

**Maximum Likelihood Estimation of
Stochastic Volatility Models**

By

**G Sandmann
and
S J Koopman**

DISCUSSION PAPER 248

FINANCIAL MARKETS GROUP
AN ESRC RESEARCH CENTRE

LONDON SCHOOL OF ECONOMICS



Any opinions expressed are those of the author and not necessarily those of the Financial Markets Group.

ISSN 0956-8549-248

Estimation of Stochastic Volatility Models via Monte Carlo Maximum Likelihood

G. Sandmann and S.J. Koopman

First Draft: June 1996

This version: May 1997

Abstract:

This paper discusses the Monte Carlo maximum likelihood method of estimating stochastic volatility (SV) models. The basic SV model can be expressed as a linear state space model with log chi-square disturbances. The likelihood function can be approximated arbitrarily accurately by decomposing it into a Gaussian part, constructed by the Kalman filter, and a remainder function, whose expectation is evaluated by simulation. No modifications of this estimation procedure are required when the basic SV model is extended in a number of directions likely to arise in applied empirical research. This compares favorably with alternative approaches. The finite sample performance of the new estimator is shown to be comparable to the Monte Carlo Markov chain (MCMC) method.

Key Words: GARCH model; Importance sampling; Kalman filter smoother; Monte Carlo simulation; Quasi-maximum likelihood; Stochastic Volatility; Unobserved components.

*Financial Markets Group and Department of Statistics, respectively, London School of Economics, Houghton Street, London WC2A 2AE, UK. We are indebted to Jim Durbin and Andrew Harvey as well as seminar participants of the ESEM 96 conference in Istanbul for helpful comments and suggestions. The comments of the associate editor and an anonymous referee helped greatly improve the quality of the paper. The second author is grateful to the ESRC for financial support as part of the project "Interrelationships in Economic Time Series", grant No. R000235330. All remaining errors remain our responsibility.

1. Introduction

The empirical distributions of financial time series differ substantially from distributions obtained from sampling independent homoskedastic Gaussian variables. Unconditional density functions exhibit leptokurtosis and skewness; time series of stock returns show evidence of volatility clustering; and squared returns exhibit pronounced serial correlation whereas little or no serial dependence can be detected in the return process itself. These empirical regularities suggest that the behaviour of financial time series may be captured by a model which recognises the time varying nature of return volatility, as follows:

$$R_t = \alpha + \beta R_{t-1} + \epsilon_t, \quad \epsilon_t \sim IID(0,1), \quad t = 1, \dots, T$$

where R_t denotes the return on an asset. One way of modelling ϵ_t is to express it as a deterministic function of lagged residuals, $\epsilon_t = \beta R_{t-1}$. Econometric specifications of this form are known as ARCH models and have achieved widespread popularity in applied empirical research (Bollerslev, Chow and Kroner, 1992; Bollerslev, Engle and Nelson, 1993; and Bera and Higgins, 1993). Alternatively, volatility may be modelled as an unobserved component following some latent stochastic process, such as an autoregression. Models of this kind are known as stochastic volatility (SV) models (Taylor, 1994; Ghysels, Harvey and Renault, 1996; Shephard, 1996).

Despite their intuitive appeal, SV models have been used less frequently than ARCH models in empirical applications. This is due to the difficulties associated with the estimation of SV models. Unlike ARCH models, where the likelihood function can be evaluated exactly, the likelihood function of an SV model is hard to construct. Existing estimation procedures can be subdivided into two groups: (i) methods that attempt to build the full likelihood function, and (ii) methods which rely on alternative, usually less efficient principles.

Several propositions have been made as to how the full likelihood function may be evaluated. Kim, Shephard and Chib (1996) show how the likelihood can be constructed when a mixture of normals is used to approximate the density of the disturbances. Jacquier, Polson and Rossi (1994) have proposed a Bayesian approach to the estimation of SV models using the Monte Carlo Markov chain (MCMC) technique. Fridman and Harris (1996) show how the extended Kalman filter can be used to perform numerical integration. Finally, Danielsson (1994a) suggested that accurate approximations to the likelihood function can be obtained by means of importance sampling.

Recently, Shephard and Pitt (1997) and Durbin and Koopman (1997a, b) designed methods for constructing the likelihood function for general state space models using Monte Carlo simulation techniques; thereafter referred to as Monte Carlo likelihood (MCL). This paper shows how the general concepts can be implemented efficiently for the standard linear SV model and for a variety of SV model extensions. The properties of the MCL estimates are compared with other approaches by using Monte Carlo experiments and empirical illustrations.

¹ The Quasi-Maximum Likelihood (QML) method of Harvey, Ruiz and Shephard (1994), and GMM methods of Andersen and Sørensen (1996) are examples of this category.

The crucial feature of the MCL approach is the formulation of the SV model in a linear state space form with $\ln(r_t^2)$ disturbances in the measurement equation. The question of how to treat zero observations - which will make $\ln(r_t^2)$ ill-defined - is addressed explicitly. The linear state space form allows very powerful algorithms for filtering and smoothing to be utilised, and more generally, to draw upon a vast body of knowledge on structural time series models (Harvey, 1989). The gain is also due to the elegant form of the likelihood specification as put forward by Durbin and Koopman (1997a) in which the log likelihood is the sum of the Gaussian log likelihood (i.e. QML) and a correction for the departures from Gaussian assumptions. Monte Carlo simulation is only employed to construct the correction part of the likelihood function.

Apart from reducing the computational effort considerably (while attaining finite sample efficiency), the algorithm has several distinct advantages. First, the sampling variation can be reduced giving arbitrarily close approximations to the true likelihood function value. Thus standard likelihood ratio tests can be constructed. This is likely to be very useful since numerical standard errors of model parameters often leave much to be desired.

Secondly, a wide range of extensions can be addressed with trivial modifications of the estimation procedure due to the fact that the state space form is retained. Thus, several well known extensions of the basic SV model can be examined: fat tailed distributions for the mean equation disturbances, SV in the mean specification, correlated return and volatility processes, as well as stochastic seasonal components, effects of dummy and exogenous explanatory variables may be explored in detail.

The paper is organised as follows. Section 2 discusses in more detail the various aspects of estimation and inference in the context of SV models. In Section 3 we describe the Monte Carlo likelihood (MCL) estimation method while Section 4 compares its finite sample performance with existing techniques by means of an extensive Monte Carlo experiment and some empirical data examples. Section 5 illustrates how the method can be applied to the estimation of models with fat-tails, explanatory variables and non-zero correlation. Section 6 concludes.

2. Stochastic Volatility Model

Consider the stochastic volatility model:

$$\begin{cases} r_t = \bar{r} e^{0.5h_t} \varepsilon_t & \varepsilon_t \sim N(0,1) \\ h_t = \bar{h} + \alpha(h_{t-1} - \bar{h}) + \beta \eta_t & \eta_t \sim N(0, \sigma^2) \end{cases} \text{Corr}(\varepsilon_t, \eta_t) = \rho \quad (1)$$

where $r_t = R_t - \bar{r}$ is the mean adjusted return on an asset. Since many financial time series exhibit little or no dynamic behaviour in the mean, but pronounced serial dependence in the variance (Bollerslev, Chow and Kroner, 1992) the estimation of ε_t will not be the subject of interest in the present context. The average volatility level is denoted by \bar{h} , and the mean and variance equation disturbances ε_t and η_t may be contemporaneously correlated.

Despite a very parsimonious representation, this model captures most of the empirical regularities found in financial time series (Ghysels, Harvey and Renault, 1996). An

attractive feature of specification (1) is the possibility of linearising the model. By taking logarithms of the squared mean adjusted returns one obtains

$$\begin{cases} \ln r_t^2 = \ln \sigma_t^2 + h_t + \varepsilon_t & \varepsilon_t = \ln \left(\frac{\sigma_t^2}{\sigma_{t-1}^2} \right) \\ h_t = h_{t-1} + \eta_t & \eta_t \sim N(0, \pi^2) \end{cases} \text{Corr}(\varepsilon_t, \eta_t) = 0 \quad (2)$$

If the original mean equation disturbance, ε_t , is standard normal, $\ln r_t^2$ follows the $\ln(\sigma_t^2)$ distribution whose mean and variance are known to be -1.27 and $\pi^2/2$, respectively. Notice, that once the transformation is accomplished, the information regarding the correlation coefficient, ρ , is lost, but can be recovered by conditioning on the signs of the original observations (Harvey and Shephard, 1996). Estimation of ρ is addressed in section 5.4.

Harvey, Ruiz and Shephard (1994) suggested a Quasi-Maximum Likelihood (QML) method of estimating the model based on the Kalman filter. Assuming joint conditional normality of (ε_t, η_t) , equation (2) represents the measurement and transition equations of the general linear state space model, details of which can be found in the Appendix. Once the model is in the state space form, the advantages of this approach become evident: (i) explanatory variables can be easily incorporated into the variance equation, (ii) more general ARMA processes can be assumed for the evolution of the latent variable, (iii) missing or irregularly spaced observations can be handled, and (iv) generalisations to the multivariate case are straightforward.

The QML method approximates the distribution of ε_t by $N(-1.27, \pi^2/2)$, while $\ln r_t^2$ is far from being Gaussian. In fact, its density is given by:

$$p_{\ln r_t^2}(z) = \frac{1}{\sqrt{2}} \exp\left[\frac{z - e^z}{2}\right] \quad (3)$$

Figure 1 below shows in how far $\ln r_t^2$ deviates from its normal approximation which implies that the QML estimator is likely to have poor small sample properties even though it is consistent.

