Realized Measures and Statistical Inference for Stochastic Volatility Models

Md. Nazmul Ahsan*

Abstract
Since realized measures of volatility are affected by measurement errors, the study considers a new class of discrete-time stochastic volatility (SV) models, which can relate many realized volatility measures to the latent conditional variance. We propose a hybrid estimator for this class of models that combines a generalized least square (GLS) type transformation and instrumental variable (IV) approach. A simulation study reveals that the hybrid estimation method has excellent finite-sample properties. We illustrate the proposed method's empirical relevance using mixed frequency IBM stock returns and options prices.

Key Words: Stochastic volatility, Realized volatility, High frequency data.

1. Introduction

Modelling the time-varying volatility of asset returns is one of the major problems of financial econometrics. To deal with such features, two main classes of parametric models have been proposed: (1) ARCH [Engle (1982)] and GARCH models [Bollerslev (1986)], where volatility is modelled as a deterministic function of past shocks; (2) stochastic volatility (SV) models [Taylor (1986)], where volatility is a latent stochastic process. Several studies have documented the superior performance of SV models over GARCH-type models for several reasons:

- SV models constitute discrete versions of continuous-time diffusion processes, which are widely used in the option-pricing literature; see Hull and White (1987), Taylor (1994), Shephard and Andersen (2009).
- SV models are flexible and relatively robust to model misspecification. GARCH models often require adding a random jump component or allowing for innovations with heavy-tailed distributions to tackle these problems. Such modifications substantially improve the performance of the standard GARCH, but do not appear to be required for SV models; see Carnero et al. (2004), Chan and Grant (2016).
- SV models perform better than GARCH-type models in volatility forecasting, which suggests that time-varying volatility is better modelled as a latent first-order autoregression; see Kim et al. (1998), Yu (2002), Poon and Granger (2003), Koopman et al. (2005).

Despite their appealing features, statistical inference is challenging in SV models due to the inherent problem of evaluating the likelihood function. The marginal likelihood of SV models is given by a high dimensional integral, which makes the estimation by conventional maximum likelihood (ML) infeasible. This is a general feature of most nonlinear latent variable models because the latent variables must be integrated out of the joint density for the observed and latent processes, leading to an integral of high dimensionality.

*McGill University and Centre interuniversitaire de recherche en analyse des organisations (CIRANO), Mailing address: Department of Economics, McGill University, Leacock Building, 855 Sherbrooke Street West, Montreal, Quebec H3A 2T7, Canada; e-mail: md.ahsan@mail.mcgill.ca.

As a result, a variety of alternative methods have been proposed to estimate SV models. Major references include: the Quasi-Maximum Likelihood (QML) [Harvey et al. (1994); Ruiz (1994)], the Generalized Method of Moments (GMM) [Melino and Turnbull (1990); Andersen and Sørensen (1996)], the Efficient Method of Moments (EMM) [Gallant and Tauchen (1996); Andersen et al. (1999)], the Maximum Likelihood Monte Carlo (MLMC) [Sandmann and Koopman (1998)], the Simulated Maximum Likelihood (SML) [Danielsson and Richard (1993); Danielsson (1994); Durham (2006); Liesenfeld and Jung (2000); Richard and Zhang (2007)], method base on linear-representation (LiR) [Francq and Zakoïan (2006)], the closed-form moment-based estimator (DV) [Dufour and Valéry (2006)], the ARMA-based winsorized estimator (W-ARMA-SV) [Ahsan and Dufour (2019)] and Bayesian methods based on Markov Chain Monte Carlo (MCMC) methods [Jacquier et al. (1994), Kim et al. (1998), Chib et al. (2002), Flury and Shephard (2011)]. For a review of the SV literature; see Ghysels et al. (1996), Broto and Ruiz (2004), Shephard (2005), Ahsan and Dufour (2020).

