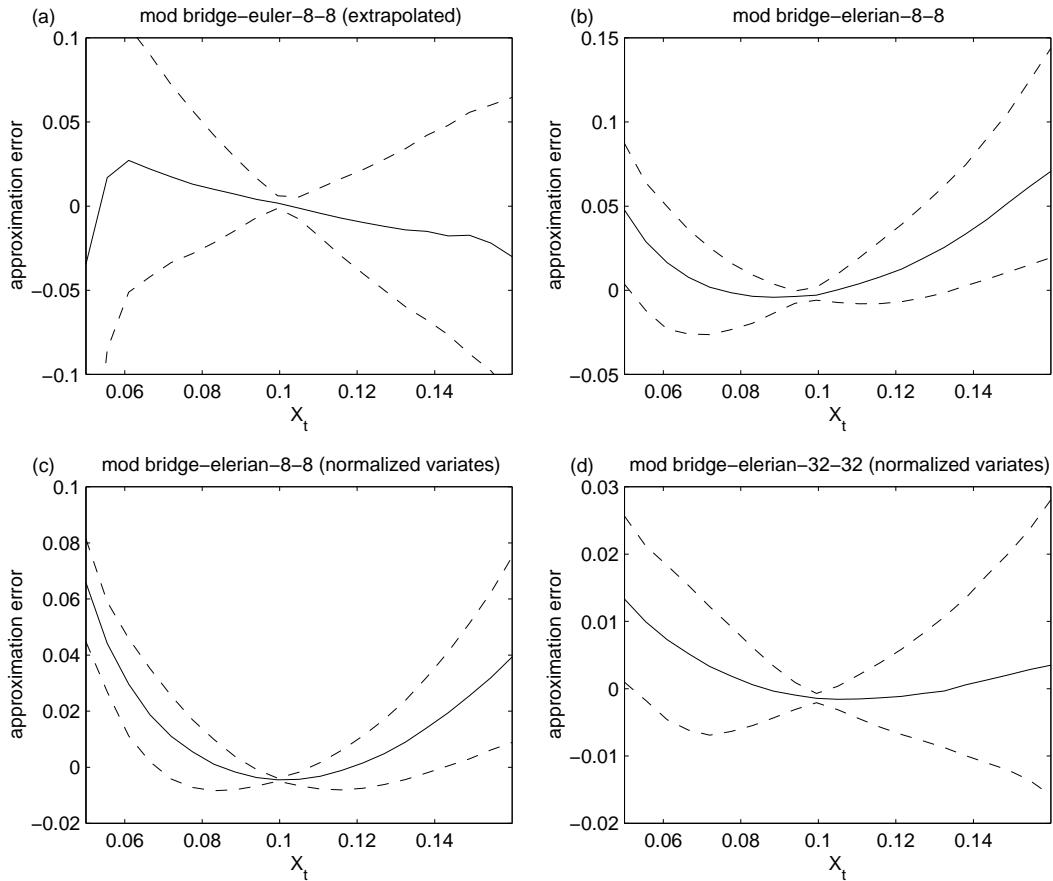


Figure 7: Approximation error, $\log \hat{p}(X_t, t; X_s, s) - \log p(X_t, t; X_s, s)$, for various simulated likelihood schemes given $\Delta = t - s = 1/12$, $X_s = 0.10$ and $\theta = (0.5, 0.06, 0.15)$. The median and interquartile range over 128 repetitions are plotted. The untransformed model is used.