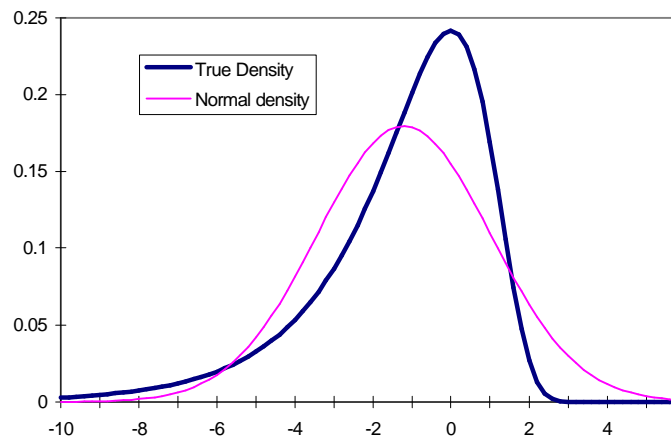


Figure 1: The $\ln(\sigma_t^2)$ density and the Normal approximation $N(-1.27, \pi^2/2)$.

² Thus the logarithm of variance follows a discrete time analogue of the Orstein-Uhlenbeck process often used in the context of option pricing (Hull and White, 1987; Renault and Touzi, 1996).

Note the high degree of skewness and the long tail in on the negative half line. Large negative values reflect inliers in r_t , which may arise in empirical applications with high frequency data.

Several other estimation techniques achieved prominent attention in the literature. First, various method of moments (MM) estimators have been suggested by Taylor (1986), Melino and Turnbull (1990), and Andersen and Sørensen (1996), among others. MM estimators avoid the problem associated with the linearisation of the model as well as the evaluation of the likelihood. They are not difficult to implement and to generalise but the efficiency of these estimators is known to be suboptimal to the likelihood-based method of inference (Jacquier, Polson and Rossi, 1994; Andersen and Sørensen, 1996). Gallant and Tauchen (1996) have developed a MM procedure which matches the scores of an auxiliary model via simulation (SMM). They claim that, if the auxiliary model is a good approximation to the distribution of the data, the MM estimator is as efficient as maximum likelihood. However, none of the MM estimation methods provide an estimate of the instantaneous volatility σ_t^2 throughout the sample, $t = 1, \dots, T$, so that an additional form of estimation is required. For instance, Andersen (1994) and Ghysels and Jasiak (1996) use MM techniques to estimate the parameters and they use the Kalman filter to obtain volatility estimates.

Secondly, Kim, Shephard and Chib (1996) suggest to approximate the distribution of σ_t by a mixture of normals. Given a particular mixture, the likelihood can be computed via the prediction error decomposition since the linear structure of the model is essentially retained. An important drawback of this method is that no matter how many mixture components are used, the mixture of normals cannot give a good approximation to the tail behaviour of the $\ln(\sigma_t^2)$ distribution. In addition, the convergence of the algorithm is likely to be very sensitive to the number and weight of individual mixture components (Jacquier, Polson and Rossi, 1994).

Thirdly, Fridman and Harris (1996) suggest that the non-Gaussianity of the measurement equation disturbances can be handled by means of a “brute force” numerical integration. In a Monte Carlo study - similar to the one presented here - the authors demonstrate how Kitagawa’s (1987) extended Kalman filter can be applied in this context. By retaining the state space form, Fridman and Harris’ (1996) estimation technique offers the same advantages as the MCL. Some of the disadvantages of this method consist of computational inefficiencies (the extended Kalman filter is known to be rather slow) and the necessity to choose *a priori* a fixed grid, over which the volatility process will be integrated. This creates a trade-off between numerical accuracy on the one hand, and computational efficiency on the other. It is conceivable, that in some instances an optimal grid may not exist. For instance, when estimating the volatility process around the stock market Crash of ’87 the grid selection procedure proposed by Fridman and Harris (1996) will either lead to a very coarse grid over the entire volatility range, or place no probability weight on the high volatility state during the Crash.

A Bayesian approach to the estimation of SV models using a Monte Carlo Markov chain (MCMC) technique was developed by Jacquier, Polson and Rossi (1994), JPR thereafter. They have performed extensive simulation experiments which demonstrate that MCMC is superior to QML and MM estimation techniques across a wide range of parameter values. However, this technique has some undesirable features. The procedure is quite involved, requiring a large amount of computer intensive

simulations. In addition, the method needs to be nontrivially modified for the extensions like the introduction of explanatory variables, alternative processes for the evolution of variance, or multivariate specifications (Jacquier, Polson and Rossi, 1995). Shephard and Pitt (1997) have constructed an efficient block MCMC algorithm for performing Bayesian inference on general non-linear and non-Gaussian state space models of which the SV model (1) is a special case. They conclude that the performance of the multi block MCMC methods outperforms the single block approach of JPR in terms of computational efficiency.

Finally, Danielsson (1994a) proposed to estimate the SV model by the Monte Carlo likelihood (MCL) estimation method. His accelerated Gaussian importance sampler (AGIS) algorithm is a simulation-based technique whose time requirement and precision is on par with MCMC. However, the method is difficult to generalise and remains computationally expensive largely due to the failure of the technique to exploit the linear structure resulting from the transformation (2).

The present paper discusses how a very efficient MCL estimator can be obtained with retaining the linear state space form. Following the work of Shephard and Pitt (1997) and Durbin and Koopman (1997a, b) we demonstrate that MCL is a viable alternative to the MCMC technique. It is shown that: (i) the finite sample performance of the MCL parameter estimators is at least as good as the one obtained by MCMC; (ii) the computational requirement for MCL is smaller than MCMC; (iii) only trivial modifications are required to the MCL when the basic model is extended in a number of interesting directions; (iv) inference in the SV model can be performed by means of Likelihood Ratio test statistics because the likelihood function can be approximated arbitrarily close.

3. Monte Carlo Maximum Likelihood Estimation.

3.1. The general algorithm

Taking logarithms of the squared residuals in (1), that is $y_t = \ln(r_t^2)$, gives the SV model in the linear state space form (2) but invokes an additional difficulty: the disturbance term in the measurement equation becomes non-Gaussian. The Appendix discusses the general linear state space model and the associated algorithms for filtering, smoothing and simulation. The Monte Carlo likelihood (MCL) approach for non-Gaussian models such as the SV model is based on importance sampling techniques (Ripley, 1987). Danielson (1994a) and Shephard and Pitt (1997) consider generating samples from an approximating Gaussian model. Durbin and Koopman (1997a) demonstrate that the log likelihood function of state space models with non-Gaussian measurement disturbances can be simply expressed as:

$$\ln L(y| \theta) = \ln L_G(y| \theta) + \ln E_G \left[\frac{p_{true}(\epsilon| \theta)}{p_G(\epsilon| \theta)} \right] \quad (4)$$

where $y = (y_1, y_2, \dots, y_T)'$, $\theta = (\theta_1, \theta_2, \dots, \theta_T)'$, $\ln L_G(y| \theta)$ is the log likelihood function of the approximating Gaussian model, $p_{true}(\epsilon| \theta)$ is the density function of the measurement disturbances, that is the $\ln(\chi^2_1)$ density in the case of the basic SV model,

$p_G(\cdot | \cdot)$ is the Gaussian density of the measurement disturbances of the approximating model and E_G refers to expectation with respect to the so-called importance density $p_G(\cdot | y, \cdot)$ associated with the approximating model. Equation (4) reveals that only importance samples are required for the measurement disturbances $\mathbf{v} = (v_1, v_2, \dots, v_T)'$. Furthermore, equation (4) shows that the non-Gaussian log likelihood function can be expressed as the log likelihood function of the Gaussian approximating model plus a correction for the departures from the Gaussian assumptions in relation to the true model. The unbiased estimate of (4) is given by:

$$\ln \hat{L}(\cdot) = \ln L_G(y | \cdot) + \ln \bar{w} + \frac{s_w^2}{2N\bar{w}^2} \quad (5)$$

where \bar{w} and s_w^2 are computed by the algorithm of Durbin and Koopman (1997a):

1. Choose a Gaussian approximating model from which a feasible sampling scheme can be deduced based on the importance density $p_G(\cdot | y, \cdot)$; see section 3.2 for details.
2. Compute (and store) $\ln L_G(y | \cdot)$ and $\hat{\cdot} = E(\cdot | y, \cdot)$ for the approximating model via the Kalman filter smoother; see Appendix for details.
3. Generate a sample $\mathbf{v}^{(i)} = (v_1^{(i)}, \dots, v_T^{(i)})'$ from the importance density $p_G(\cdot | y, \cdot)$ referring to the approximating model of step 1. A specific version of the simulation smoother of de Jong and Shephard (1995) is used to generate this sample; see Appendix for details.
4. Construct an antithetic sample: $\tilde{\mathbf{v}}^{(i)} = 2\hat{\cdot} - \mathbf{v}^{(i)}$.
5. Compute (and store) $w^{(i)} = \frac{w(\mathbf{v}^{(i)}) + w(\tilde{\mathbf{v}}^{(i)})}{2}$ where $w(\cdot) = \frac{p_{\ln \hat{L}}(\cdot | \cdot)}{p_G(\cdot | \cdot)}$.
6. Repeat steps 3 to 5 until N samples are drawn.
7. Calculate \bar{w} and s_w^2 as the sample mean and variance of $w^{(i)}$, $i=1, \dots, N$.