This paper considers a new class of discrete-time SV models, which is proposed by Ahsan (2020). This class of models extends the SV model's usual state-space representation by adding an additional measurement equation. This additional measurement equation is viewed as an instrument equation, and it can relate many realized volatility measures to the latent log volatility. In this paper, we propose a hybrid estimator for this class of models. The hybrid method combines a generalized least square (GLS) type transformation and instrumental variable (IV) approach.

We use RV measures as instruments for the latent volatility, in contrast with recent studies, where RV has been incorporated in traditional volatility models (GARCH or SV) by adding a measurement equation that connects the low-frequency volatility measure and realized volatility, these are: (1) Realized SV [Takahashi et al. (2009), Koopman and Scharth (2012)], (2) Multiplicative Error Model [Engle and Gallo (2006)], (3) HEAVY model [Shephard and Sheppard (2010), Noureldin et al. (2012)], (4) Realized GARCH [Hansen et al. (2012)].

We present some simulation evidence on the performance of the Hybrid estimator. The hybrid estimation method has excellent finite-sample properties in terms of bias and root mean square error. Finally, we illustrate the hybrid inference method's empirical relevance using mixed frequency IBM stock returns and options prices.

2. Framework

The process $\{s_t : t \in \mathbb{N}_0\}$ follows an SV model of the type:

$$s_t = \sigma_t z_t, \tag{1}$$

$$\log(\sigma_t^2) = \mu + \phi \log(\sigma_{t-1}^2) + v_t, \tag{2}$$

where s_t is the return observed at time t, and σ_t is the corresponding volatility. $\mathbb{N}_0$ refers to the non-negative integers. The z_t's and v_t's, are i.i.d. $\mathrm{N}(0,1)$ and $\mathrm{N}(0,\sigma_v^2)$ random variables, respectively and ϕ, μ, σ_v are the fixed parameters of the model. Further, the process $l_t = \big(s_t, \log(\sigma_t^2)\big)'$ is strictly stationary.

The linear state space representation for the above SV model can be written as follows

$$\text{State Transition Equation:} \quad w_t = \mu + \phi w_{t-1} + v_t \tag{3}$$

$$\text{Measurement Equation:} \quad y_t = w_t + \epsilon_t \tag{4}$$

where

$$y_t := \log(s_t^2) - \mathbb{E}\left[\log(z_t^2)\right], \quad w_t := \log(\sigma_t^2), \tag{5}$$

and

$$\epsilon_t := \log(z_t^2) - \mathbb{E}\left[\log(z_t^2)\right]. \tag{6}$$

Under the standard normality assumption for z_t, the transformed errors ϵ_t are i.i.d. according to the distribution of a centered $\log(\chi^2_{(1)})$ random variable with $\mathbb{E}\left[\log(z_t^2)\right] \simeq -1.2704$, $\sigma_\epsilon^2 := \mathbb{E}[\epsilon_t^2] = \text{Var}\left(\log(z_t^2)\right) = \pi^2/2$ and $\mathbb{E}[\epsilon_t^4] = \pi^4 + 3\sigma_\epsilon^2$ [see Abramowitz and Stegun (1970)].

It is evident from (3)-(4) that using any proxy for latent volatility (*e.g.*, replacing w_t by y_t) will induce a measurement error problem. Further, the latent volatility process introduces a moving average of measurement errors. We could alleviate this type of problem by using an IV regression where we replace the unobserved variables by their proxies. In below, we introduce a new class of stochastic volatility models, where the instrument equation can relate many realized volatility measures $\bar{Z}_{t-2}$ to the latent log volatility w_{t-1}. The process $\{y_t : t \in \mathbb{N}_0\}$ satisfies the following equations:

$$\begin{aligned}
\text{State Transition Equation:} \quad w &= \phi w_{-1} + X\beta + v && (7)\\
\text{Measurement Equation:} \quad y &= w + \epsilon && (8)\\
\text{Instrument Equation:} \quad w_{-1} &= \bar{Z}_{-2}\bar{\pi} + u_{-1} && (9)
\end{aligned}$$

