Do We Need Stochastic Volatility and Generalised Autoregressive Conditional Heteroscedasticity? Comparing Squared End-Of-Day Returns on FTSE

David E. Allen [1,2,3,*] **and Michael McAleer** [2,4,5,6,7,8]

[1] School of Mathematics and Statistics, University of Sydney, Sydney 2006, Australia
[2] Department of Finance, Asia University, Taichung 41354, Taiwan; michael.mcaleer@gmail.com
[3] School of Business and Law, Edith Cowan University, Joondalup 6027, Australia
[4] Discipline of Business Analytics, University of Sydney Business School, Sydney, NSW 2006, Australia
[5] Econometric Institute, Erasmus School of Economics, Erasmus University Rotterdam, 3062 Rotterdam, The Netherlands
[6] Department of Economic Analysis and ICAE, Complutense University of Madrid, 28040 Madrid, Spain
[7] Department of Mathematics and Statistics, University of Canterbury, Christchurch 8041, New Zealand
[8] Institute of Advanced Sciences, Yokohama National University, Kanagawa 240-8501, Japan
[*] Correspondence: profallen2007@gmail.com

Abstract: The paper examines the relative performance of Stochastic Volatility (SV) and Generalised Autoregressive Conditional Heteroscedasticity (GARCH) (1,1) models fitted to ten years of daily data for FTSE. As a benchmark, we used the realized volatility (RV) of FTSE sampled at 5 min intervals taken from the Oxford Man Realised Library. Both models demonstrated comparable performance and were correlated to a similar extent with RV estimates when measured by ordinary least squares (OLS). However, a crude variant of Corsi's (2009) Heterogeneous Autoregressive (HAR) model, applied to squared demeaned daily returns on FTSE, appeared to predict the daily RV of FTSE better than either of the two models. Quantile regressions suggest that all three methods capture tail behaviour similarly and adequately. This leads to the question of whether we need either of the two standard volatility models if the simple expedient of using lagged squared demeaned daily returns provides a better RV predictor, at least in the context of the sample.

Keywords: stochastic volatility; GARCH (1,1); FTSE; RV 5 min; HAR model; demeaned daily squared returns.

JEL Classification: C22; G12.

1. Introduction

The paper explores the performance of Stochastic Volatility (SV) and Generalised Autoregressive Conditional Heteroscedasticity (GARCH) (1,1) models as estimators of the volatility of the FTSE Index. The volatilities estimated by these models are compared with realised volatility estimates for FTSE, obtained from the Oxford Man Realised Library and sampled at 5 min intervals, as described in Heber et al. (2009). Their volatility forecasts were further compared with those derived from a simple historical volatility model.

We used the stochvol R package, which uses Markov chain Monte Carlo (MCMC) samplers, to conduct inference by obtaining draws from the posterior distribution of parameters and latent variables, which could then be used for predicting future volatilities. This is done within the context of a fully Bayesian implementation of heteroscedasticity modelling within the framework of stochastic

volatility. For more information, see the discussion of the method by Kastner and Frühwirth-Schnatter (2014), and of the stochvol package by Kastner (2016).

Taylor (1982) suggested probabilistically modelling volatility, in effect through a state-space model where the logarithm of the squared volatilities, the latent states, follows an autoregressive process of order one. Over time, this specification became known as the SV model. A series of early papers by Jacquier et al. (1994), Ghysels et al. (1996), and Kim et al. (1998) provided evidence in support of the application of stochastic-volatility models, but their practical use has been infrequent. The reasons for this were considered by Bos (2012), who suggested that the empirical application of SV models was limited by two major factors: the variety (and potential incompatibility) of estimation methods for SV models and the lack of standard software packages for implementing these methods. The situation for multivariate SV is even more problematic.

Taylor (1994) provides a review of both stochastic volatility and the ARCH/GARCH literature. More recent reviews of the use of inference in the context of this literature were provided by McAleer (2005), and a review of multivariate stochastic-volatility models was provided by Asai et al. (2006). Granger and Poon (2003, p. 485), in their review of volatility-forecasting methods, noted the difficulties in the application of the SV model: "the SV model has no closed form, and hence cannot be estimated directly by maximum likelihood". They also noted that the quasimaximum-likelihood method may also be inefficient if the volatility proxies are non-Gaussian. The advantage of the stochvol R package is that it incorporates an efficient MCMC estimation scheme for SV models, as discussed by Kastner and Frühwirth-Schnatter (2014). This facilitated analysis in this paper, which features a direct comparison of the volatility predictions of a SV model, a GARCH (1,1) model, and a simple application of a historical-volatility-based estimation method, as applied to the FTSE index.

The paper is divided into four sections: Section 2 reviews the literature and employed econometric method. Section 3 presents the results, and Section 4 concludes the paper.

2. Previous Work and Econometric Models

2.1. Stochastic Volatility

There have been numerous empirical studies of changes in volatility in various stock, currency, and commodity markets. Findings in volatility research have implications for option pricing, volatility estimation, and the degree to which volatility shocks persist. These research questions have been approached by means of different models and methodologies.

Taylor (1982) suggested a novel SV approach, and Taylor (1994) implemented the SV model as follows: if P_t denotes the prices of an asset at time t,, and assuming no dividend payments, returns on the asset can be defined, in the context of discrete time periods, as:

$$X_t = ln[P_t/P_{t-1}]. \tag{1}$$

Volatility is customarily indicated by σ, and prices are described by a stochastic differential equation:

$$d(lnP) = \mu dt + \sigma dW, \tag{2}$$

with W being a standard Weiner process. If μ and σ are constants, X_t has normal distribution, and

$$X_t = \mu + \sigma U_t, \tag{3}$$

with U_t being independent and identically distributed ($i.i.d$).

Equation (3) can be generalised by replacing σ with positive random variable σ_t to give:

$$X_t = \mu + \sigma_t U_t, \tag{4}$$

where $U_t \sim N(0,1)$.

In circumstances where returns process $\{X_t\}$ can be presented by Equation (4), Taylor (1994) called σ_t the stochastic volatility for period t. His definition assumed that $(X_t - \mu)/\sigma_t$ follows normal distribution.

Stochastic process $\{\sigma_t\}$ generates realised volatilities $\{\sigma_t^*\}$ that, in general, are not observable. For any realisation, σ_t^*

$$X_t \mid \sigma_t = \sigma_t^* \sim N(\mu, \sigma_t^{*2}). \tag{5}$$

The mixture of these conditional normal distributions defines the unconditional distribution of X_t, which has excess kurtosis whenever σ_t has positive variance and is independent of U_t.

In the empirical section that follows, we used the RV of the FTSE sampled at 5 min intervals, provided by Oxford Man, as a proxy for true realised volatility. We then compared the estimates of volatility obtained from the SV and GARCH (1,1) models by using the RV estimates as a benchmark.

Taylor (1994, p. 3) suggested using "capital letters to represent random variables and lowercase letters to represent outcomes", and we follow that convention. Given the observed returns of $I_{t-1} = \{x_1, x_2, \ldots x_{t-1}\}$, the conditional variance for period t is:

$$h_t = var(X_t \mid I_{t-1}). \tag{6}$$

Taylor (1994) notes that, in general, random variable H_t, which generates observed conditional variance h_t, is not, in general, equal to σ_t^2. A convenient way to use economic theory to motivate changes in volatility is to assume that returns are generated by a number of intraperiod price revisions, as in the manner of Clark (1973) and Tauchen and Pitts (1983), to mention two of many studies.

It is assumed that there are N_t price revisions during trading day t, each caused by unpredictable information. Let event i on day t change the logarithmic price by ω_{it}, with

$$X_t = \mu + \sum_{i=1}^{N_t} \omega_{it}. \tag{7}$$

If we assume that $\omega_{it} \sim i.i.d.$ and is independent of random variable N_t, with $\omega_{it} \sim N(0, \sigma_\omega^2)$, then

$$\sigma_t^2 = \sigma_\omega^2 N_t \ \ and \ \ H_t = \sigma_\omega^2 E[N_t \mid I_{t-1}]. \tag{8}$$

The above model suggests that squared volatility is proportional to the amount of price information.

The lack of standard software to estimate such a model is addressed by Kastner and Frühwirth-Schnatter (2014), who proposed an efficient MCMC estimation scheme that is implemented in the R stochvol package (Kastner (2016)). Kastner (2016, p. 2) proceeded by "letting $y = (y_1, y_2, \ldots y_n)^|$ be a vector of returns, with mean zero. An intrinsic feature of the SV model is that each observation y_t is assumed to have its 'own' contemporaneous variance, e^{h_t}, which relaxes the usual assumption of homoscedasticity". It was assumed that the logarithm of this variance follows an autoregressive process of order one. This assumption is fundamentally different to GARCH models, where time-varying conditional volatility is assumed to follow deterministic instead of stochastic evolution.