The MCL estimates of model parameters, $\hat{\cdot}$, are obtained by numerical optimisation of (5). The log likelihood function of the approximating model, $\ln L_G(y | \cdot)$, can be used to obtain starting values. The choice of N governs the accuracy of the approximation to the likelihood function: as N increases, the approximation becomes more accurate. For practical purposes we find that $N=5$ is sufficient; see the detailed discussion in Section 4 below.

3.2. Implementation

The importance sampling density $p_G(\cdot | y, \cdot)$, could be based on the approximating SV model as given by (2) with $v_t \sim N(0, H_t)$ where $H_t = \sigma^2/2$, for $t = 1, \dots, T$. However, Durbin and Koopman (1997a) argue that a better importance density is obtained from (2) with $v_t \sim N(0, \tilde{H}_t)$ where the scalar variances \tilde{H}_t 's are chosen so as to make the

differences between the log densities $\ln p_{\ln \hat{y}_t}(\cdot | \hat{y}_t)$ and $\ln p_G(\cdot | \hat{y}_t)$ as constant as possible in the neighbourhood of $\hat{y}_t = E(\hat{y}_t | y_t)$ ³. The smoothed error vector $\hat{\epsilon} = (\hat{\epsilon}_1, \dots, \hat{\epsilon}_T)'$ is calculated by the Kalman filter smoother; details of which can be found in the Appendix. Intuitively, large negative values of $\hat{\epsilon}_t$ would require high values of \tilde{H}_t in order for the slopes of the densities in **Figure 1** to be roughly equal, for $t = 1, \dots, T$. The choice of the variances \tilde{H}_t 's of the approximating model is determined by equalising the derivatives of the log densities at:

$$\frac{1}{2} \left(1 - e^{-\hat{\epsilon}_t}\right) \Big|_{\hat{\epsilon}_t} = - \frac{1}{\tilde{H}_t} \Big|_{\hat{\epsilon}_t},$$

where the left hand side is the derivative of (3) and the right hand side is the derivative of the Gaussian density. This leads to a set of T equations:


$$\tilde{H}_t = \frac{2\hat{\epsilon}_t}{e^{\hat{\epsilon}_t} - 1} \quad t = 1, \dots, T. \quad (6)$$

Note that equations (6) ensure the non-negativity of \tilde{H}_t since the nominator and denominator are of the same sign for any value of $\hat{\epsilon}_t$. The set of nonlinear equations (6) is solved for \tilde{H}_t by iteration, starting at $\tilde{H}_t^{(0)} = \pi^2/2 \forall t$. Given a parameter vector θ , we iterate K times between computing $\hat{\epsilon}_t$ (using the Kalman filter smoother) and computing $\tilde{H}_t^{(k)}$ by (6), for $k = 1, \dots, K$. This procedure is performed only once in Step 1 of the algorithm of section 3.1. Naturally, when the parameter vector changes, the algorithm needs to be repeated.

Table 1 shows rapid convergence results across a range of parameter values for the simulated SV model. Choosing the metric $d(k) = T^{-1} \sum_t |\tilde{H}_t^{(k)} - \tilde{H}_t^{(k-1)}|$ to describe successive changes in the variance vector, \tilde{H} we find that after about six to eight iterations the elements of the variance vector cease to fluctuate, i.e. $\tilde{H}_t^{(k)} \rightarrow \bar{H}_t$. The individual elements of the variance vector \bar{H} are now different across $t=1, \dots, T$.

Figure 2 is a histogram of \bar{H}_t 's from a simulated SV process ($T=1,000$) governed by a set of parameters $\theta_5 = (\bar{\alpha}, \beta, \sigma^2, \gamma, \delta)_5$, the numerical values of which are discussed in Section 4. It is the mirror image of the density of $\hat{\epsilon}_t$ and reconfirms the intuition behind the method: large negative, but infrequent values of $\hat{\epsilon}_t$ require high values of variance parameter \bar{H}_t in order to compensate for the difference in density slopes in this region. The converse is represented by large probability mass $d\bar{H}_t$ or less than $\pi^2/2$.

³ A similar suggestion is made by Shephard and Pitt (1997) but their argument is based on finding an approximating Gaussian model which produces posterior mode estimates for the true model.



The linear Kalman filter smoother does not explicitly take account of the non-Gaussianity in the measurement equation. However, when the Kalman filter smoother is applied to the approximating Gaussian SV model, it is effectively computing the posterior mode estimates of the volatility (Shephard and Pitt, 1997; Durbin and Koopman, 1997b). If the posterior mean is required, Durbin and Koopman (1997b) show that a computationally efficient algorithm is given by:

$$\hat{\sigma}_{t|T} = y_t - \sum_{i=1}^N f_t \hat{\sigma}_t^{(i)}, \quad f_t = \frac{w(\hat{\sigma}_t^{(i)})}{N\bar{w}_t} \quad (8)$$

where $\hat{\sigma}_{t|T}$ is the volatility estimate, $\hat{\sigma}_t^{(i)}$ is a draw from the importance density as computed by the simulation smoother (see Step 3 of the algorithm of section 3.1), $w(\hat{\sigma}_t^{(i)}) = \exp\{l(\hat{\sigma}_t^{(i)})\}$ is defined in (7) and $\bar{w}_t = \frac{1}{N} \sum_{i=1}^N w(\hat{\sigma}_t^{(i)})$. The weights f_t can be interpreted as the corrections for non-Gaussianity in the measurement equation. It is worth stressing that equation (8) requires little additional computing because the quantities $\hat{\sigma}_t^{(i)}$ and $w(\hat{\sigma}_t^{(i)})$ are already evaluated by the algorithm of section 3.1.

Finally, the estimation error, $\hat{\sigma}_{t|T} - \sigma_t$ (where σ_t denotes the true volatility) is $O(1)$ which implies that treating $\exp(\hat{\sigma}_{t|T})$ as lognormal may lead to distortions (Harvey and Shephard, 1993). These authors consider the following estimate of the volatility process:

$$\hat{\sigma}_t^2 = \frac{\sigma_t^2}{T} e^{-\epsilon_t}, \quad \sigma_t^2 = T^{-1} \sum_{i=1}^T r_i^2 e^{-\epsilon_t} \quad t = 1, \dots, T \quad (9)$$

where the smoothed signal, $\hat{\sigma}_{t|T}$ is obtained from (8). Thus the estimation of model parameters and the values of the latent volatility process are addressed simultaneously in the MCL framework.

4. Finite Sample Performance.

4.1. Simulation experiment

To assess the performance of the new method we conducted simulation experiments following the design of Jacquier, Polson and Rossi (1994), thus facilitating direct comparison with the MCMC method. The range of parameter values $\theta = (\rho, \sigma^2)$ is selected in the following manner.⁵ First, the values of the autoregressive parameter are set to 0.90, 0.95, and 0.98. This choice is motivated by empirical studies which reported the values of the autoregressive coefficient close to unity, ranging between 0.9 and 0.995. Secondly, for each value of ρ , the values of σ^2 are selected so that the coefficient of variation:

$$CV = \frac{\text{var}(h)}{E[h]^2} = \exp\left(\frac{\sigma^2}{1 - \rho^2}\right) - 1$$

⁵ The correlation coefficient, ρ is set to zero.

takes the values 10, 1, and 0.1. High values of the ratio of volatility variance to its squared mean indicate pronounced relative strength of the stochastic volatility process while low values of CV signify that the model is close to the one of constant volatility. In fact, if preliminary exploratory analysis of the data from a model with low CV ($CV=0.1$) was based only on the autocorrelation structure of r_t^2 or $\ln(r_t^2)$ the practitioner without a strong prior belief that the SV model is the correct specification will be unable to distinguish between the SV and a homoskedastic model. Nevertheless, the parameter triplets $\gamma- \rho$ are included for completeness. The focus of interest is thus centered around parameter triplets $\alpha- \beta$ which correspond to the coefficient of variation close to unity. Most of the empirical studies surveyed by JPR report parameter estimates in this range. Finally, the values of the location parameter, μ are chosen such that the expected variance:

$$E[h] = \mu^{-2} \exp\left(\frac{\sigma^2}{2(1-\rho^2)}\right)$$

is set to 0.0009. If the simulated data are regarded as weekly returns, this corresponds to approximately 22% annualised variance. Note that JPR chose a slightly different parameterisation of the SV model ($r_t = e^{\frac{1}{2}h_t} \epsilon_t$, $h_t = \alpha + \beta h_{t-1} + \sigma \epsilon_t$) but the intercept parameter μ is mapped into α via $\ln(\mu^{-2}) = (1-\beta)^{-1}$. This gives nine parameter triplets, $\theta_i = (\alpha, \beta, \rho)_i$, $i=1, \dots, 9$ the values of which are reported in **Table 2** in the row labelled “True”. For each θ_i we generate samples of length $T=500$, we estimate the model via MCL and via QML and we compute means, standard deviations and mean squared errors of the parameter estimates over $K=500$ simulated realisations of the process. In these calculations the number of draws N used in the algorithm of section 3.1 is set to $N=5$.

Results from the sampling experiments are presented in **Table 2** which is divided into three panels in accordance with the coefficient of variation CV. Within each panel the true parameter values are displayed first. The results for the Quasi-Maximum Likelihood estimator (QML) are reported second, followed by the MCL estimator. Finally, JPR’s simulation results for the Bayes (MCMC) estimator, their Table 7, are reproduced in the final row. The starting parameter values for both, the QML and MCL optimisation routines are obtained from a two dimensional grid search procedure which searches for an optimum across the surface of the Gaussian log likelihood function $\ln L_G(\theta)$.