where $w = (w_1, \ldots, w_T)'$, $w_{-1} = (w_0, \ldots, w_{T-1})'$, $y = (y_1, \ldots, y_T)'$ are $T \times 1$ vector, $X = [X_1', \ldots, X_T']'$ is a $T \times k$ matrix of exogenous explanatory variables which may predict the latent volatility as well as capture the leverage effect, $\bar{Z}_{-2} = [\bar{Z}_{-1}', \ldots, \bar{Z}_{T-2}']'$ is a $T \times m$ matrix of of variables related to w_{-1}, while $\epsilon = (\epsilon_1, \ldots, \epsilon_T)'$, $v = (v_1, \ldots, v_T)'$, $u_{-1} = (u_0, \ldots, u_{T-1})'$ are $T \times 1$ vector of disturbances. The matrices of unknown coefficients ϕ, β , and $\bar{\pi}$ have dimensions respectively 1×1, $k \times 1$, and $m \times 1$. Note that model (7)-(8) with $X\beta = \mu\mathbf{1}$ corresponds to a standard SV model.

3. A Hybrid Estimator

In this section, we propose a hybrid estimator for the model given in (7)-(9). We assume that the ϵ_t's and v_t's are i.i.d. $\text{N}(0, \sigma_\epsilon^2)$ and $\text{N}(0, \sigma_v^2)$ random variables, and $X\beta = \mu\mathbf{1}$. Substituting (8) into (7), we have:

$$y = \mu + \phi y_{-1} + v + \epsilon - \phi\epsilon_{-1} = \mu + \phi y_{-1} + \xi \tag{10}$$

where $\xi := v + \epsilon - \phi\epsilon_{-1}$ is an MA(1) process with $\xi \sim \text{N}\left[0, \sigma_\xi^2\Sigma(\rho)\right]$ where

$$\Sigma(\rho) := \begin{pmatrix}
1 & -\rho & 0 & \cdots & \cdots & \cdots & \cdots & 0\\
-\rho & 1 & -\rho & 0 & & & & \vdots\\
0 & -\rho & 1 & -\rho & \ddots & & & \vdots\\
\vdots & 0 & \ddots & \ddots & \ddots & \ddots & & \vdots\\
\vdots & & \ddots & \ddots & \ddots & \ddots & 0 & \vdots\\
\vdots & & & \ddots & -\rho & 1 & -\rho & 0\\
\vdots & & & & 0 & -\rho & 1 & -\rho\\
0 & \cdots & \cdots & \cdots & \cdots & 0 & -\rho & 1
\end{pmatrix}, \tag{11}$$

$$\sigma_\xi^2 := (1 + \phi^2)\sigma_\epsilon^2 + \sigma_v^2, \quad \rho := \frac{-\text{Cov}(\xi_t\xi_{t-1})}{\text{Var}(\xi_t)} = \frac{\phi\sigma_\epsilon^2}{(1 + \phi^2)\sigma_\epsilon^2 + \sigma_v^2}. \tag{12}$$

Clearly, ρ is a function of ϕ, σ_v^2, and σ_ϵ^2. $\Sigma(\rho)$ is a *Toeplitz* matrix (or diagonal-constant matrix) with dimension $T \times T$. Because $\Sigma(\rho)$ is a symmetric positive-definite matrix, there exists a $T \times T$ matrix C, such that $C\Sigma(\rho)C' = I_T$.

Table 1: Comparison of different estimation methods with respect to bias and RMSE for the SV model using simulated data. Bayes is the Bayesian estimator based on Markov Chain Monte Carlo methods proposed by Jacquier et al. (1994). We used R package *stochvol* of Kastner (2016) for the Bayesian estimation. Hybrid is the simple estimator proposed in Section 3.