The centered parameterization of the SV model can be given as:

$$y_t \mid h_t \sim N(0, exp^{h_t}), \tag{9}$$

$$h_t \mid h_{t-1}, \mu, \phi, \sigma_n \sim N(\mu, +\phi(h_{t-1} - \mu), \sigma_n^2), \tag{10}$$

$$h_0 \mid \mu, \phi, \sigma_n \sim N(\mu, \sigma_n^2/(1 - \phi^2)), \tag{11}$$

where $N(\mu, \sigma_n^2)$, denotes normal distribution with mean μ and variance σ_n^2. $\theta = (\mu, \phi, \sigma_n)^|$ is the vector of parameters that consists of the level of log variance μ, the persistence of log variance, ϕ, and volatility of log variance σ_n. Process $h = (h_0, h_1, \ldots, h_n)^|$, which features in Equations (10) and (11), is the unobserved or latent time-varying volatility process.

Kastner (2016, p. 4) remarked that: "A novel and crucial feature of the algorithm implemented in stochvol is the usage of a variant of the "ancillarity–sufficiency interweaving strategy" (ASIS) which has been brought forward in the general context of state-space models by Yu and Meng (2011). ASIS exploits the fact that, for certain parameter constellations, sampling efficiency improves substantially when considering a non-centered version of a state-space model".

Another key feature of the algorithm used in stochvol is the joint sampling of all instantaneous volatilities "all without a loop" (AWOL), a technique with links to Rue (2001) that is discussed in McCausland et al. (2011). The combination of these features enables the stochvol R package to efficiently estimate SV models even when large datasets are involved.

2.2. ARCH and GARCH

Engle (1982) developed the Autoregressive Conditional Heteroscedasticity (ARCH) model that incorporates all past error terms. It was generalised to GARCH by Bollerslev (1986) to include lagged term conditional volatility. In other words, GARCH predicts that the best indicator of future variance is a weighted average of long-run variance, the predicted variance for the current period, and any new information in this period, as captured by the squared residuals.

The framework was developed as follows: consider a time series $y_t = E_{t-1}(y_t) + \varepsilon_t$, where $E_{t-1}(y_t)$ is the conditional expectation of y_t at time $t-1$, and ε_t is the error term. The basic GARCH model has the following specification:

$$\varepsilon_t = \sqrt{h_t \eta_t}, \quad \eta_t \sim N(0,1) \tag{12}$$

$$h_t = \omega + \sum_{j=1}^{p} \alpha_j \varepsilon_{t-j}^2 + \sum_{j=1}^{q} \beta_j h_{t-j} \tag{13}$$

in which $\omega > 0$, $\alpha_j \geq 0$ and $\beta_j \geq 0$ (usually a positive fraction), to ensure a positive conditional variance, $h_t \geq 0$ (see Tsay (1987)). The ARCH effect is captured by parameter α_j, which represents the short-run persistence of shocks to returns, β_j captures the GARCH effect that contributes to long-run persistence, and $\alpha_j + \beta_j$ measures the persistence of the impact of shocks to returns to long-run persistence. A GARCH (1,1) process is weakly stationary if $\alpha_j + \beta_j \leq 1$ (see the discussion in **?**).

We contrasted the estimates of volatility from the SV model with those from a GARCH (1,1) model, and assessed which better explained the behaviour of the RV of FTSE sampled at 5 min intervals.

2.3. Realised Volatility

We used the 5 min RV estimates from Oxford Man for FTSE as the RV benchmark (see: https://realized.oxford-man.ox.ac.uk/data). Their database contains "daily (close-to-close) financial returns, and a corresponding sequence of daily realised measures $rm_1, rm_2, \ldots, rm_T$. Realised measures are theoretically sound, high-frequency, nonparametric-based estimators of the variation of the price path of an asset during the times in which the asset trades frequently on an exchange. Realised measures ignore the overnight variation of prices, and sometimes the variation in the first few minutes of the trading day, when recorded prices may contain large errors". The metrics were developed by Andersen et al. (2001), Andersen et al. (2003), and Barndorff-Nielsen and Shephard (2002). Shephard and Sheppard (2010) provide an account of RV measures used in the Oxford Man Realised Library.

The simplest realised metric is realised variance (RV):

$$RV_t = \sum x_{j,t}^2, \tag{14}$$

where $x_{j,t} = X_{t_{j,t}} - X_{t_{j-1,t}}$. The $t_{j,t}$ are the times of trades or quotes on the t-th day. The theoretical justification of this measure is that, if prices are observed without noise, then, as $\min j \mid t_{j,t} - t_{j-1,t} \mid \downarrow 0$, it consistently estimates the quadratic variation of the price process on the t-th day. If sampling is

reduced to very small time intervals, market-microstructure noise may become a contaminant. In order to avoid this issue, we used RV estimates from Oxford Man, sampled at 5 min intervals, hereafter $RV5$.

2.4. Historical-Volatility Model

Poon and Granger (2005) discussed various practical issues encountered in attempts to forecast volatility. They suggested that the HISVOL model has the following form:

$$\hat{\sigma}_t = \phi_1 \sigma_{t-1} + \phi_2 \sigma_{t-2} + .. + \phi_\tau \sigma_{t-\tau}, \tag{15}$$

where $\hat{\sigma}_t$ is the expected standard deviation at time t, ϕ is the weight parameter, and σ is the historical standard deviation for periods indicated by the subscripts. Poon and Granger (2005) suggested that this group of models include random walk, historical averages, autoregressive (fractionally integrated) moving average, and various forms of exponential smoothing that depend on weight parameter ϕ.

We used a simple form of this model in which the estimate of σ is the previous day's demeaned squared return. Poon and Granger (2005) noted that, in a review of 66 previous studies, implied standard deviations appeared to perform best, followed by historical volatility and GARCH, which have roughly equal performance. They also noted that, at the time of writing, there were insufficient studies of SV models to come to any conclusions about this class of models. This is a motivation for the current study that assesses the performance of all three classes of models.

Barndorff-Nielsen and Shephard (2003) pointed out that taking the sums of squares of increments of log prices has a long tradition in the financial economics literature. Early examples are Poterba and Summers (1986), Schwert (1989), Taylor and Xu (1997), Christensen and Prabhala (1998), Dacorogna et al. (1998), and Andersen et al. (2001). Shephard and Sheppard (2010, p. 200, footnote 4) noted that: "Of course, the most basic realised measure is the squared daily return". We utilised this approach as the basis of our historical-volatility model.

2.5. Heterogeneous Autoregressive Model (HAR)

Corsi (2009, p. 174) suggested "an additive cascade model of volatility components defined over different time periods. The volatility cascade leads to a simple AR-type model in the realized volatility with the feature of considering different volatility components realized over different time horizons and which he termed as a Heterogeneous Autoregressive model of Realized Volatility." Corsi (2009) suggested that the model successfully achieves the purpose of reproducing the main empirical features of financial returns (long memory, fat tails, and self-similarity) in a parsimonious way. He wrote his model as:

$$\sigma_{t+1d}^{(d)} = c + \beta^{(d)} RV_t^{(d)} + \beta^{(w)} RV_t^{(w)} + \beta^{(m)} RV_t^{(m)} + \tilde{\omega}_{t+1d}^{(d)}, \tag{16}$$

where $\sigma^{(d)}$ is daily integrated volatility, and $RV_t^{(d)}$, $RV_t^{(w)}$ and $RV_t^{(m)}$ are the daily, weekly, and monthly (ex post) observed realized volatilities, respectively.

We appealed to Corsi (2009) to inspire our HISVOL model, which uses lags of historical RV estimates, but we used lags of squared demeaned daily returns.

2.6. Quantile Regression

We used the RV5 estimates as a benchmark to assess how the SV and GARCH(1,1) models performed. We were interested in behaviour in the extreme tails as well as in the centre of distribution. We therefore used quantile regression to assess behaviour in the tails.

Koenker and Hallock (2001, p. 145) provided an introduction to quantile regression, and noted that "quantiles seem inseparably linked to the operations of ordering and sorting the sample observations that are usually used to define them". They added: "The symmetry of the piecewise linear absolute value function implies that the minimization of the sum of absolute residuals must equate the number

of positive and negative residuals, thus assuring that there are the same number of observations above and below the median".

They asked: What about the other quantiles? As the symmetry of the absolute value yields the median, it follows that minimising the sum of asymmetrically weighted absolute residuals by simply giving differing weights to positive and negative residuals provides the other quantiles. The solution to

$$\underset{\xi \in \mathbb{R}}{Min} \sum \rho_\tau (y_i - \xi), \tag{17}$$

where function $\rho_\tau(\cdot)$ is the titled absolute value function, shown in Figure 1, gives the τth sample quantile function.
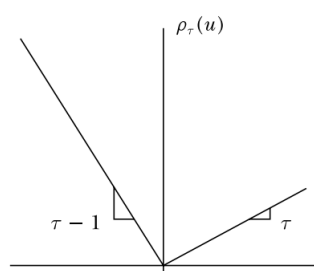

Figure 1. Quantile regression ρ function.