The figures presented in **Table 2** allow several conclusions to be drawn. The experiment demonstrates that the MCL estimator exhibits similar efficiency as the MCMC estimator across most parameter values. For the $CV=10$ and $CV=1$ the mean squared errors on all parameters (except for ρ when $CV=1$) are smaller for the MCL. When $CV=1$ the mean squared errors on α are slightly larger, but the estimator exhibits lower bias (alas larger sample variance).

The QML estimator is found to be inefficient, thereby confirming the results of JPR. **Figures 3** and **4** present the smoothed densities of the estimates of α and β for two triplets θ_4 and θ_5 . The MCL estimator is shown to exhibit a much tighter sampling distribution than the QML estimator, a property also indicated by smaller standard errors of the estimates.




Figure 3: Sampling distributions of the MCL and the QML estimators;₄




Figure 4: Sampling distributions of the MCL and the QML estimators;₅

Across the entire parameter space the standard errors of the QML estimator are at least twice the size of the fully efficient MCL estimator while the bias is non-negligible. The efficiency of the QML estimator increases as the strength of the SV process becomes more pronounced. For instance, for $CV=10$ the sample standard error on $\omega=0.95$ is 0.04 while in the case of $CV=1$ the standard error on $\omega=0.95$ increases fourfold to 0.18.

However, although we find QML to be inefficient, its performance is nowhere as near as bad as reported by JPR. Same conclusion was reached by Breidt and Carrquiry (1996) who also re-examined the final sample performance of the QML estimator. Since **Figures 3** and **4** were constructed so as to correspond to JPR's Figures 4 and 5 respectively, direct comparison reveals dramatic differences in the performance of the same estimation technique. This raises the question of possible inefficiencies in JPR's QML estimation method such as poor starting values, different convergence criteria, or inefficient implementation of the algorithm.

We also find that the performance of *all three estimators* deteriorates as CV decreases. Comparison of the MCL and the MCMC estimators in this region ($CV=0.1$) reveals that the MCL estimates of the long run volatility level ω (i.e. ω) are more efficient, but those of α and β are less efficient than the corresponding MCMC estimates. This can be explained by the fact that the dynamic properties of the model are so weak (or the signal-to-noise ratio so small) in this region that the data appear almost indistinguishable from a constant volatility model.

The performance of the estimator is examined in cases when larger data sets are available, results of which are presented in **Table 3**. For each of the parameter triplets (α, β, ω) (corresponding to $CV=1$) $K=500$ samples of length $T=2,000$ were drawn and the simulation experiment repeated. Across all parameters the standard errors are smaller by a factor of two, when compared to the relevant entries in Panel 2, **Table 2**. The

mean squared errors are similarly reduced. Comparison with MCMC (also based only on μ) reveals that the finite sample efficiency is very similar. The mean squared errors are of the same order of magnitude: with the MSE on σ larger, the MSE on ρ identical, and MSE on ω (i.e. σ^2) being smaller.

As the number of draws N increases, the expectation of the MCL likelihood function (5) can be calculated more precisely, thus leading to increased performance. On the other hand the computational burden needs to be taken into account. In our experience, a very small number of draws ($N=5$) is sufficient to produce results comparable with the MCMC estimator while little can be gained by increasing N from 20 to 40. This is illustrated in **Table 4**. We revert to the original simulation study design of $T=500$ observations (with $K=500$ realisations of the process) and re-estimate the model for the central parameter triplet, μ with a varying number of draws⁶: $N=5$, $N=20$, $N=40$. Across all three parameters, μ , σ and ρ (i.e. σ^2) the comparison in terms of absolute bias and standard deviation reveals a small improvement in accuracy when the number of draws, N is increased from 5 to 20. However, further increasing N to 40 leads to negligible improvement. Since the increase in accuracy occurs in the third decimal place for model parameters we recommend setting $N=5$ in empirical applications.

These results are very encouraging. They demonstrate that the MCL estimator exhibits very satisfactory small sample performance which is directly comparable to the fully efficient Bayesian MCMC method. The evidence also suggests that these results can be achieved by using a very small number of draws.

4.2. S&P500 return series

The performance of the MCL estimator is further illustrated by fitting the model to the seasonally adjusted (Gallant, Rossi and Tauchen, 1992) S&P 500 returns. JPR and Danielsson (1994b) has already utilised a subset of the data (2/1/80-30/12/87, $T=2,022$ observations) to fit the basic SV model. Re-estimation by MCL allows not only for a comparison of the point estimates but also for the comparison of the computational requirement. The results of the estimation are reported in the table below:

				<i>time in min</i>
<i>MCL</i>	-0.00	0.96	0.16	1:21
<i>JPR's MCMC</i>	-0.00	0.97	0.15	7:15
<i>Danielsson's MCL</i>	-0.00	0.97	0.15	10:45

The parameter estimates are almost identical across the three estimation methods. The time requirement was calculated in the following manner.⁸ First, Danielsson's code was executed on our machine. The algorithm required 5:59 minutes to converge. On the other hand, starting at the same initial value, the MCL estimation method required 0:58 minutes to achieve convergence. The figure of 1:21 minutes in the above table for

⁶ The results on $N=5$ are reproduced from **Table 2**, also in greater detail.

⁷ More detailed estimation results are reported in **Table 7**.

⁸ We are grateful to Jon Danielsson for pointing out this procedure.

MCL was obtained by calibrating it to the time requirement reported in Danielsson (1994b). These results suggest that estimation via the MCL method is computationally more efficient than the MCMC and the alternative importance sampling methods.

Secondly, the basic SV model was fitted to the entire data-set (4/1/28-31/12/87, $T=16,127$), thus providing for a comparison with the SMM (Gallant and Tauchen, 1996). The parameter estimates are reported in **Table 5**. The resulting point estimates are very close to the SMM values obtained by Gallant, Hsieh and Tauchen (1995) except for the estimate of σ^2 . The sub-period analysis reported in the remainder of **Table 5** illustrates that the parameter estimates, and in particular the implied CV are not stable across the sub-periods. This may explain why Gallant, Hsieh and Tauchen (1995) found the SV model to be incapable of capturing the time series dynamics of the S&P500 index. It is not surprising that, when a data set of roughly 60 years of daily observations is used, some regime switches may be present.

This application also demonstrates that large data-sets present no difficulty for the estimation by MCL.

5. Further issues

Having shown that the MCL estimator exhibits satisfactory finite sample performance we would now like to turn to the practical issues in SV model estimation and indicate some of the interesting extensions of the basic SV model.

5.1. The inlier problem.

Since our method, as much as QML, relies on the use of the linear state space, taking the logarithms of squared mean adjusted returns becomes a problem when zero, or small values are encountered. In particular, if the drift in of the asset can be assumed to be zero ($\mu=0$) or prices are recorded discretely then it is possible that some returns will be zero. In many practical applications, however, equality of prices at two successive observations in time, leading to zero returns, arise due to data irregularities. For instance, properly accounting for holidays eliminates many “zero” returns in most daily exchange rate series. Deleting such observations from the sample eliminates the inlier problem. Alternatively, the updating equations of the Kalman filter can be modified so as to handle missing values (Harvey, 1989, pg. 143).

If an inlier cannot be assumed to be an irregular observation there are three alternatives of dealing with the problem. First, the sample mean of the series (or some general ARMA specification) may be subtracted from the observations. While the method may be feasible numerically (the resulting series may be devoid of entries identically equal to zero) it does not solve the problem conceptually. Secondly, Fuller (1996) suggests a transformation which mitigates the inlier problem by shifting some probability mass towards the centre of the distribution:

$$\ln^*(r_t^2) = \ln(r_t^2 + s_r^2) - \frac{s_r^2}{r_t^2 + s_r^2}$$

where α is some subjectively chosen constant, e.g. 0.02, and s_r^2 is the sample variance of returns r_t . Breidt and Carriquiry (1996) show that this transformation improves the

performance of the QML estimator and mitigates the inlier problem. This method, however, remains inefficient.

Finally, one may cut off the inliers by setting the observation at some value

$$\ln^*(r_t^2) = \ln(r_t^2 I_{\{r_t^2 \geq \epsilon\}}) + \ln(\epsilon^2 I_{\{r_t^2 < \epsilon\}}) \quad (10)$$

where $I_{\{\cdot\}}$ is the indicator function, and ϵ is a small number. Invariably, the choice of ϵ is subjective but it is demonstrated below that (10) leads to reasonably good MCL estimates for very small ϵ .

To assess the performance of the MCL and QML methods across various values of ϵ we designed the following Monte Carlo experiment. For the parameter triplet θ_0 we generated the basic SV model as before, except that the mean equation disturbances, ϵ_t in (1) have now a 10% chance of taking the value zero and 90% chance of being drawn from $N(0,1)$. It is rarely the case in practical applications that 10% of the sample are identically equal to zero but the experiment has been designed to illustrate the behaviour of the estimator in extreme situations. The generated series was then transformed according to (10) with $\ln(r_t^2)$ taking the values of -20, -30, -100, and -200. The results of the simulations for QML and MCL are presented in **Table 6** and compared to those of Section 4.

It is apparent that the performance of QML leaves much to be desired. The bias and the standard errors are very sensitive to the choice of ϵ . As ϵ decreases the performance deteriorates rapidly, leading to enormous biases in all three parameters. However, the decline in precision is not homogenous across the three model parameters. Interestingly, for small ϵ (e.g. $\epsilon = 3.72 \times 10^{-44}$) the bias in the estimate of the autoregressive parameter ρ disappears, while the biases in the estimates of σ^2 (i.e. σ^2) and β remain very large.