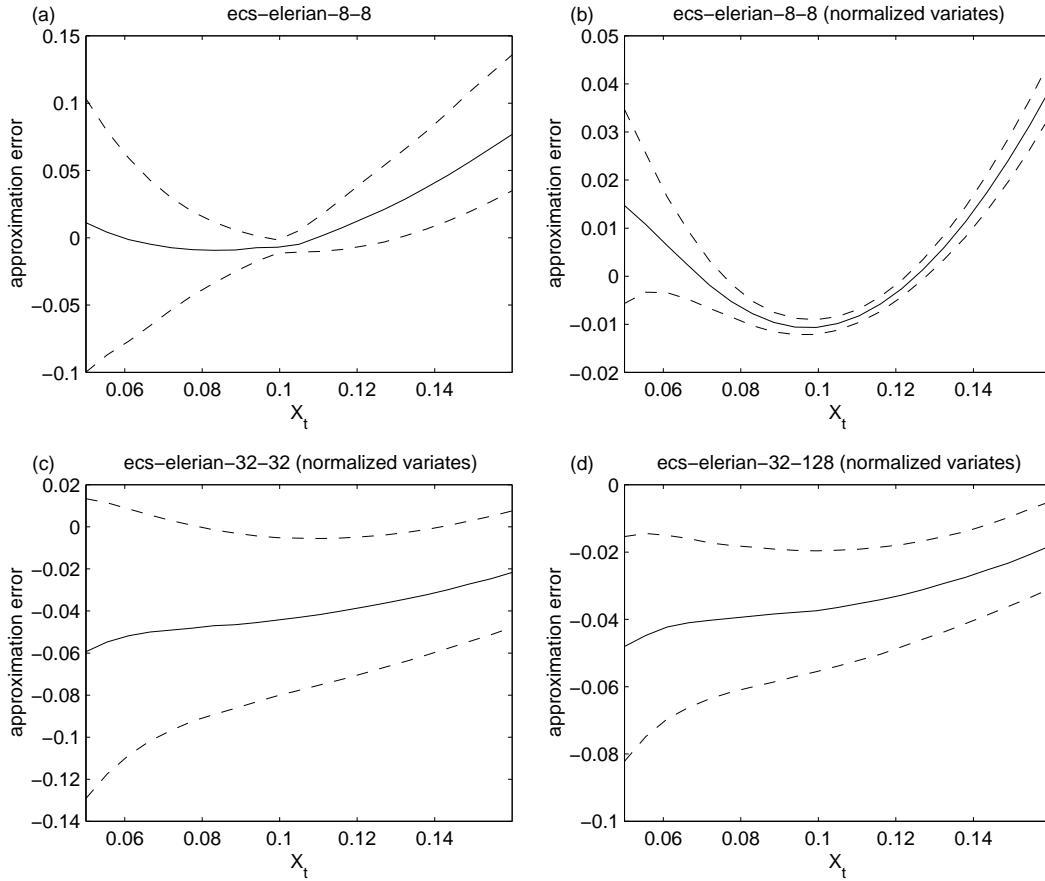


Figure 8: Approximation error, $\log \hat{p}(X_t, t; X_s, s) - \log p(X_t, t; X_s, s)$, for various simulated likelihood schemes given $\Delta = t - s = 1/12$, $X_s = 0.10$ and $\theta = (0.5, 0.06, 0.15)$. The median and interquartile range over 128 repetitions are plotted. The untransformed model is used.