	$T = 1000$			$T = 2000$		
	ϕ	μ	σ_v	ϕ	μ	σ_v
True value	0.98	-0.25	0.25	0.98	-0.25	0.25
Bias						
Bayes	-0.006	-0.028	0.013	-0.003	0.040	0.005
Hybrid ($l = 1$)	-0.011	0.141	0.125	-0.007	0.154	0.114
Hybrid ($l = 3$)	-0.011	0.141	0.126	-0.007	0.154	0.112
Hybrid ($l = 5$)	-0.013	0.133	0.142	-0.007	0.154	0.114
Hybrid ($l = 10$)	-0.011	0.139	0.129	-0.006	0.156	0.111
RMSE						
Bayes	0.011	0.388	0.037	0.007	0.293	0.025
Hybrid ($l = 1$)	0.016	0.148	0.144	0.012	0.158	0.132
Hybrid ($l = 3$)	0.016	0.148	0.146	0.012	0.158	0.129
Hybrid ($l = 5$)	0.020	0.143	0.163	0.012	0.158	0.132
Hybrid ($l = 10$)	0.017	0.146	0.148	0.011	0.159	0.128

Since $\xi = v + \epsilon - \phi\epsilon_{-1} = \tilde{\epsilon}_t - \theta\tilde{\epsilon}_{-1} \sim \mathrm{N}\big[0, \sigma_\xi^2\Sigma(\rho)\big]$, we can fit an ARMA(1,1) model in (10) and obtain an estimate of θ. It is easy to see that

$$\hat{\rho} = \frac{-\hat{\theta}}{1 + \hat{\theta}^2},$$

where $\hat{\theta}$ is the estimated MA average parameter. Given $\hat{\rho}$, we can have $C\Sigma(\hat{\rho})C' = I_T$ and we can consider the following transformed model:

$$Cy = \mu C\mathbf{1} + \phi Cy_{-1} + C\xi \tag{13}$$

$$Cy_{-1} = \alpha C\mathbf{1} + \phi CZ_{-2} + C\eta_{-1} \tag{14}$$

where the variance-covariance matrix of $\xi^* := C\xi$ is now an i.i.d. $\mathrm{N}(0, \sigma_\xi^2 I_T)$ distribution and $\eta_{-1} := \epsilon_{-1} + u_{-1}$.

Given (13)-(14), an IV/2SLS estimator of $\hat{\beta}_H = (\hat{\mu}, \hat{\phi})'$ is as follows:

$$\hat{\beta}_H = (Z'X)^{-1}Z'Y,$$

where $X = Cy_{-1}, Y = Cy$ and $Z = C\bar{Z}_{-2}$. We are also interested in volatility of volatility innovation parameter, which can be obtain as follows:

$$\hat{\sigma}_v^2 = \frac{\pi^2}{-2\hat{\rho}}[\hat{\phi} + (1 + \hat{\phi}^2)\hat{\rho}].$$

The asymptotic theory for the hybrid estimator is standard. However, for inference we may consider the parametric bootstrap; see Andrews (1997), Bai (2003).

4. Simulation Study

In this section, we compare the statistical performance of the proposed hybrid estimator with the Bayesian estimator based on the MCMC technique. We consider a standard SV model where parameter values of (ϕ, μ, σ_v) are $M = (0.98, -0.25, 0.25)$. The parameters

Table 2: Hybrid estimates with different instruments. Ticker: IBM, January 2009 - December 2013, $T = 1258$. The instrument set consists of a constant and lags of an instrument, $l = 1,\ 3$. The average precision of an instrument set i over the proposed inference methods is measured by $\bar{d}_{i,s} := S^{-1}\sum_{i=1}^{S} d_i$, where $s \in S$ and S is the set of identification-robust inference methods, and $d_i := 1 - (ub_i - lb_i)$, where ub and lb are the upper and lower bound of the confidence set, and $ub - lb$ is the length of the confidence set.