The results of the MCL estimator are considerably better. The bias and the standard errors on all three model parameters decrease with the cut off value ϵ . Comparison with the estimation results for the full sample, reveals that less precision can be achieved when 10% of observations are zero (the mean squared errors on all parameters are larger). This is not surprising, and results from the fact that the likelihood function is ‘flatter’ in cases when many zero observations are present.

5.2. Extensions of the basic SV model: Heavy Tails.

The unconditional density of many financial series exhibits larger kurtosis than can be captured by simply incorporating conditional heteroscedasticity into a Gaussian process. The basic SV model can be generalised so as to allow the mean equation disturbances, ϵ_t in (1) to follow a *Student-t* distribution with ν degrees of freedom. In this case the density of the transformed disturbances, $z_t = \ln(r_t^2)$ is given by:

$$p_{\ln r^2}(z) = C \left(1 + \frac{e^z}{\nu}\right)^{-\frac{\nu+1}{2}} e^{-\frac{z}{2}} \quad C = \frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\nu} \Gamma(\frac{\nu}{2})} \quad (3')$$

The limit of (3'), as $\nu \rightarrow \infty$, is of course, the $\ln(r_t^2)$ density (3) which can be verified by taking logarithms of (3), and expanding $\ln(1+x)$ as a Taylor series.

A suitable importance sampling density is found by equalising density slopes as described in section 3.1. The first derivative of the log density in (3) is:

$$d_1(z) = \frac{1}{2} \left\{ 1 - \left(1 + \frac{e^z}{1 + e^z} \right) \right\}$$

so that the updating equations for \tilde{H}_t become:


$$\tilde{H}_t = \frac{2 \hat{\sigma}_t^2}{e^{\hat{\sigma}_t^2} \left[\frac{1 + 1}{1 + e^{\hat{\sigma}_t^2}} \right] - 1} \quad t = 1, \dots, T \quad (6')$$

Again, the Gaussian equations (6) are obtained in the limit, as $\hat{\sigma}_t^2 \rightarrow \infty$ in (6'). Moreover, (6') automatically ensures the non-negativity of \tilde{H}_t , which can be verified by observing that the signs of the nominator and denominator are identical for any value of $\hat{\sigma}_t^2$. The computation of the MCL likelihood (5) involves the quantities $l(\hat{\sigma}_t^2)$ in (7) which are now constructed via:

$$l(\hat{\sigma}_t^2) = \frac{1}{2} \left(\ln 2 - \tilde{H}_t + \hat{\sigma}_t^2 + \frac{\hat{\sigma}_t^4}{\tilde{H}_t} + 2 \ln C - \left(1 + \frac{e^{\hat{\sigma}_t^2}}{1 + e^{\hat{\sigma}_t^2}} \right) \right) \quad (7')$$

where \tilde{H}_t is given by (6') after convergence. The number of degrees of freedom, ν , enters the parameter vector θ , over which the likelihood function is maximised.

To illustrate the validity of the method, we proceed to fitting the SV- t model to the S&P500 return series, results of which are presented in **Table 7**. For ease of reference the results of the basic SV are reproduced in the upper panel. The estimated number of degrees of freedom is 7.634, well in the range of empirical estimates reported by Bollerslev (1987) using the GARCH- t model: 6.211-13.889. The likelihood ratio test statistic takes the value 26.6 which is significant at the 1% level when compared to the relevant critical value of the χ^2 distribution. Similarly, the standard error on ν indicates the significance of this parameter. The introduction of the *Student-t* distributed mean equation disturbances reduces the value of the implied coefficient of variation, CV from 0.389 to 0.255. Intuitively, lower variance of the latent process is sufficient to account for the variability in the series.



⁹ This is different from the Harvey, Ruiz and Shephard's (1994) QML set-up where the variance of the measurement equation, H in the state space formulation is treated as a parameter.

Figure 5: S&P500: unconditional density and the density of the SV model

Finally, **Figure 5** demonstrates that the unconditional density of the S&P500 returns is closely approximated by the unconditional density from the estimated SV- t model.¹⁰ By contrast, the unconditional density of the basic SV model (with normal ϵ_t) does not capture as well the unconditional distribution of asset's returns. Thus the MCL estimator can be easily adjusted so as to incorporate heavy tailed distributions.

5.3. Extensions of the basic SV model(III): Explanatory variables.

The impact of exogenous explanatory variables on volatility has been examined in the context of the GARCH model by several authors (Baillie and Bollerslev, 1989; Lamoureux and Lastrapes, 1990, 1993). Such explanatory variables could be intervention dummies, seasonal components, or regressors like option implied volatility, trade volume data, etc. The empirical validity of the SV model with explanatory variables has been examined elsewhere (Ghysels and Jasiak, 1996; Board *et. al.*, 1997) but it deserves to be mentioned here that the MCL method need not be modified to handle this extension. More importantly, since the explanatory variables enter the state vector (see Appendix) the dimensionality of the optimisation problem is unchanged. For instance, the basic SV model with k explanatory variables requires the optimisation in only two directions, σ^2 and ω . This is very useful since multidimensional non-linear optimisation is a formidable task.

5.4. Extensions of the basic SV model(IV): Non-zero correlation.

When the linearising transformation - which transforms the basic SV model into the linear state space form - is applied, information regarding the correlation between the return and the (log)variance process is lost. In the context of QML estimation Harvey and Shephard (1996) show that this information can be recovered by conditioning on the signs of the returns. The augmented model takes the form:

$$\begin{cases} \ln r_t^2 = \ln \sigma^2 + h_t + \epsilon_t & \epsilon_t = \ln \frac{\sigma_t^2}{\sigma^2} \\ h_t = h_{t-1} + A s_t + \tilde{\epsilon}_t & \tilde{\epsilon}_t \sim N(0, \sigma^2 - A^2) \end{cases} \quad Cov(\epsilon_t, \tilde{\epsilon}_t) = B s_t \quad (11)$$

where $A = \sqrt{2 \rho^{-1}}$, $B = 1.1061 \rho$, s_t is the sign of r_t , and $\rho = Corr(\epsilon_t, \tilde{\epsilon}_t)$ is the correlation between the two original residuals in (1).

No modifications to the MCL estimation procedure are required since the filtering and smoothing algorithms - see Appendix - are all written for the correlated state space model. The correlation coefficient, ρ , enters the parameter vector $\theta = (\omega, \sigma^2, \rho)$ over which the likelihood function is optimised. Fitting the model to the CRSP data-set¹¹ used by Nelson (1991) and Harvey and Shephard (1996) gives the results in **Table 8**.

¹⁰ Parameter estimates of **Table 2.5** were used to draw two samples of the SV process the density of which is presented in the figure. The t -distributed random numbers were constructed in accordance with the Bailey (1994) algorithm.

¹¹ The data-set consists of the returns on the CRSP value-weighted US market index for the period 3/7/62-31/12/87 resulting in $T=6,409$ observations.

Not surprisingly, the MCL parameter estimates are not identically equal to the QML estimates reported by Harvey and Shephard (1996). However, they reflect the general pattern of high volatility persistence and large negative correlation between the return and the (log)variance processes. Both, the likelihood ratio test statistic and the standard error on $\hat{\rho}$ indicate the statistical significance of the correlation coefficient. The MCL estimates are less sensitive to the presence of correlation than the QML estimators. Omitting ρ does not significantly alter the estimates of the remaining parameters (ω , σ^2 , γ). Furthermore, the standard errors on all parameters estimates are smaller than the QML errors, reflecting the increased efficiency of the estimator.

6. Conclusion.

In this paper, the Monte Carlo likelihood (MCL) method of estimating stochastic volatility (SV) models is implemented successfully. The representation of the SV model is in a linear state space form so the Kalman filter can be employed to compute the Gaussian likelihood function via the prediction error decomposition. However, due to the log chi-square disturbances in the measurement equation of the SV model, the Gaussian likelihood will only make up a part of the true likelihood function. The MCL estimator proposed here approximates the remainder term via Monte Carlo simulation. As the number of simulations N increases, the approximation becomes more accurate. The finite sample performance of the MCL is examined in a simulation experiment. The results indicate full efficiency of the estimator across a range of possible parameter values even for very moderate simulation sizes such as $N = 5$.

Apart from the computational efficiency of the MCL method, we also have exploited the linear representation of the SV model. The state space formulation allows the SV model to be extended in a number of directions likely to arise in empirical research. We have examined and implemented the following extensions in detail: fat-tailed distribution of the mean equation disturbances, inclusion of explanatory variables and the non-zero correlation model. The illustrations have shown that all these extensions can be handled by the MCL straightforwardly. Finally, it is noticeable that these modifications do not require any substantial changes to the methodology of MCL or any changes to the Kalman filter smoother algorithms.

Appendix.