An estimate of the conditional median function can be obtained by replacing scalar ξ in Equation (10) by parametric function $\xi(x_i, \beta)$, and setting τ to $1/2$. Estimates of the other conditional quantile functions can be obtained by replacing absolute values by $\rho_\tau(\cdot)$, and solving Expression (11) by linear programming:

$$\underset{\beta \in \mathbb{R}p}{min} \sum \rho_\tau (y_i - \xi(x_i, \beta)). \tag{18}$$

We applied quantile regression to investigate the relationship between RV5 estimates on the FTSE Index and risk values predicted by the SV and GARCH(1,1) models.

3. Analysis Results

3.1. Preliminary Analysis

The sample dataset consisted of approximately 10 years of daily data of adjusted continuously compounded returns for FTSE, from 24 April 2009 through to 16 April 2019, sourced from Yahoo Finance via the R quantmod library. However, we retained the last 20 observations from the series for forecast purposes, leaving a total of 2454 observations. A matching set of daily RV5 estimates for FTSE was obtained from the Oxford Man Realised Library. Summary statistics for the two series is provided in Table 1 and plots of the two series in Figure 2.

FTSE has a mean daily return of 0.02% and a standard deviation of 0.97%. It has positive excess kurtosis and does not conform to Gaussian distribution, as can be seen from the QQ plot in Figure 3. RV5 has a mean of 8.6261×10^{-5} and a standard deviation of 0.00016097. However, it is measured as a variance, and if we take the square root of its value and multiply it by 100, it is on a common scale with the FTSE returns. We undertook this transformation in some of the comparison plots in subsequent figures. It has very high skewness and kurtosis, which is also evident in the QQ plots in Figure 3.

3.2. SV and GARCH Estimates

We used R library stochvol to fit a stochastic-volatility model to the FTSE return series, and used Gaussian distributions to fit the SV model. Some of the initial parameters for SV model estimation are shown in Table 2. The SV model applied to FTSE produced the volatility estimates shown in

Figure 4, while Figure 5 displays a kernel-density estimate for parameters contained in ϑ, producing posterior-density estimates. Figure 4 plots the estimated volatilities produced by the SV model at the 5%, 50%, and 95% posterior quantiles.

We also estimated GARCH (1,1) to obtain conditional volatilities. Coefficients are shown in Figure 3. The diagnostic tests suggested that there were no ARCH effects in the residuals, but rejected the null hypothesis that residuals conform to normal distribution. Plots of the volatilities obtained from both the SV and GARCH (1,1) estimates are shown in Figure 6.

Table 1. Summary Statistics of FTSE Continuously Compounded Returns and RV5.

Summary statistics using 27 April 2009–16 April 2019 observations for FTSERET variable (2454 valid observations)			
Mean	Median	Minimum	Maximum
0.00022266	0.00050198	−0.047798	0.050323
Std. Dev.	C.V.	Skewness	Ex. kurtosis
0.0097220	43.663	−0.13603	2.1367
5%	95%	IQ Range	Missing obs.
−0.015687	0.016056	0.010418	0
Summary Statistics, using 27 April 2009–16 April 2019 observations for RV5 variable (2454 valid observations)			
Mean	Median	Minimum	Maximum
8.6261×10^{-5}	5.1600×10^{-5}	1.3300×10^{-6}	0.0057390
Std. Dev.	C.V.	Skewness	Ex. kurtosis
0.00016097	1.8661	19.805	633.01
5%	95%	IQ Range	Missing obs.
1.4275×10^{-5}	0.00025567	6.8100×10^{-5}	0

Table 2. Stochastic-volatility estimates.

Summary of 1000 Markov chain Monte Carlo (MCMC) draws after burn-in of 1000					
Prior distributions					
$mu \sim normal$	mean = 0	S.D. = 100			
$(phi + 1)/2 \sim beta$	$a_0 = 5$	$b_0 = 1.5$			
$sigma^2 \sim 1 * Chisq(df = 1)$					
Posterior draws thinning = 1					
	Mean	S.D.	5%	50%	95%
mu	−9.5938	0.11914	−9.7831	−9.5944	−9.3963
phi	0.9577	0.01051	0.9389	0.9586	0.9734
$sigma$	0.2342	0.02868	0.1915	0.2321	0.2851
$exp(mu/2)$	0.0083	0.00049	0.0075	0.0083	0.0091
$sigma^2$	0.0557	0.01378	0.0367	0.0539	0.0813

Figure 2. Series plots.