	$l = 1$				$l = 3$			
Instruments	$\bar{d}_{i,s}$	$\hat{\phi}$	$\hat{\mu}$	$\hat{\sigma}_v$	$\bar{d}_{i,s}$	$\hat{\phi}$	$\hat{\mu}$	$\hat{\sigma}_v$
RSVN-5m-ss	0.8860	0.9896	-0.0003	0.1830	0.8618	0.9906	-0.0003	0.1937
ImV-C-mean	0.8830	0.9972	-0.0005	0.1851	0.8218	0.9936	-0.0004	0.1946
MinRV-5m	0.8828	0.9885	-0.0003	0.1826	0.8493	0.9896	-0.0003	0.1934
RV-5m-ss	0.8825	0.9884	-0.0003	0.1826	0.8560	0.9902	-0.0003	0.1936
BV-5m-ss	0.8823	0.9884	-0.0003	0.1826	0.8508	0.9898	-0.0003	0.1934
MedRV-5m	0.8823	0.9882	-0.0002	0.1825	0.8493	0.9895	-0.0003	0.1933
1-day	0.4255	0.9695	0.0002	0.1696	0.7490	0.9672	0.0002	0.1782

were selected to represent values often found in empirical applications of hourly or daily returns. The simulations use 50 replications and we present results for two different sample sizes ($T = 1000, 2000$). In simulations, the hybrid estimator uses past lags as instruments ($l = 1, 3, 5, 10$).

Table 1 reports the estimation results for model M. From this table, we see that the hybrid method yields the smallest RMSE for μ and the Bayesian estimation yields the smallest RMSEs for ϕ and σ_v. However, the hybrid estimator produces competitive parameter estimates in cases of ϕ and σ_v. Further, the hybrid method is consistent across different instrument sets and yields similar results. The results for the two sample sizes are qualitatively similar ($T = 1000, 2000$) and indicate that estimator precision increases with the sample size.

5. Applications to Stock Price Volatilities

We apply our proposed estimator to IBM's price and option data (2009-2013, 1258 trading days). The low-frequency daily prices are obtained from the CRSP database. The raw series p_t is converted to returns by the transformation $r_t := 100[\log(p_t) - \log(p_{t-1})]$ and the returns are converted to residual returns by $s_t := r_t - \hat{\mu}_r$, where $\hat{\mu}_r$ is the sample average of returns. The daily volatility proxy is constructed by the transformation $y_t = \log(s_t^2) + 1.2704$. IBM's tick price data are taken from the TAQ (Trade and Quote) database and option (American) data are sourced from the OptionMetrics database. Using these data, we construct seven different instruments, these include: RSVN-5m-ss (5-minute realized negative semivariance with 1-minute subsampling, Barndorff-Nielsen et al. (2010)), ImV-C-mean (average implied volatility extracted from call options), MinRV-5m (5-minute minimum RV, Andersen et al. (2012)), RV-5m-ss (5-minute realized volatility with 1-minute subsampling, Andersen et al. (2001)), BV-5m-ss (5-minute bipower variation with 1-minute subsampling, Barndorff-Nielsen and Shephard (2004)), MedRV-5m (5-minute median RV, Andersen et al. (2012)), 1-day (1-day realized volatility). In below, we analyze these instruments' ability to describe the low-frequency volatility.

For empirical analysis, we consider the following model:

$$w_t = \mu + \phi w_{t-1} + v_t\,, \quad y_t = w_t + \epsilon_t\,, \tag{15}$$

$$y_{t-1} = \bar{\pi}_0 + Z_{t-2}'\bar{\pi}_1 + \eta_{t-1}\,, \quad \eta_{t-1} := \epsilon_{t-1} + u_{t-1}\,, \tag{16}$$

$$v_t \sim \text{i.i.d. } N(0, \sigma_v^2)\,, \quad \epsilon_t \sim \text{N}(0, \sigma_\epsilon^2)\,, \quad u_t \sim \text{i.i.d. } N(0, \sigma_u^2)\,, \tag{17}$$

where $w_t = \log(\sigma_t^2)$, $y_t = \log(s_t^2) + 1.2704$ with $s_t := r_t - \mu_r$ is residual return of an asset with μ_r is the mean of return $r_t = 100[\log(p_t) - \log(p_{t-1})]$ and Z_{t-2} is the set of IV's. Note that the constant term $\bar{\pi}_0$ captures the bias in the RV estimate due to the non-trading hours and microstructure noise. If the bias-correction term $\bar{\pi}_0$ is negative, RV has an upward bias that may be due to the market microstructure noise; see Takahashi et al. (2009). Further, if $\bar{\pi}_0$ is positive, RV has a downward bias due to the non-trading hours.