(i) The general univariate state space model (Harvey, 1989) is:

$$\begin{aligned} y_t &= Z_t' \alpha_t + \varepsilon_t & \alpha_t &\sim N(0, H_t) \\ \alpha_t &= T_t \alpha_{t-1} + c_t + \eta_t & \eta_t &\sim N(0, Q_t) \end{aligned} \quad \text{Cov}(\varepsilon_t, \eta_t) = G_t, \quad t = 1, \dots, T$$

where y_t is a scalar observation, α_t is the $(m \times 1)$ state vector, the covariance matrices H_t (1×1) and Q_t ($m \times m$) are non-singular, while the measurement and transition equation disturbances may be contemporaneously correlated with an $(m \times 1)$ non-zero covariance matrix G_t . In case of the SV model, the elements (or functions of the elements) of the parameter vector $\alpha = (\sigma^2, \rho)$ enter into the appropriate elements of Q_t , T_t , G_t and c_t . The long run volatility level, σ^2 (along with any explanatory variables) enters the state vector α_t which reduces the dimensionality of the nonlinear optimization problem of maximizing the likelihood function with respect to the parameter vector; see below. For instance, an correlated SV model with k explanatory variables, z_t^k grouped into $z_t = (z_t^1, \dots, z_t^k)'$ is put into the state space form by defining the matrices

$$\begin{aligned} \alpha_t &= \begin{bmatrix} h_t \\ \alpha_t \end{bmatrix}, \quad T_t = \begin{bmatrix} e_k & e_k' \\ e_k & I_k \end{bmatrix}, \quad Q_t = \begin{bmatrix} \sigma^2 - A^2 & e_k' \\ e_k & e_k e_k' \end{bmatrix}, \quad Z_t = [1 \quad z_t'] \\ H_t &= \frac{\sigma^2}{2}, \quad G_t = \begin{bmatrix} B s_t \\ e_k \end{bmatrix}, \quad c_t = \begin{bmatrix} A s_t \\ e_k \end{bmatrix}, \quad A = 0.7979, \quad B = 1.1061 \end{aligned}$$

where α_t is $(k \times 1)$, e_k is a $(k \times 1)$ vector of zeros, I_k is the $(k \times k)$ identity matrix, and s_t is the sign of the return at time t . The basic SV model is obtained as a special case by setting $\sigma^2 = 0$ and $k = 1$, $z_t^1 = 1$, $\forall t$.

The Kalman filter is given by:

$$\begin{aligned} \hat{\alpha}_t &= y_t - Z_t' a_t & F_t &= Z_t P_t Z_t' + H_t \\ K_t &= (T_{t+1}' P_t Z_t' + G_t) F_t^{-1} & L_t &= T_{t+1} - K_t Z_t \\ a_{t+1} &= T_{t+1} a_t + c_t + K_t' r_t & P_{t+1} &= T_{t+1}' P_t L_t' + Q_{t+1} - G_t K_t' \end{aligned}$$

for $t = 1, \dots, T$. The recursions are initialised with $\alpha_0 = N(a_0, P_0)$ where P_0 is the unconditional variance matrix of the state vector which may contain diffuse elements. The parameter estimates for α are obtained by numerically optimizing the Gaussian log likelihood function as given by

$$\ln L_G(\alpha) = -\frac{1}{2} \sum_{t=1}^T (\ln 2 + \ln |F_t| + \alpha_t' F_t^{-1} \alpha_t)$$

The estimate of $\ln \sigma^2$ is given by the relevant element of a_T while the standard errors are obtained from the relevant diagonal elements of P_T (Harvey, 1989, pg. 367).

(ii) The Kalman smoother (deJong, 1988; Koopman, 1993) is used to construct:

$$\hat{\alpha}_t = E(\alpha_t | y) = H_t e_t + G_t' r_t$$

$$C_t = \text{Var}(y_t | y) = H_t - H_t D_t H_t - G_t' N_t G_t + J_t + J_t'$$

where y is the vector of all observations, $J_t = H_t K_t' N_t G_t$, and the remaining quantities are obtained from the backwards recursions:

$$\begin{aligned} e_t &= F_t^{-1} v_t - K_t' r_t & D_t &= F_t^{-1} + K_t' N_t K_t \\ r_{t-1} &= Z_t' F_t^{-1} v_t + L_t' r_t & N_{t-1} &= Z_t' F_t^{-1} Z_t + L_t' N_t L_t \end{aligned}$$

for $t = T, \dots, 1$, with $r_T = 0$ and $N_T = 0$. The prediction errors, v_t , their variances, F_t , and the Kalman gain matrix K_t are outputs of the Kalman filter.

(iii) A special version of de Jong and Shephard's (1995) simulation smoother is used to give draws of (i) from $p_G(y, \cdot)$:

$$y_t^{(i)} = H_t \tilde{v}_t + G_t' \tilde{r}_t + u_t^{(i)}, \quad u_t^{(i)} \sim N(0, \tilde{C}_t)$$

where the quantities \tilde{v}_t and \tilde{C}_t are obtained from:

$$\begin{aligned} \tilde{v}_t &= F_t^{-1} v_t - K_t' \tilde{r}_t \\ \tilde{D}_t &= F_t^{-1} + K_t' \tilde{N}_t K_t \\ J_t &= H_t K_t' \tilde{N}_t G_t \\ M_t &= H_t (\tilde{D}_t Z_t - K_t' \tilde{N}_t T_{t+1}) + G_t' \tilde{N}_t L_t \\ \tilde{C}_t &= H_t - H_t \tilde{D}_t H_t - G_t' \tilde{N}_t G_t + J_t + J_t' \end{aligned}$$

and the backwards recursions:

$$\begin{aligned} \tilde{r}_{t-1} &= Z_t' F_t^{-1} v_t - M_t' \tilde{C}_t^{-1} u_t^{(i)} + L_t' \tilde{r}_t \\ \tilde{N}_{t-1} &= Z_t' F_t^{-1} Z_t + M_t' \tilde{C}_t^{-1} M_t + L_t' \tilde{N}_t L_t \end{aligned}$$

for $t = T, \dots, 1$ with $\tilde{r}_T = 0$ and $\tilde{N}_T = 0$. Note that when a set of samples is required, the Kalman filter and the recursions for $\tilde{D}_t, M_t, \tilde{C}_t, \tilde{N}_{t-1}$ need only be applied once since these quantities remain the same for each sample.

Table 1: Equalising density slopes: a recursive solution

	1	2	3	4	5	6	6	8	9
$d(1)$	3.13	3.18	3.22	3.22	3.24	3.26	3.28	3.28	3.29
$d(2)$	1.31	1.42	1.52	1.49	1.55	1.60	1.63	1.65	1.68
$d(3)$	0.26	0.28	0.29	0.30	0.30	0.31	0.31	0.31	0.31
$d(4)$	0.14	0.13	0.11	0.13	0.11	0.09	0.09	0.08	0.07
$d(5)$	0.05	0.04	0.03	0.04	0.03	0.02	0.02	0.02	0.01
$d(6)$	0.02	0.02	0.01	0.01	0.01	0.01	0.01	4e-3	3e-3
$d(7)$	0.01	0.01	0.00	0.01	3e-3	2e-3	1e-3	1e-3	7e-4
$d(8)$	5e-3	3e-3	1e-3	2e-3	1e-3	4e-4	3e-4	2e-4	2e-4
$d(9)$	2e-3	1e-3	4e-4	7e-4	3e-4	1e-4	8e-5	6e-5	3e-5
$d(10)$	1e-3	5e-4	1e-4	3e-4	1e-4	3e-5	2e-5	1e-5	8e-6

This table reports the values of the metric $d(k) = T^{-1} \sum_t |\tilde{H}_t^{(k)} - \tilde{H}_t^{(k-1)}|$ for $k=1, \dots, 10$ iterations of the simulated SV model with $T=1,000$ and across several parameter triplets $\theta_i = (\alpha_i, \beta_i, \gamma_i)$, the values of which are reported in **Table 2**. Small values of $d(k)$ indicate that the individual elements of the variance vector \tilde{H} are not changing considerably across further iterations, i.e. $\tilde{H}_i^{(k)} \rightarrow \bar{H}_i$.

Table 2: Comparison of MCMC, QML and MCL estimators

<i>CV=10</i>	<i>1</i>			<i>2</i>			<i>3</i>		
True	0.675	0.900	-0.821	0.484	0.950	-0.411	0.308	0.980	-0.164
QML	0.703 (0.17) (0.029)	0.884 (0.06) (0.003)	-0.821 (0.03) (0.001)	0.502 (0.13) (0.017)	0.938 (0.04) (0.002)	-0.410 (0.02) (0.001)	0.321 (0.10) (0.010)	0.970 (0.03) (0.001)	-0.164 (0.01) (0.000)
MCL	0.579 (0.07) (0.014)	0.915 (0.02) (0.001)	-0.837 (0.03) (0.001)	0.436 (0.06) (0.005)	0.953 (0.02) (0.000)	-0.417 (0.02) (0.001)	0.290 (0.05) (0.002)	0.977 (0.02) (0.000)	-0.166 (0.01) (0.000)
MCMC	0.562 (0.12) (0.027)	0.916 (0.03) (0.001)	-0.679 (0.22) (0.069)	0.460 (0.06) (0.004)	0.940 (0.02) (0.001)	-0.464 (0.16) (0.028)	0.350 (0.06) (0.005)	0.980 (0.01) (0.000)	-0.190 (0.08) (0.007)

<i>CV=1</i>	<i>4</i>			<i>5</i>			<i>6</i>		
True	0.363	0.900	-0.736	0.260	0.950	-0.368	0.166	0.980	-0.147
QML	0.417 (0.21) (0.049)	0.845 (0.18) (0.035)	-0.736 (0.02) (0.000)	0.302 (0.17) (0.030)	0.906 (0.18) (0.033)	-0.368 (0.01) (0.000)	0.203 (0.15) (0.025)	0.942 (0.16) (0.029)	-0.147 (0.01) (0.000)
MCL	0.325 (0.07) (0.006)	0.897 (0.10) (0.010)	-0.745 (0.02) (0.000)	0.233 (0.07) (0.003)	0.930 (0.10) (0.011)	-0.372 (0.01) (0.000)	0.161 (0.05) (0.002)	0.970 (0.07) (0.004)	-0.148 (0.01) (0.000)
MCMC	0.350 (0.07) (0.005)	0.880 (0.05) (0.003)	-0.870 (0.34) (0.134)	0.280 (0.07) (0.005)	0.920 (0.05) (0.003)	-0.560 (0.34) (0.152)	0.230 (0.08) (0.011)	0.970 (0.02) (0.001)	-0.220 (0.14) (0.025)