Q-Q plot for FTSERET
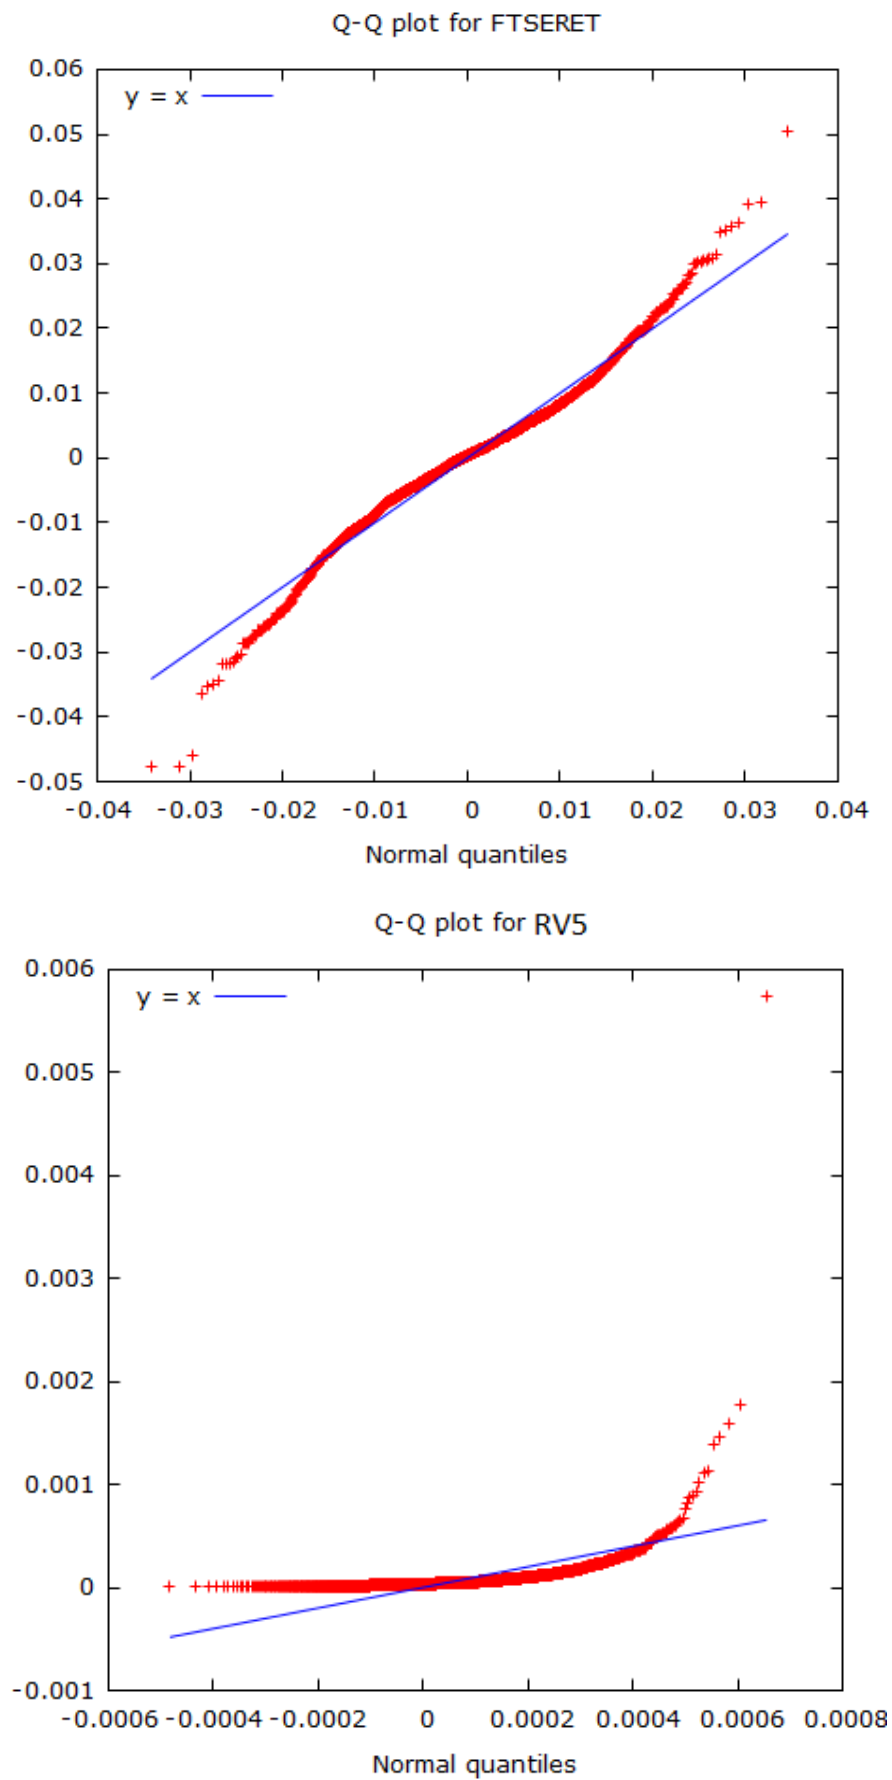
Q-Q plot for RV5
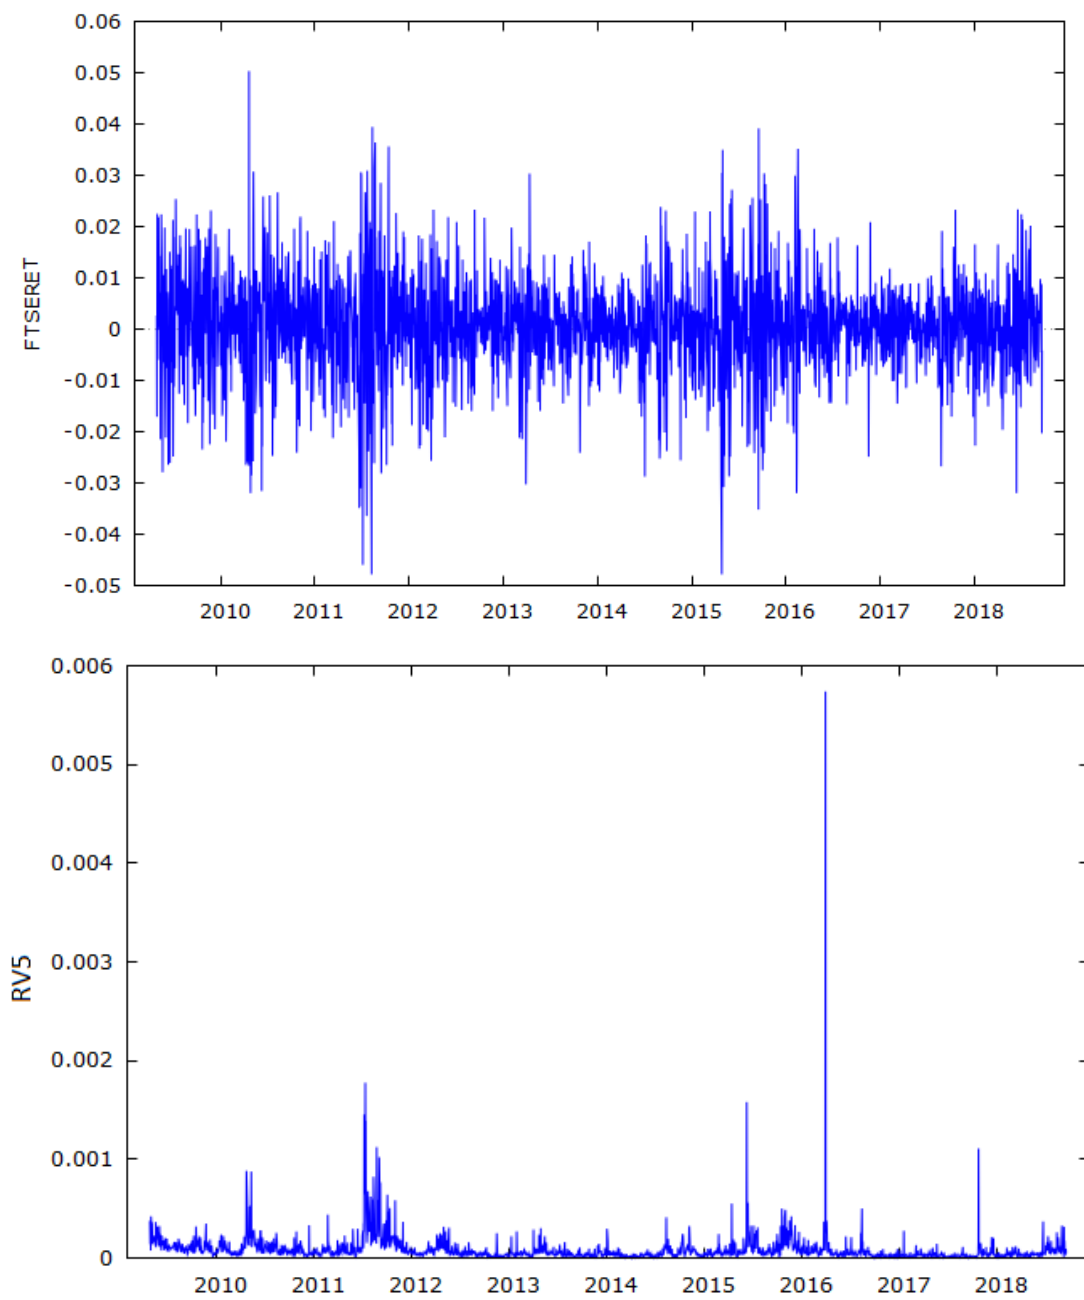
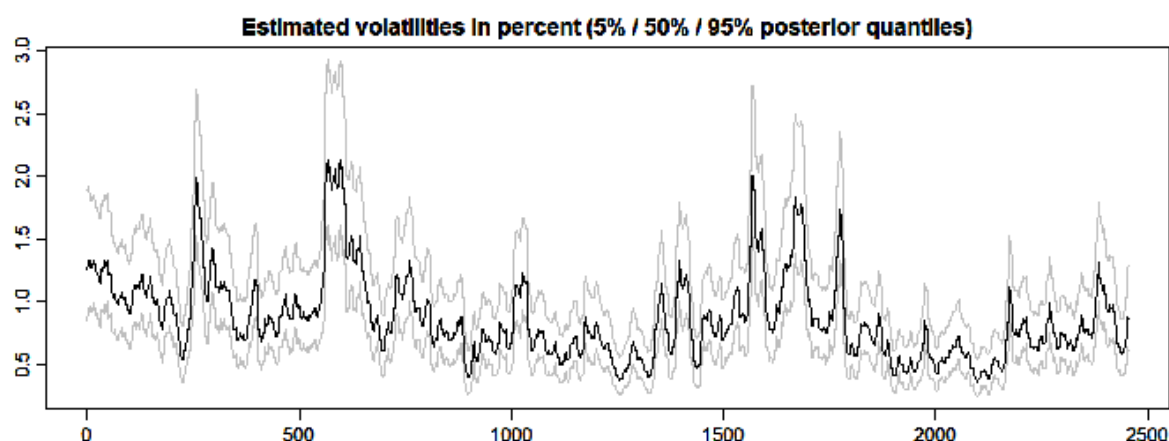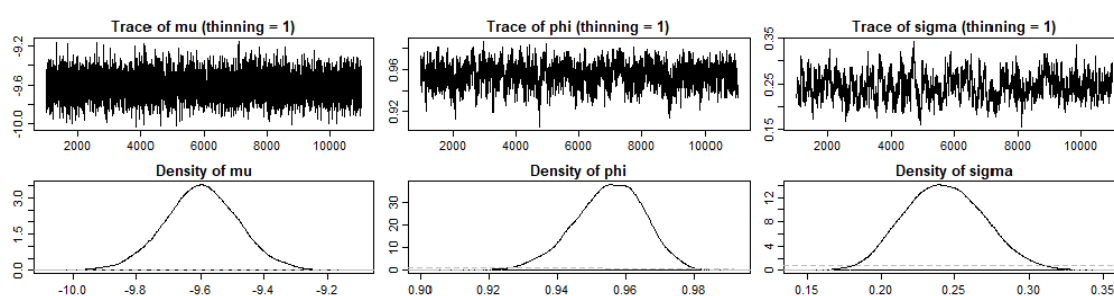**Figure 3.** QQ plots.

Figure 4. SV estimates.

Figure 5. Posterior-density plots of parameters in θ.

Table 3. Generalised Autoregressive Conditional Heteroscedasticity (GARCH) (1,1) fitted to FTSE returns.