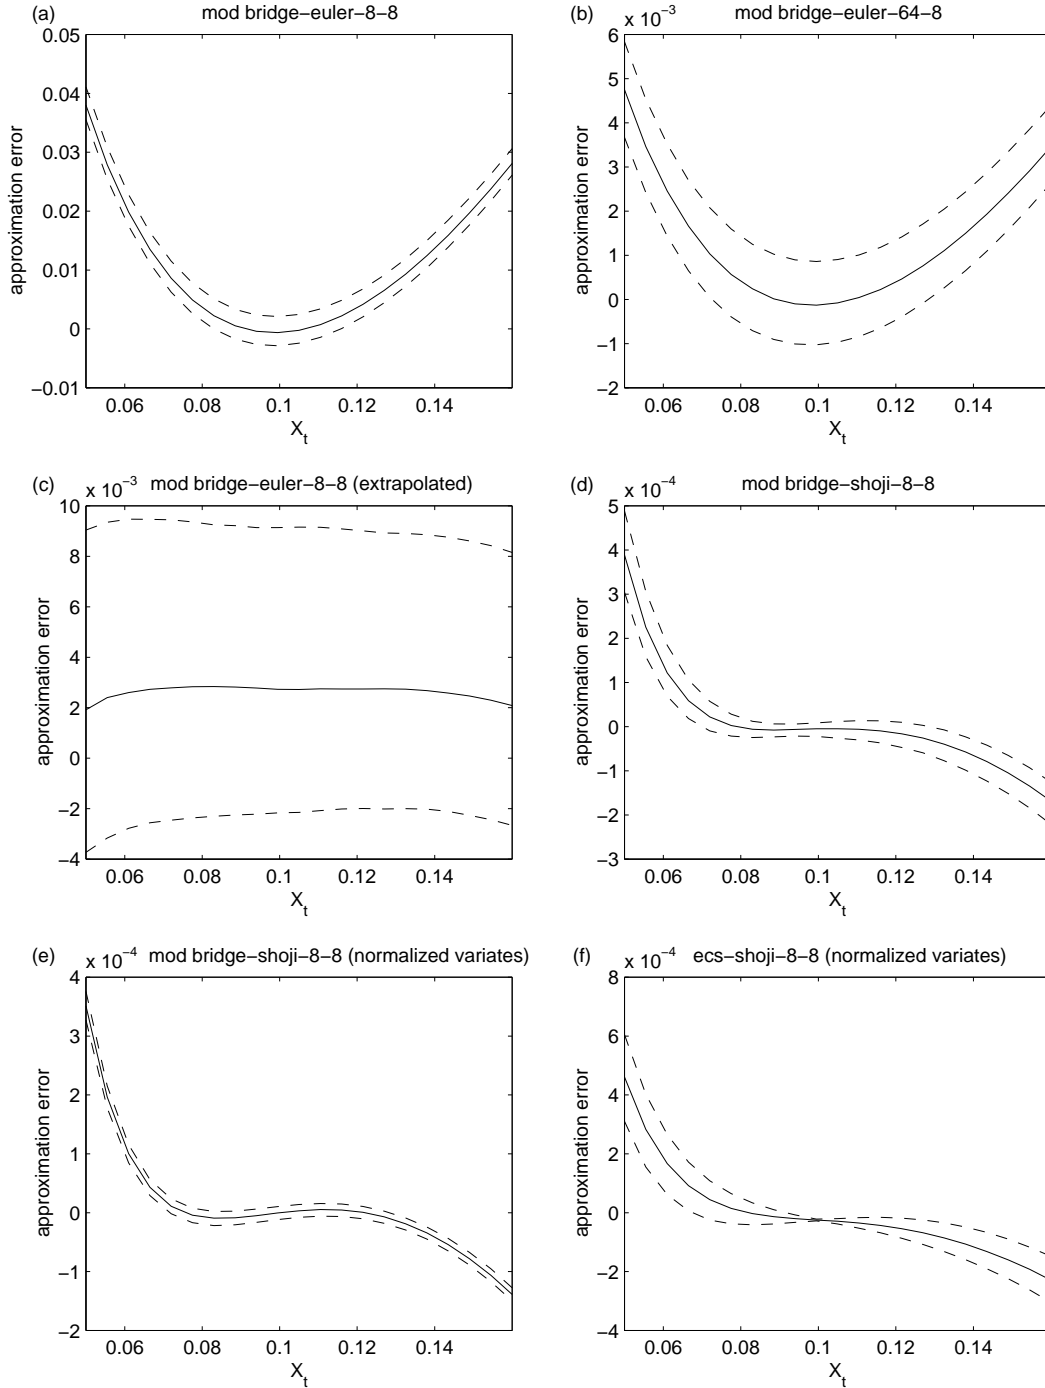


Figure 9: Approximation error, $\log \hat{p}(X_t, t; X_s, s) - \log p(X_t, t; X_s, s)$, for various simulated likelihood schemes given $\Delta = t - s = 1/12$, $X_s = 0.10$ and $\theta = (0.5, 0.06, 0.15)$. The median and interquartile range over 128 repetitions are plotted. The transformed model is used.

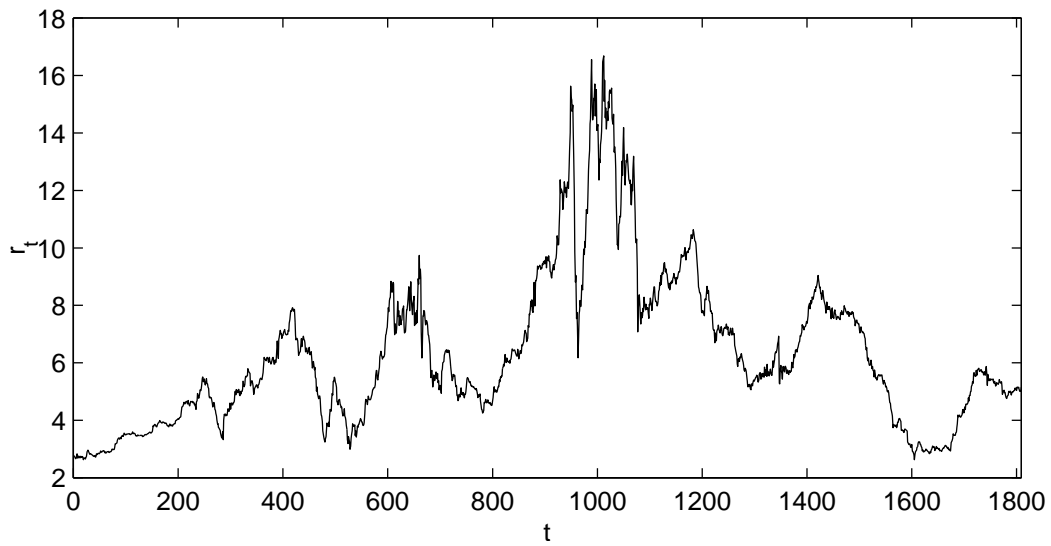


Figure 10: Weekly observations of the 3-month Treasury bill rate, 1/5/1962–8/30/1996, $n = 1809$.