As pointed out by Dufour (1997), when IV's are arbitrary weak, then confidence sets with correct coverage probability must have an infinite length with positive probability. As a result, the length of a weak instrument robust confidence interval can summarize the identification strength of the corresponding instrument. If we restrict $\phi \in [0,\ 1]$, then an irrelevant (no identification) instrument for the regressor should produce a confidence interval with length equal to 1. In our context, using the set of identification-robust confidence intervals, Ahsan (2020) define the notion of the average precision of an instrument set i over the identification-robust inference methods $[AR$, AR^*, SS, $SS^*]$ by

$$\bar{d}_{i,s} := \frac{1}{S} \sum_{i=1}^{S} d_i \tag{18}$$

where $s \in S$ and S is the set of identification-robust inference methods, and $d_i := 1 - (ub_i - lb_i)$, where ub and lb are the upper and lower bound of the confidence set, and $ub - lb$ is the length of the confidence set. The definition $\bar{d}_{i,s}$ implies that if i is a weak instrument then it will produce $\bar{d}_{i,s}$ close to 0 and if i is a strong instrument then it will produce $\bar{d}_{i,s}$ close to 1. For example, a large value of $\bar{d}_{i,s}$ implies that the corresponding instrument set is highly informative about the parameter ϕ. For further details about the the identification-robust inference methods $[AR$, AR^*, SS, $SS^*]$, see Ahsan (2020).

We estimate the model given in (15)-(17) using the proposed hybrid estimator. In Table 2, we report the estimated parameters, where the instrument set includes a constant and several lags of an instrument, $l = 1, 3$. Several conclusions emerge from the results. *First*, in most cases, we find that the estimates of ϕ are close to unity, implying that the volatility process is highly persistent. *Second*, all high-frequency instruments are highly informative compare to the daily instrument. These high-frequency instruments produce similar parameter estimates. This is consistent even when $l = 3$.

6. Conclusion

This paper has proposed a computationally simple hybrid estimator for a class of SV models, which can utilize information from high-frequency data. Compared with existing alternative procedures for a standard SV model, the proposed estimators enjoy a considerable advantage in computation time and match the standard Bayesian estimator in terms of bias and RMSE. Due to its simplicity, the hybrid estimators allow one to build reliable simulation-based inference for SV models.

We fitted the SV model using our hybrid estimator to IBM stock return time series, using various instruments. We found that the volatility process of IBM is highly persistent and close to the unit root. High-frequency realized volatility measures are more informative compared to the low-frequency volatility measures.

References

Abramowitz, M. and Stegun, I. A. (1970), *Handbook of Mathematical Functions*, Vol. 55, Dover Publications Inc., New York.

Ahsan, M. N. (2020), High-frequency instruments and identification-robust inference for stochastic volatility models, Technical report, McGill University.

Ahsan, M. N. and Dufour, J.-M. (2019), A simple efficient moment-based estimator for the stochastic volatility model, *in* I. Jeliazkov and J. Tobias, eds, 'Topics in Identification, Limited Dependent Variables, Partial Observability, Experimentation, and Flexible Modeling', Vol. 40 of *Advances in Econometrics*, Emerald, Bingley, U.K., pp. 157–201.

Ahsan, M. N. and Dufour, J.-M. (2020), Simple estimators and inference for higher-order stochastic volatility models, Technical report, McGill University.

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P. (2001), 'The distribution of realized exchange rate volatility', *Journal of the American statistical association* **96**(453), 42–55.

Andersen, T. G., Chung, H.-J. and Sørensen, B. E. (1999), 'Efficient method of moments estimation of a stochastic volatility model: A Monte Carlo study', *Journal of Econometrics* **91**, 61–87.