	Coefficients	Standard Error	T Statistic
mu	1.360×10^{-4}	1.601×10^{-4}	0.850
omega	3.406×10^{-6}	8.087×10^{-7}	4.211 ***
alpha1	1.199×10^{-1}	1.728×10^{-2}	6.939 ***
beta1	8.443×10^{-1}	2.240×10^{-2}	37.688 ***

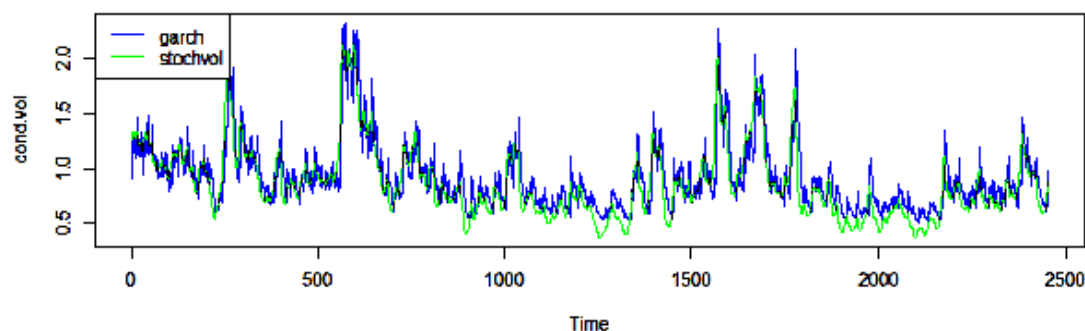Note: *** significance at 1% level.

Figure 6. Conditional volatilities: SV and GARCH (1,1).

Table 4 confirms that SV estimates were slightly lower than GARCH(1,1) estimates. Mean SV estimate was 0.87152, while mean GARCH(1,1) volatility estimate was 0.92405. This difference is significant, as revealed by a nonparametric sign test, the results of which are available from the authors on request. The interquartile range for the SV model is 0.38324, while that for GARCH(1,1) was slightly smaller with a value of 0.32905. This is reflected in the standard deviations of the estimates, with a

value of 0.32600 for the SV model, and 0.30490 for the GARCH(1,1). (Details of the GARCH model fitted are provided in Table 3).

Table 4. Summary statistics: conditional volatilities from SV and GARCH(1,1) models plus SQRV5L.

Summary statistics using 27 April 2009–16 April 2019 observations for SV variable (2454 valid observations)			
Mean	Median	Minimum	Maximum
0.87152	0.80112	0.37146	2.1388
Std. Dev.	C.V.	Skewness	Ex. kurtosis
0.32600	0.37405	1.2754	1.9519
5%	95%	IQ Range	Missing obs.
0.45723	1.5217	0.38324	0
Summary statistics using 27 April 2009–16 April 2019 observations for garchh variable_t (2454 valid observations)			
Mean	Median	Minimum	Maximum
0.92405	0.84738	0.51430	2.3305
Std. Dev.	C.V.	Skewness	Ex. kurtosis
0.30490	0.32996	1.6654	3.3669
5%	95%	IQ Range	Missing obs.
0.59248	1.5934	0.32905	0
Summary statistics using 27 April 2009–16 April 2019 observations for SQRV5L variable (2454 valid observations)			
Mean	Median	Minimum	Maximum
0.81989	0.71837	0.11543	7.5756
Std. Dev.	C.V.	Skewness	Ex. kurtosis
0.43644	0.53231	3.3736	29.966
5%	95%	IQ Range	Missing obs.
0.37737	1.5990	0.44676	0

The third panel in Table 4 provides summary statistics for SQRV5L which is the square root of realised volatility series RV5, scaled by 100, to transform it into the same dimension as the two conditional-volatility series. Mean value was markedly greater at 0.81989, standard deviation was also larger at 0.43644, while the interquartile range was 0.44676. Excess kurtosis was much greater, with a value of 29.966, given that the SV and GARCH(1,1) models had excess kurtosis of 1.9519 and 3.3669, respectively.

These statistics suggest that the modelfree estimate of volatility, provided by RV5, was more volatile and had larger spikes than the conditional volatilities produced by the two models.

We also estimated some regressions to explore the linear correlations between the estimates of the two volatility models and our base RV5 volatility estimates. In order to keep the values in the same dimension, we used the SQRV5L values. Results are shown in Table 5. If we used the RV5 values in the form of SQRV5L as the benchmark dependent variable, then conditional volatilities from the SV model had an adjusted R-square of 0.193842, and a highly significant slope coefficient with a value of 0.589930 and t-statistic of 24.31. The GARCH(1,1) model performs slightly better, with an adjusted R-squared of 0.202789 and a slope coefficient of 0.645116, with a t statistic of 25. The Durbin–Watson statistic for both models was around 1, suggesting that there iswasa borderline problem of serial correlation in the residuals.

As a further cross-check of the effectiveness of the two models, we used a further crude estimate of volatility with the demeaned squared daily returns on FTSE in the context of a HISVOL model, which was motivated by Corsi's (2009) Heterogeneous Autoregressive model of Realized Volatility (HAR-RV).

Table 5. Regression analysis of three volatility models: SV, GARCH(1,1), and lagged demeaned squared returns.

Ordinary least squares (OLS) using 27 April 2009–16 April 2019 observations (T = 2454). Dependent variable: SQRV5L.				
	Coefficient	Std. Error	t-ratio	p-value
const	0.305752	0.0225825	13.54	0.0000
SV	0.589930	0.0242700	24.31	0.0000
Mean dependent var		0.819888	S.D. dependent var	0.436435
Sum squared resid		376.5133	S.E. of regression	0.391859
R^2		0.194171	Adjusted R^2	0.193842
$F(1, 2452)$		590.8287	p-value(F)	4.1×10^{-117}
Log-likelihood		-1182.037	Akaike criterion	2368.075
Schwarz criterion		2379.686	Hannan–Quinn	2372.294
$\hat{\rho}$		0.503365	Durbin–Watson	0.991133

OLS using 27 April 2009–16 April 2019 observations (T = 2454). Dependent variable: SQRV5L.				
	Coefficient	Std. Error	t-ratio	p-value
const	0.223766	0.0251093	8.912	0.0000
garchh_t	0.645116	0.0258051	25.00	0.0000
Mean dependent var		0.819888	S.D. dependent var	0.436435
Sum squared resid		372.3346	S.E. of regression	0.389678
R^2		0.203114	Adjusted R^2	0.202789
$F(1, 2452)$		624.9784	p-value(F)	4.6×10^{-123}
Log-likelihood		-1168.344	Akaike criterion	2340.687
Schwarz criterion		2352.298	Hannan–Quinn	2344.906
$\hat{\rho}$		0.490865	Durbin–Watson	1.015052

OLS using 4 June 2009–16 April 2019 observations (T = 2426). Dependent variable: SQRV5L.				
	Coefficient	Std. Error	t-ratio	p-value
const	0.544629	0.0111890	48.68	0.0000
LSQDMFTSE_1	0.433934	0.414977	1.046	0.2958
LSQDMFTSE_2	0.515250	0.419102	1.229	0.2190
LSQDMFTSE_3	1.18029	0.420051	2.810	0.0050
LSQDMFTSE_4	0.839649	0.421807	1.991	0.0466
LSQDMFTSE_5	1.35951	0.423013	3.214	0.0013
LSQDMFTSE_6	0.102526	0.422970	0.2424	0.8085
LSQDMFTSE_7	0.663505	0.423177	1.568	0.1170
LSQDMFTSE_8	0.242477	0.421954	0.5747	0.5656
LSQDMFTSE_9	1.58351	0.423249	3.741	0.0002
LSQDMFTSE_10	0.715148	0.422994	1.691	0.0910
LSQDMFTSE_11	1.72472	0.421790	4.089	0.0000
LSQDMFTSE_12	2.43111	0.422496	5.754	0.0000
LSQDMFTSE_13	2.11699	0.422470	5.011	0.0000
LSQDMFTSE_14	0.667181	0.422379	1.580	0.1143
LSQDMFTSE_15	2.04957	0.422271	4.854	0.0000
LSQDMFTSE_16	0.438659	0.422071	1.039	0.2988
LSQDMFTSE_17	0.994273	0.421786	2.357	0.0185
LSQDMFTSE_18	0.430727	0.420974	1.023	0.3063
LSQDMFTSE_19	1.06033	0.420799	2.520	0.0118
LSQDMFTSE_20	0.0170270	0.421031	0.04044	0.9677
LSQDMFTSE_21	1.03750	0.419658	2.472	0.0135
LSQDMFTSE_22	0.303899	0.420813	0.7222	0.4703
LSQDMFTSE_23	0.842299	0.420028	2.005	0.0450
LSQDMFTSE_24	1.12694	0.419905	2.684	0.0073
LSQDMFTSE_25	0.834749	0.418637	1.994	0.0463
LSQDMFTSE_26	0.746605	0.416243	1.794	0.0730
LSQDMFTSE_27	1.92539	0.415125	4.638	0.0000
LSQDMFTSE_28	2.02831	0.411577	4.928	0.0000
Mean dependent var		0.812434	S.D. dependent var	0.432184
Sum squared resid		312.0871	S.E. of regression	0.360831
R^2		0.310989	Adjusted R^2	0.302940
$F(28, 2397)$		38.63918	p-value(F)	6.2×10^{-171}
Log-likelihood		-954.8254	Akaike criterion	1967.651
Schwarz criterion		2135.677	Hannan–Quinn	2028.745
$\hat{\rho}$		0.429918	Durbin–Watson	1.139917

In the third panel, we used the attribute that the squared demeaned returns on the FTSE are highly persistent, or have long memory, and regressed the square root of RV5 scaled by 100 on 28 lags of the squared demeaned FTSE return, also scaled by 100.