CV=0.1	7			8			9		
True	0.135	0.900	-0.706	0.096	0.950	-0.353	0.061	0.980	-0.141
QML	0.319 (0.31) <i>(0.132)</i>	0.350 (0.63) <i>(0.702)</i>	-0.706 (0.01) <i>(0.000)</i>	0.295 (0.30) <i>(0.131)</i>	0.420 (0.62) <i>(0.669)</i>	-0.353 (0.01) <i>(0.000)</i>	0.266 (0.30) <i>(0.133)</i>	0.449 (0.64) <i>(0.692)</i>	-0.141 (0.00) <i>(0.000)</i>
MCL	0.156 (0.11) <i>(0.012)</i>	0.443 (0.62) <i>(0.592)</i>	-0.709 (0.01) <i>(0.000)</i>	0.136 (0.10) <i>(0.012)</i>	0.526 (0.60) <i>(0.545)</i>	-0.355 (0.01) <i>(0.000)</i>	0.113 (0.10) <i>(0.013)</i>	0.572 (0.60) <i>(0.524)</i>	-0.142 (0.00) <i>(0.000)</i>
MCMC	0.150 (0.08) <i>(0.007)</i>	0.780 (0.19) <i>(0.051)</i>	-1.540 (1.35) <i>(2.518)</i>	0.120 (0.07) <i>(0.006)</i>	0.840 (0.16) <i>(0.038)</i>	-1.120 (1.15) <i>(1.911)</i>	0.140 (0.10) <i>(0.016)</i>	0.910 (0.12) <i>(0.019)</i>	-0.660 (0.83) <i>(0.958)</i>

This table reports the results of the simulation experiments. For each set of parameter triplets $\theta_i = (\alpha, \beta, \gamma)_i$, samples of length $T=500$ of the basic SV model are generated $K=500$ times. The model is then estimated by QML and MCL and the average estimated parameter values are presented. The standard deviations and mean squared errors (in italic) are reported in parenthesis below. The results for the MCMC estimator are reproduced from Jacquier, Polson and Rossi (1994), Table 7.

Table 3: Finite sample performance at sample length of T=2,000

<i>CV=1</i>	<i>4</i>			<i>5</i>			<i>6</i>		
<i>True</i>	<i>0.363</i>	<i>0.900</i>	<i>-0.736</i>	<i>0.260</i>	<i>0.950</i>	<i>-0.368</i>	<i>0.166</i>	<i>0.980</i>	<i>-0.147</i>
<i>QML</i>	0.381 (0.10) <i>(0.010)</i>	0.888 (0.05) <i>(0.003)</i>	-0.736 (0.01) <i>(0.000)</i>	0.268 (0.06) <i>(0.004)</i>	0.945 (0.03) <i>(0.001)</i>	-0.368 (0.01) <i>(0.000)</i>	0.167 (0.04) <i>(0.001)</i>	0.978 (0.01) <i>(0.000)</i>	-0.147 (0.00) <i>(0.000)</i>
<i>MCL</i>	0.317 (0.03) <i>(0.003)</i>	0.913 (0.02) <i>(0.000)</i>	-0.745 (0.01) <i>(0.000)</i>	0.239 (0.03) <i>(0.001)</i>	0.954 (0.01) <i>(0.000)</i>	-0.372 (0.01) <i>(0.000)</i>	0.158 (0.02) <i>(0.000)</i>	0.980 (0.01) <i>(0.000)</i>	-0.148 (0.00) <i>(0.000)</i>
<i>MCMC</i>	0.359 (0.03) <i>(0.001)</i>	0.896 (0.02) <i>(0.000)</i>	-0.762 (0.15) <i>(0.023)</i>						

For each set of parameter triplets $\theta_i = (\alpha_i, \beta_i, \gamma_i)$, samples of length $T=2,000$ of the basic SV model are generated $K=500$ times. The model is then estimated by QML and MCL and the average estimated parameter values are presented. The standard deviations and mean squared errors (in italic) are reported in parenthesis below. The results for the MCMC estimator are reproduced from Jacquier, Polson and Rossi (1994), Table 9.

Table 4: The Effect of varying the number of draws,

		δ		
	δ_0	0.26	0.95	-0.368
<i>MCL(N=5)</i>	$\bar{\delta}$	0.233	0.930	-0.372
	$se(\bar{\delta})$	(0.0664)	(0.1027)	(0.0126)
	$ \bar{\delta} - \delta_0 $	0.0272	0.0202	0.0456
	$mse(\bar{\delta})$	(0.0052)	(0.0110)	(0.0022)
<i>MCL(N=20)</i>	$\bar{\delta}$	0.256	0.936	-0.372
	$se(\bar{\delta})$	(0.0583)	(0.1003)	(0.0126)
	$ \bar{\delta} - \delta_0 $	0.0044	0.0142	0.0452
	$mse(\bar{\delta})$	(0.0034)	(0.0103)	(0.0022)
<i>MCL(N=40)</i>	$\bar{\delta}$	0.257	0.937	-0.372
	$se(\bar{\delta})$	(0.0490)	(0.0993)	(0.0126)
	$ \bar{\delta} - \delta_0 $	0.0029	0.0133	0.0452
	$mse(\bar{\delta})$	(0.0024)	(0.0100)	(0.0022)

This table reports the results of the simulation experiment on a single set of parameter values, δ . Samples of length $T=500$ of the basic SV model are simulated $K=500$ times and estimated by MCL. Values 5, 20, and 40 in parenthesis behind the MCL label signify the number of draws, N , employed by taking the expectation in equation (5). For each estimator the average parameters estimates, the sample standard deviations, the absolute bias and the mean squared error are reported in the rows labelled $\bar{\delta}$, $se(\bar{\delta})$, $|\bar{\delta} - \delta_0|$, $mse(\bar{\delta})$ respectively. The results for MCL($N=5$) are reproduced from **Table 2**

Table 5: S&P500 1928-1987: Estimation results

	a_1	r_w^2	\ln^{-2}	<i>CV</i>	<i>LogLik</i>
<i>SMM</i>	0.976	0.009	-9.362	0.210	-
<i>Full sample</i>	0.977	0.030	-9.367	0.952	-34,610
<i>Sub I</i>	0.984	0.039	-9.315	2.441	-8,689
<i>Sub II</i>	0.876	0.117	-9.454	0.654	-8,799
<i>Sub III</i>	0.970	0.034	-9.579	0.760	-8,554
<i>Sub IV</i>	0.989	0.010	-9.219	0.630	-8,512

This table reports the parameter estimates of the basic SV model when fitted to the adjusted daily observations on the S&P500 stock index level, 1928-87. The row labelled “SMM” reproduces the estimates of Gallant, Hsieh and Tauchen (1995), Table2, row1. The correspondence to the notation of the present paper is established by: a_1 , r_w^2 , $\ln^{-2}=2\ln(10^{-2}r_y)$. The factor 10^{-2} appears because we choose to work with percentage returns. The subsequent rows report the MCL estimates for the full sample (4/1/28-31/12/87, $T=16,127$), and four sub-samples of equal length, $T=4,030$: *Sub I* (4/1/28-7/7/41), *Sub II* (8/7/41-30/11/55), *Sub III* (1/12/55-7/1/72), *Sub IV*(10/1/72-21/12/87).

Table 6: Sensitivity of QML and MCL estimators w.r.t. cut-off value

	<i>True δ_5</i>	0.26	0.95	-0.368
<i>QML</i>	$\delta_1 = 4.54 \cdot 10^5$	0.966 (0.29) (0.584)	0.836 (0.13) (0.030)	-1.393 (1.14) (2.359)
	$\delta_2 = 3.06 \cdot 10^7$	0.888 (0.12) (0.409)	0.912 (0.09) (0.010)	-0.835 (0.88) (0.987)
	$\delta_3 = 1.93 \cdot 10^{22}$	0.848 (0.10) (0.356)	0.948 (0.04) (0.002)	-0.851 (0.65) (0.654)
	$\delta_4 = 3.72 \cdot 10^{44}$	0.849 (0.10) (0.358)	0.949 (0.06) (0.004)	-1.36 (1.73) (3.980)
<i>MCL</i>	$\delta_1 = 4.54 \cdot 10^5$	0.396 (0.14) (0.039)	0.869 (0.32) (0.109)	-0.845 (0.70) (0.715)
	$\delta_2 = 3.06 \cdot 10^7$	0.411 (0.13) (0.040)	0.867 (0.30) (0.097)	-0.854 (0.55) (0.537)
	$\delta_3 = 1.93 \cdot 10^{22}$	0.370 (0.12) (0.026)	0.906 (0.21) (0.046)	-0.601 (0.29) (0.136)
	$\delta_4 = 3.72 \cdot 10^{44}$	0.365 (0.11) (0.023)	0.910 (0.20) (0.042)	-0.575 (0.28) (0.121)

This table reports the results of the simulation experiment on a single set of parameter values, δ_5 . Samples of length $T=500$ of the basic SV model with 10% zero values are generated $K=500$ times and estimated by QML and MCL. Inliers are cut-off at δ_i where the cut-off constants $\delta_1 - \delta_4$ were chosen so as to correspond to $\ln(\delta_i^2) = -20, -30, -100, \text{ and } -200$.