Andersen, T. G., Dobrev, D. and Schaumburg, E. (2012), 'Jump-robust volatility estimation using nearest neighbor truncation', *Journal of Econometrics* **169**(1), 75–93.

Andersen, T. G. and Sørensen, B. E. (1996), 'GMM estimation of a stochastic volatility model: A Monte Carlo study', *Journal of Business and Economic Statistics* **14**(3), 328–352.

Andrews, D. W. (1997), 'A conditional kolmogorov test', *Econometrica* **65**(5), 1097–1128.

Bai, J. (2003), 'Testing parametric conditional distributions of dynamic models', *Review of Economics and Statistics* **85**(3), 531–549.

Barndorff-Nielsen, O. E., Kinnebrock, S. and Shephard, N. (2010), Measuring downside risk—realised semivariance, *in* T. Bollerslev, J. Russell and M. Watson, eds, 'Volatility and Time Series Econometrics: Essays in Honour of Robert F. Engle', Oxford University Press, Oxford, U.K., pp. 117–136.

Barndorff-Nielsen, O. E. and Shephard, N. (2004), 'Power and bipower variation with stochastic volatility and jumps', *Journal of Financial Econometrics* **2**(1), 1–37.

Bollerslev, T. (1986), 'Generalized autoregressive conditional heteroskedasticity', *Journal of Econometrics* **31**, 307–327.

Broto, C. and Ruiz, E. (2004), 'Estimation methods for stochastic volatility models: a survey', *Journal of Economic Surveys* **18**(5), 613–649.

Carnero, M. A., Peña, D. and Ruiz, E. (2004), 'Persistence and kurtosis in GARCH and stochastic volatility models', *Journal of Financial Econometrics* **2**(2), 319–342.

Chan, J. C. and Grant, A. L. (2016), 'Modeling energy price dynamics: GARCH versus stochastic volatility', *Energy Economics* **54**, 182–189.

Chib, S., Nardari, F. and Shephard, N. (2002), 'Markov chain Monte Carlo methods for stochastic volatility models', *Journal of Econometrics* **108**, 281–316.

Danielsson, J. (1994), 'Stochastic volatility in asset prices estimation with simulated maximum likelihood', *Journal of Econometrics* **61**, 375–400.

Danielsson, J. and Richard, J.-F. (1993), 'Accelerated Gaussian importance sampler with application to dynamic latent variable models', *Journal of Applied Econometrics* **8**, S153–S173.

Dufour, J.-M. (1997), 'Some impossibility theorems in econometrics, with applications to structural and dynamic models', *Econometrica* **65**, 1365–1389.

Dufour, J.-M. and Valéry, P. (2006), 'On a simple two-stage closed-form estimator for a stochastic volatility in a general linear regression', *Advances in Econometrics* **20**, 259–288.

Durham, G. B. (2006), 'Monte Carlo methods for estimating, smoothing, and filtering one- and two-factor stochastic volatility models', *Journal of Econometrics* **133**, 273–305.

Engle, R. F. (1982), 'Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation', *Econometrica* **50**(4), 987–1008.

Engle, R. F. and Gallo, G. M. (2006), 'A multiple indicators model for volatility using intra-daily data', *Journal of Econometrics* **131**(1), 3–27.

Flury, T. and Shephard, N. (2011), 'Bayesian inference based only on simulated likelihood: particle filter analysis of dynamic economic models', *Econometric Theory* **27**(05), 933–956.

Francq, C. and Zakoïan, J.-M. (2006), 'Linear-representation based estimation of stochastic volatility models', *Scandinavian Journal of Statistics* **33**(4), 785–806.

Gallant, A. R. and Tauchen, G. (1996), 'Which moments to match?', *Econometric Theory* **12**, 657 – 681.

Ghysels, E., Harvey, A. and Renault, E. (1996), Stochastic volatility, *in* G. S. Maddala and C. R. Rao, eds, 'Handbook of Statistics: Statistical Methods in Finance', Vol. 14, North-Holland, Amsterdam, pp. 119–191.