This crude model produced better results than those using SV and GARCH(1,1) conditional volatilities as explanatory variables. The adjusted R-squared value was 0.302940, and coefficients on lags of 3, 4, 5, 9, 11, 12, 13, 15, 17, 19, 21, 23, 24, 25, 27, and 28 of LSQDMFTSE were highly significant. The F statistic was highly significant, and the Durbin Watson statistic was slightly improved at 1.1399.

Plots of the volatility series are shown in Figure 7. In the upper panel of Figure 7, SQRV5 is plotted on a different scale to the two conditional volatilities so they were not superimposed on each other. In the lower panel, correspondence between RV5 and the crude measure of demeaned squared daily returns on the FTSE is evident.

Figure 8 provides a plot of actual versus fitted values when SQRV5L was regressed on 28 lags of LSQDMFTSE. Again, the effectiveness of the model is apparent. Figure 9 shows the predicted values of the conditional volatilities from the SV and GARCH, indexed on the L.H.S. of Figure 9, and the lagged values DMSQFTSERET plus the actual RV5 value for the day. The latter two series are indexed on the R.H.S. of Figure 9. If we rescaled the values of the latter two variables, it is apparent that there was closer proximity between them than from the predictions of SV and GARCH, even in the case of the first observation in the graph.
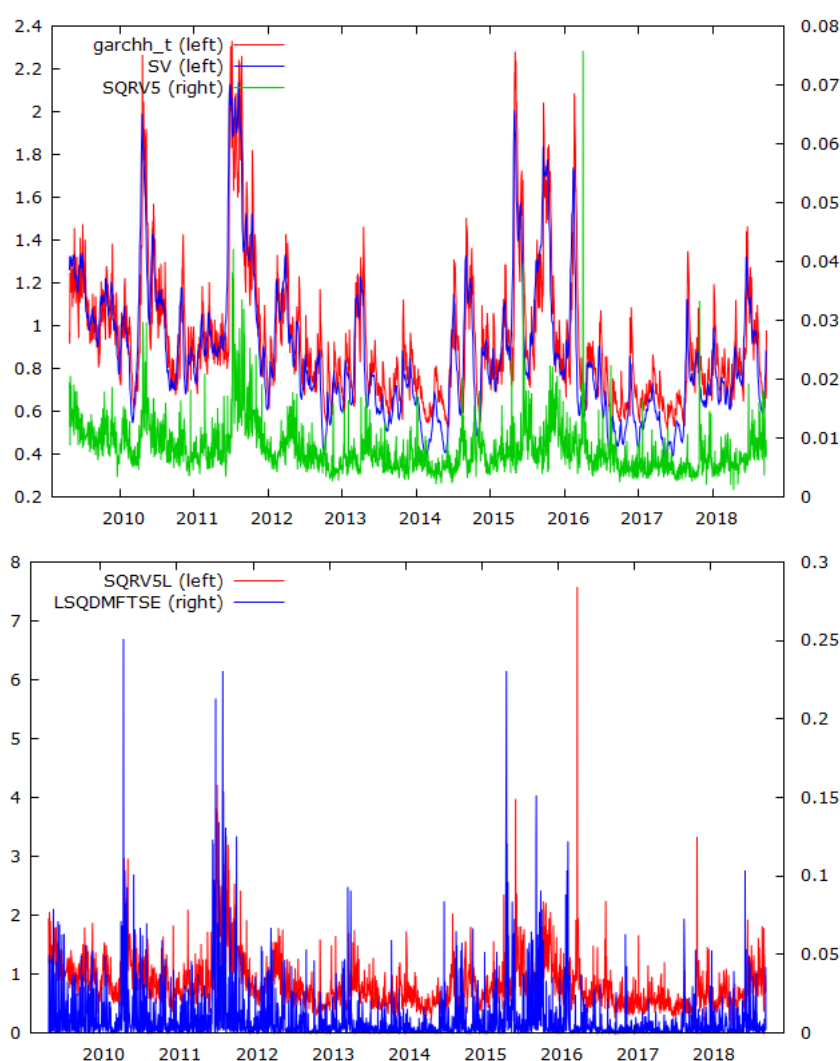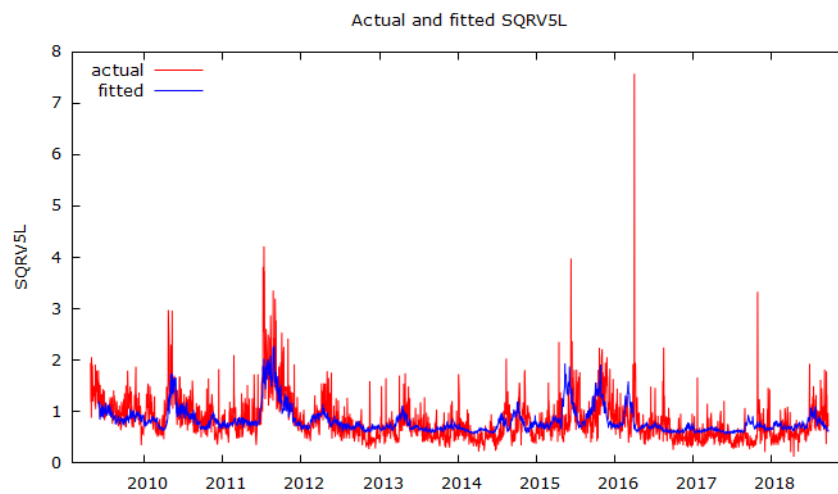

Figure 7. Plots of volatility estimates.

Figure 8. Plots of actual and fitted volatilities using lagged DMSQFTSERET as explanator.
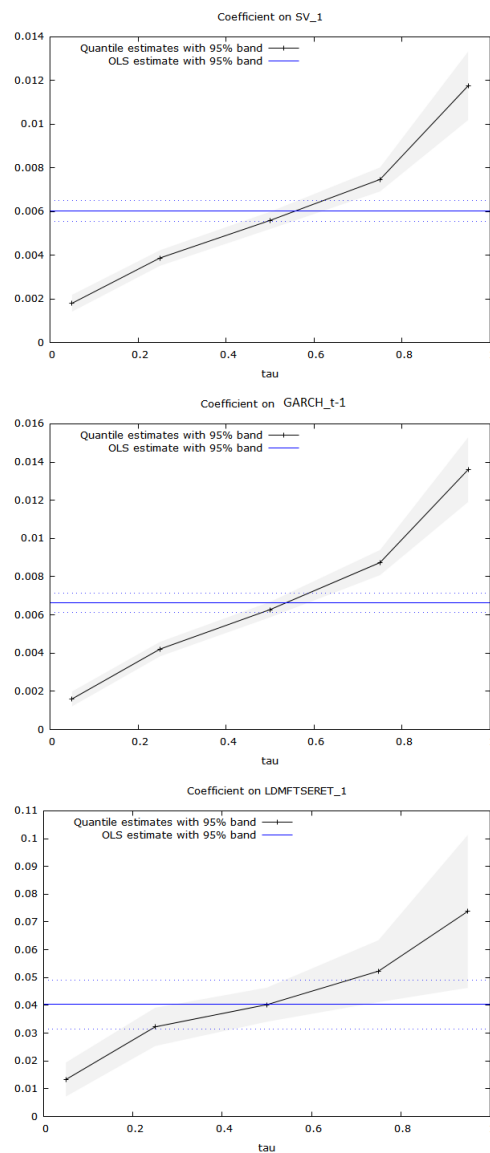

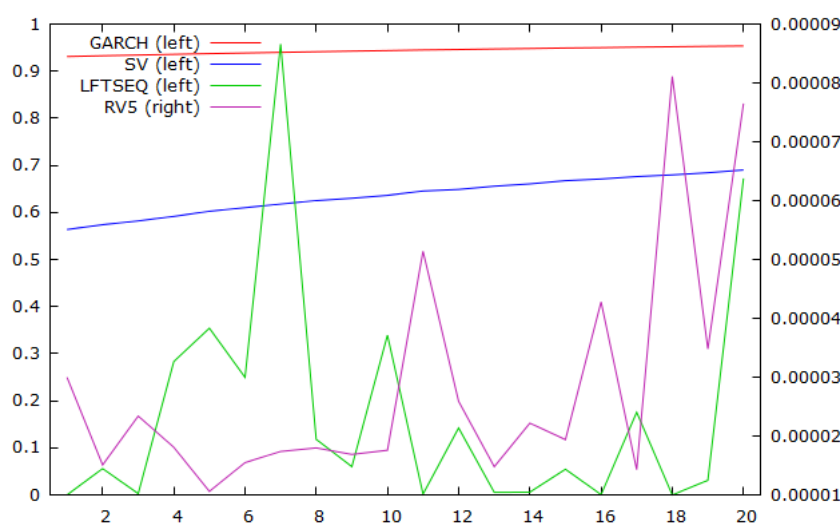**Figure 9.** Tau plots of quantile-regression coefficients.