Table 7: Estimates of the SV model with fat-tailed disturbances

	σ^2	\ln^{-2}	CV	$LogLik$	LR		
$\hat{\sigma}^2$	0.960	0.026	-9.313	-	0.389	-4,311.6	26.6
$s.e.(\hat{\sigma}^2)$	(0.018)	(0.009)	(0.094)	-	-	-	-
$\hat{\ln}^{-2}$	0.984	0.007	-9.498	7.634	0.255	-4,298.3	-
$s.e.(\hat{\ln}^{-2})$	(0.010)	(0.003)	(0.122)	(0.003)	-	-	-

This table reports the estimation results of the SV model where the mean equation disturbances follow a *Student-t* distribution with degrees of freedom. The data-set consists of daily observations on the S&P500 stock index level in the period 2/1/80-30/12/87. The return series are prefiltered to remove the calendar effects as documented in Gallant, Rossi and Tauchen (1992). The sample length is $T=2,022$ observations. The standard errors of (σ^2, \ln^{-2}) are obtained from the numerical approximation to the Hessian, while the standard errors of the estimate of \ln^{-2} are taken from the corresponding diagonal element of the state covariance matrix, P_T . The likelihood ratio test statistic follows the χ^2 distribution.

Table 8: Estimates of the SV model with non-zero correlation

	σ^2	\ln^{-2}	CV	$LogLik$	LR		
$\hat{\sigma}^2$	0.988	0.018	-10.017	-	1.11	-13,680	124
$s.e.(\hat{\sigma}^2)$	(0.003)	(0.003)	(0.136)	-	-	-	-
$\hat{\ln}^{-2}$	0.985	0.021	-9.924	-0.375	0.99	-13,618	-
$s.e.(\hat{\ln}^{-2})$	(0.003)	(0.003)	(0.109)	(0.004)	-	-	-

This table reports the estimation results of the SV model where the mean and variance equation disturbances are allowed to be correlated. The data-set consists of the returns on the CRSP value-weighted US market index for the period 3/7/62-31/12/87 used by Nelson (1991) and Harvey and Shephard (1996). The sample size is $T=6,409$ observations. The standard errors of (σ^2, \ln^{-2}) are obtained from the numerical approximation to the Hessian, while the standard errors of the estimate of \ln^{-2} are taken from the corresponding diagonal element of the state covariance matrix, P_T . The likelihood ratio test statistic follows the χ^2 distribution.

References

- Andersen, T.G. (1994), "Stochastic Autoregressive Volatility: A Framework for Volatility Modeling," *Mathematical Finance*, 4, 75-102.
- Andersen, T. G. and B. Sørensen (1996), "GMM Estimation of a Stochastic Volatility Model: a Monte Carlo Study," *Journal of Business and Economic Statistics*, 14, 328-52.
- Bailey, R.W. (1994), "Polar generation of random variates with the t-distribution," *Mathematics of Computation*, 62, 779-81.
- Baillie, R.T. and T. Bollerslev (1989), "The Message in Daily Exchange Rates: A Conditional Variance Tale," *Journal of Business and Economic Statistics*, 7, 297-305.
- Bera, A.K. and M.L. Higgins (1993), "ARCH Models: Properties, Estimation and Testing," *Journal of Economic Surveys*, 7, 305-66.
- Board, J.L.G., G. Sandmann and C.M.S. Sutcliffe (1997), "The Effect of Contemporaneous Futures Market Volume on Spot Market Volatility," mimeo, London School of Economics.
- Bollerslev, T. (1987), "A Conditional Heteroskedastic Time Series Model for Speculative Prices and Rates of Return," *Review of Economics and Statistics*, 69, 542-547.
- Bollerslev, T., R.Y. Chow and K.F. Kroner (1992), "ARCH Modelling in Finance: A Review of the Theory and Empirical Evidence," *Journal of Econometrics*, 52, 5-59.
- Bollerslev, T., R.F. Engle and D.B. Nelson (1993), "ARCH Models" in R.F. Engle and D.L. McFadden (eds.), *Handbook of Econometrics, Vol. IV*, Amsterdam: North-Holland.
- Breidt, F.J. and A.L. Carriquiry (1996), "Improved Quasi-Maximum Likelihood estimation for Stochastic Volatility Models," *Modelling and Prediction: Honouring Seymour Geisel*, A. Zellner and J.S. Lee (eds). New York: Springer Verlag.
- Danielsson, J. (1994a), "Stochastic Volatility in Asset Prices: Estimation with Simulated Maximum Likelihood," *Journal of Econometrics*, 61, 375-400.
- Danielsson, J. (1994b), "Comment on Jacquier, Polson and Rossi," *Journal of Business and Economic Statistics*, 12, 389-392.
- de Jong, P. (1988), "A cross-validation Filter for Time Series Models," *Biometrika*, 75, 594-600.
- de Jong, P. and N. Shephard (1995), "The Simulation Smoother for Time Series Models," *Biometrika*, 82, 339-350.
- Durbin, J. and S.J. Koopman (1997a), "Monte Carlo Maximum Likelihood Estimation for non-Gaussian State Space Models," *Biometrika*, forthcoming.
- Durbin, J. and S.J. Koopman (1997b), "Time Series Analysis for Non-Gaussian Observations Based on State Space Models," mimeo, London School of Economics.
- Fridman, M. and L. Harris (1996), "A Maximum Likelihood Approach for Non-Gaussian Stochastic Volatility Models," mimeo, University of Southern California.
- Fuller, W.A. (1996), *Introduction to Statistical Time Series*. 2nd Ed., New York: Wiley.

- Gallant, A.R., D. Hsieh and G. Tauchen (1995), "Estimation of Stochastic Volatility Models with Diagnostics," mimeo, Duke University.
- Gallant, A.R., P.E. Rossi and G. Tauchen (1992), "Stock Prices and Volume," *Review of Financial Studies*, 5, 199-242.
- Gallant, A.R. and G. Tauchen (1996), "Which moments to match," *Econometric Theory*, forthcoming.
- Ghysels, E., A.C. Harvey and E. Renault (1996), "Stochastic Volatility," in C.R. Rao and G.S. Maddala (eds.), *Statistical Methods in Finance* Amsterdam: North-Holland.
- Ghysels, E. and J. Jasiak (1996), "Stochastic Volatility and Time Deformation: An Application to Trading Volume and Leverage Effects," mimeo, Université de Montréal.
- Harvey, A.C. (1989), *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge: Cambridge University Press.
- Harvey, A.C. and N. Shephard (1993), "Estimation and Testing of Stochastic Variance Models," mimeo, London School of Economics.
- Harvey, A.C. and N. Shephard (1996), "Estimation of an Asymmetric Stochastic Volatility Model for Asset Returns," *Journal of Business and Economic Statistics*, 14, 429-34.
- Harvey, A.C., E. Ruiz and N. Shephard (1994), "Multivariate Stochastic Variance Models," *Review of Economic Studies*, 61, 247-64.
- Hull, J. and A. White (1987), "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, 42, 281-300.
- Jacquier, E., N.G. Polson and P.E. Rossi (1994), "Bayesian Analysis of Stochastic Volatility Models," *Journal of Business and Economic Statistics*, 12, 371-417 (with discussion)
- Jacquier, E., N.G. Polson and P.E. Rossi (1995), "Stochastic Volatility: Univariate and Multivariate Extensions," R.L. White Centre for Financial Research, DP #19-95. Wharton School, University of Pennsylvania.
- Kim, S., N. Shephard and S. Chib (1996), "Stochastic Volatility: Likelihood Inference and Comparison with ARCH models," mimeo, Nuffield College, Oxford.
- Kitagawa, G. (1987), "Non-Gaussian State Space Modelling of Nonstationary Time Series," *Journal of the American Statistical Association*, 82, 1032-63 (with discussion)
- Koopman, S.J. (1993), "Disturbance Smoother for State Space Models," *Biometrika*, 80, 117-126.
- Lamoureux, C.G. and W.D. Lastrapes (1990), "Heteroskedasticity in Stock Return Data: Volume versus GARCH Effects," *Journal of Finance*, 45, 221-29.
- Lamoureux, C.G. and W.D. Lastrapes (1993), "Forecasting Stock-Return Variance: Toward an Understanding of Stochastic Implied Volatilities," *Review of Financial Studies*, 6, 293-326.
- Melino, A. and S.M. Turnbull (1990), "Pricing Foreign Currency Options with Stochastic Volatility," *Journal of Econometrics*, 45, 239-65.

Nelson, D.B. (1991), "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica*, 59, 347-70.

Renault, E. and N. Touzi (1996), "Option hedging and Implied Volatilities in a Stochastic Volatility Model," *Mathematical Finance*, 6 279-302.

Shephard, N. (1996), "Statistical Aspects of ARCH and Stochastic Volatility," in: D.R. Cox, D.V. Hinkley and O.E. Barndorff-Nielsen, eds., *Time Series Models*, London, Chapman and Hall.

Shephard, N. and M.K. Pitt (1997), "Likelihood Analysis of Non-Gaussian Measurement-Time Series," *Biometrika*, forthcoming.

Taylor, S.J. (1986), *Modelling Financial Time Series*, Chichester: John Wiley.

Taylor, S.J. (1994), "Modelling Stochastic Volatility: a Review and Comparative Study," *Mathematical Finance*, 4, 183-204.