Hansen, P. R., Huang, Z. and Shek, H. H. (2012), 'Realized GARCH: a joint model for returns and realized measures of volatility', *Journal of Applied Econometrics* **27**(6), 877–906.

Harvey, A., Ruiz, E. and Shephard, N. (1994), 'Multivariate stochastic variance models', *The Review of Economic Studies* **61**, 247–264.

Hull, J. and White, A. (1987), 'The pricing of options on assets with stochastic volatilities', *The Journal of Finance* **42**(2), 281–300.

Jacquier, E., Polson, N. and Rossi, P. (1994), 'Bayesian analysis of stochastic volatility models (with discussion)', *Journal of Economics and Business Statistics* **12**, 371–417.

Kastner, G. (2016), 'Dealing with stochastic volatility in time series using the R package stochvol', *Journal of Statistical Software* **69**(5), 1–30.

Kim, S., Shephard, N. and Chib, S. (1998), 'Stochastic volatility: Likelihood inference and comparison with ARCH models', *The Review of Economic Studies* **65**, 361–393.

Koopman, S. J., Jungbacker, B. and Hol, E. (2005), 'Forecasting daily variability of the S&P 100 stock index using historical, realised and implied volatility measurements', *Journal of Empirical Finance* **12**(3), 445–475.

Koopman, S. J. and Scharth, M. (2012), 'The analysis of stochastic volatility in the presence of daily realized measures', *Journal of Financial Econometrics* **11**(1), 76–115.

Liesenfeld, R. and Jung, R. C. (2000), 'Stochastic volatility models: conditional normality versus heavy-tailed distributions', *Journal of Applied Econometrics* **15**(2), 137–160.

Melino, A. and Turnbull, S. M. (1990), 'Pricing foreign currency options with stochastic volatility', *Journal of Econometrics* **45**(1), 239–265.

Noureldin, D., Shephard, N. and Sheppard, K. (2012), 'Multivariate high-frequency-based volatility (HEAVY) models', *Journal of Applied Econometrics* **27**(6), 907–933.

Poon, S.-H. and Granger, C. W. (2003), 'Forecasting volatility in financial markets: A review', *Journal of Economic Literature* **41**(2), 478–539.

Richard, J.-F. and Zhang, W. (2007), 'Efficient high-dimensional importance sampling', *Journal of Econometrics* **141**(2), 1385–1411.

Ruiz, E. (1994), 'Quasi-maximum likelihood estimation of stochastic variance models', *Journal of Econometrics* **63**, 284–306.

Sandmann, G. and Koopman, S. J. (1998), 'Estimation of stochastic volatility models via Monte Carlo maximum likelihood', *Journal of Econometrics* **87**(2), 271–301.

Shephard, N. (2005), *Stochastic volatility: selected readings*, Oxford University Press, Oxford, U.K.

Shephard, N. and Andersen, T. G. (2009), Stochastic volatility: Origins and overview, *in* T. Mikosch, J.-P. Kreiß, A. R. Davis and G. T. Andersen, eds, 'Handbook of Financial Time Series', Springer-Verlag, Berlin, Heidelberg, pp. 233–254.

Shephard, N. and Sheppard, K. (2010), 'Realising the future: forecasting with high-frequency-based volatility (HEAVY) models', *Journal of Applied Econometrics* **25**(2), 197–231.

Takahashi, M., Omori, Y. and Watanabe, T. (2009), 'Estimating stochastic volatility models using daily returns and realized volatility simultaneously', *Computational Statistics and Data Analysis* **53**(6), 2404–2426.

Taylor, S. J. (1986), *Modelling Financial Time Series*, John Wiley & Sons, New York.

Taylor, S. J. (1994), 'Modeling stochastic volatility: A review and comparative study', *Mathematical Finance* **4**(2), 183–204.

Yu, J. (2002), 'Forecasting volatility in the New Zealand stock market', *Applied Financial Economics* **12**(3), 193–202.