Figure 10. Prediction plots for last 20 observations by SV and GARCH, and values of lagged DMSQFTSERET and RV5.

A potential issue was whether the relationship was significant in the quantiles, as this is an issue for risk measurement. In Table 6, we report the results of estimating quantile regressions for the same series.

Table 6. Quantile-regression results of SQRV5L regression on lagged SV, lagged GARCH and lagged DMSQFTSERET Quantile estimates, using 27 April 2009–16 April 2019 observations (T = 2454). Asymptotic standard errors assuming independent and identically distributed (IID) errors.

Variable	Tau	Coefficient	Std. Error	t-Ratio
SV(-1)	0.05	0.00179384	0.000194807	9.20827 ***
SV(-1)	0.25	0.00386922	0.000179883	21.5097 ***
SV(-1)	0.50	0.00558958	0.000199581	28.0066 ***
SV(-1)	0.75	0.00746445	0.000279415	26.7146 ***
SV(-1)	0.95	0.0117519	0.000798124	14.7245 ***
GARCHh_t(-1)	0.05	0.00159755	0.000195806	8.15884 ***
GARCHh_t(-1)	0.25	0.00420579	0.000191819	21.9259 ***
GARCHh_t(-1)	0.50	0.00627205	0.000204113	30.7283 ***
GARCHh_t(-1)	0.75	0.00873417	0.000334193	26.1351 ***
GARCHh_t(-1)	0.95	0.0135944	0.000860269	15.8025 ***
DMSQFTSERET(-1)	0.05	0.0132912	0.00310350	4.28266 ***
DMSQFTSERET(-1)	0.25	0.0322416	0.00350232	9.20578 ***
DMSQFTSERET(-1)	0.50	0.0402209	0.00311836	12.8981 ***
DMSQFTSERET(-1)	0.75	0.0522737	0.00566966	9.21991 ***
DMSQFTSERET(-1)	0.95	0.0737864	0.0140149	5.26486 ***

Note: *** significance at 1% level.

3.3. Quantile-Regression Results

The results of SQRV5L quantile regression on the lags of the three explanatory series are shown in Table 6. All three quantile regressions used SQRV5L as the dependent variable, adopting the modelfree RV5 estimation as the benchmark. We adopted tau values of 0.05, 0.25, 0.50, 0.75, and 0.95. Our concern was whether there was evidence of significant relationships across quantiles. All t statistics were significant at the 1% level across all quantiles. In the case of the third panel in Table 6, we only used one lag of the squared demeaned FTSE return, rather than the full model, to avoid excessive numbers of entries in the Tables. Even so, coefficients on all quantiles were significant at the 1% level.

Plots of these relationships are provided in Figure **??**. The horizontal blue lines in the panels in Figure **??** depict least-squares regression coefficients, and the dotted lines above and below them depict the error bands. The black line links the coefficients estimated by quantile regression, and the grey area around the line shows the error bands. An interesting feature of the quantile-regression coefficients was that all three black lines in Figure **??** have positive slopes. This is consistent with the existence of more elastic responses to estimates in higher quantiles of RV5 volatility, in all three cases of SV, GARCH(1,1), and squared demeaned returns on FTSE, as predictors of volatility. Results suggested that none of these models was fully successful in capturing volatility peaks in RV5 in the higher quantiles.

3.4. Rolling-Regression Analysis

As a further check, we ran some rolling regressions of our three daily volatility estimates from the three models of STOCHVOL, GARCH (1,1) and HISVOL as the explanatory variable for the realised RV5 estimates, using window sizes of 60 and 500 days.These window sizes were chosen to approximate three months and two years of daily returns. Graphs of the coefficient estimates from these three regressions, plus error bands of plus and minus two standard deviations, are shown in the appendix.

These plots suggest that, when we used the longer windows of 500 observations, the three models were significant around 50% of the time, in that the error bands did not span zero. None of the three models performed well in the 60 observation windows, and spanned zero for the majority of the considered period. The HISVOL model performed poorly in these short windows.

4. Conclusions

The paper featured an examination of the effectiveness of SV and GARCH(1,1) models as explainers of modelfree estimates of FTSE volatility using RV samples at 5 min intervals as a benchmark, as provided by Oxford Man Institute's Realised Library. In order to provide further contrast, we also used lags of squared demeaned daily returns on FTSE to provide a simple alternative estimate of daily volatility. The effectiveness of these three methods was explored via the application of ordinary-least-squares (OLS) regression and quantile-regression analysis. Poon and Granger (2005) provided motivation in their analysis of 66 studies of this topic in which they noted that, at that time, there were an insufficient number of SV studies to provide a comparison between the GARCH and HISVOL models. Our intention was to address this sparsity in the literature.

We used vanilla estimates of SV and GARCH(1,1) by adopting Gaussian distributions. If measured by adjusted R-squared values, GARCH(1,1) appeared to produce slightly better OLS predictions than the SV model. However, both performed relatively poorly as compared with the simple expedient of using squared demeaned daily returns on FTSE to predict RV5 volatility. Our results support Poon and Granger (2005), in that neither the GARCH nor the SV model outperformed a simple form of a HISVOL model on this FTSE sample when RV samples at 5 min intervals were used as a benchmark. However, all three models performed poorly within short 60 observation windows in the rolling regressions. This suggests that the HISVOL model performs better over a longer observation window and is likely to be of more practical applications over longer holding periods.

Appendix A
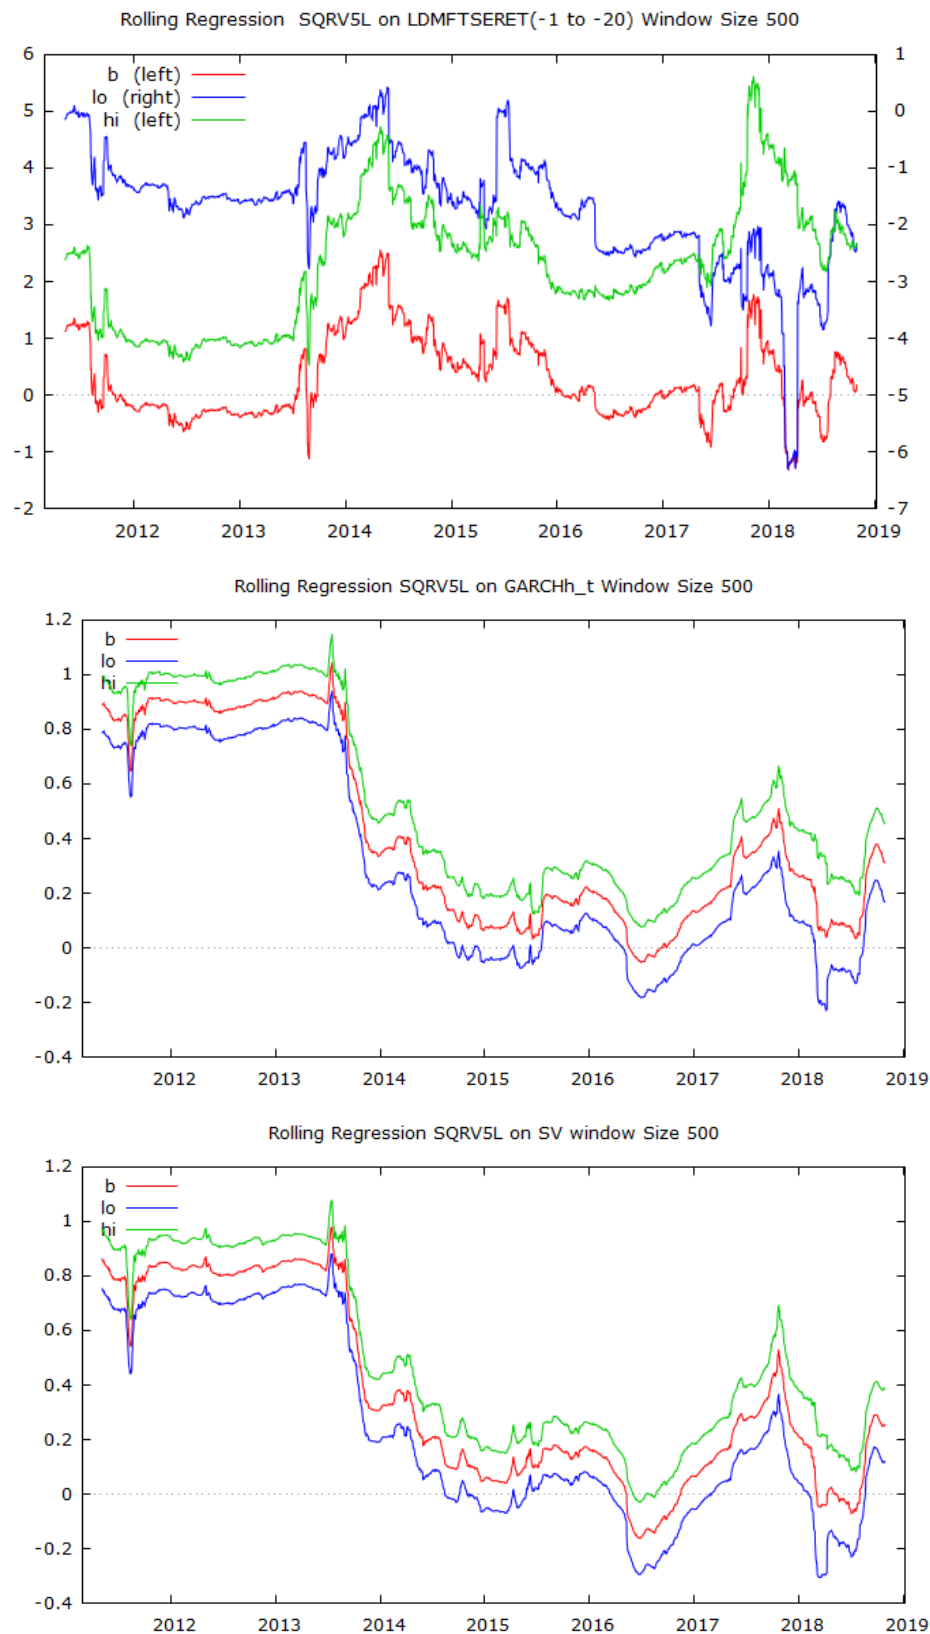
Note: Diagram middle line is regression-coefficient estimate; upper and lower lines are two standard error bands.

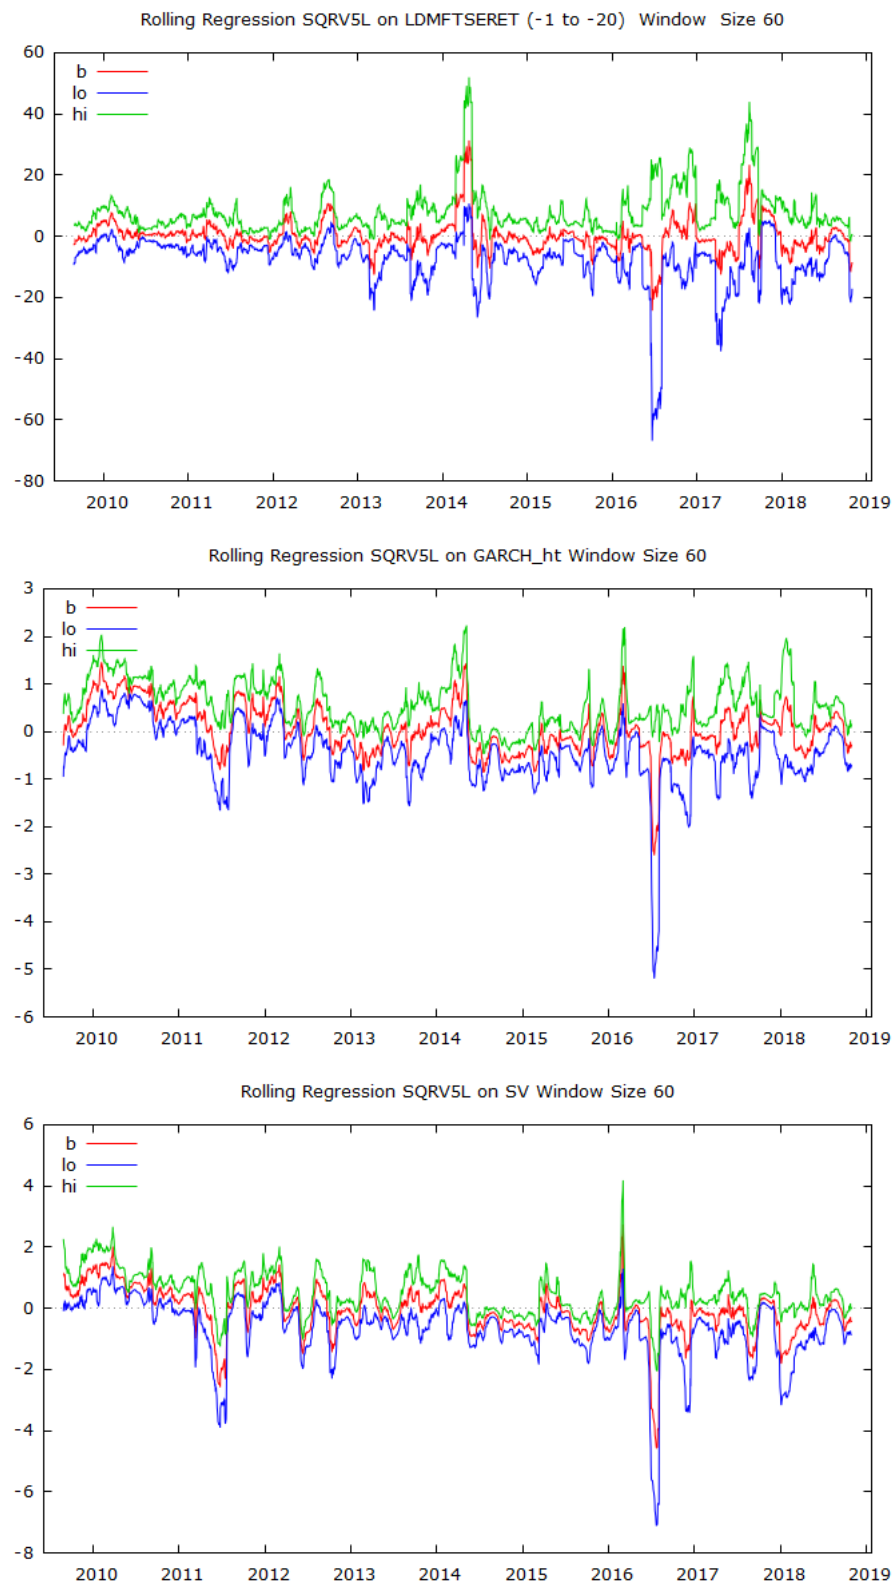**Figure A1.** Rolling-regression—window size 500.

Figure A2. Rolling-regression—window size 60.

References

Allen, David E., Ron Amram, and Michael McAleer. 2013. Volatility spillovers from the Chinese stock market to economic neighbours. *Mathematics and Computers in Simulation* 94: 238–57.

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Heiko Ebens. 2001. The distribution of realized stock return volatility. *Journal of Financial Economics* 61: 43–76.

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys. 2003. Modeling and forecasting realized volatility. *Econometrica* 71: 529–626.

Asai, Manabu, Michael McAleer, and Jun Yu. 2006. Multivariate stochastic volatility: A review. *Econometric Reviews* 25: 145–75.

Barndorff-Nielsen, Ole E., and Neil Shephard. 2002. Econometric analysis of realised volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society, Series B* 63: 253–80.

Barndorff-Nielsen, Ole E., and Neil Shephard. 2003. Realized power variation and stochastic volatility models. *Bernoulli* 9: 243–65.

Bollerslev, Tim. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31: 307–27.

Bos, Charles S. 2012. Relating stochastic volatility estimation methods. In *Handbook of Volatility Models and Their Applications*. Edited by L. Bauwens, C. Hafner, and S. Laurent. John Wiley & Sons, New York: pp. 147–74.

Christensen, Bent J., and Nagpurnanand R. Prabhala. 1998. The relation between implied and realized volatility. *Journal of Financial Economics* 37: 125–50

Clark, Peter K. 1973. A subordinated stochastic process model with finite variance for speculative prices. *Econometrica* 41: 135–55.

Corsi, Fulvio. 2009. A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics* 7: 174–96.

Dacorogna, Michel M., Ulrich A. Müller, Olivier V. Pictet, and Richard B. Olsen. 1998. Modelling short term volatility with GARCH and HARCH. In *Nonlinear Modelling of High Frequency Financial Time Series*. Edited by C. Dunis and B. Zhou. Wiley: Chichester, UK.

Engle, Robert F. 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* 50: 987–1007.

Ghysels, Eric, Andrew C. Harvey, and Eric Renault. 1996. Stochastic volatility. In *Statistical Methods in Finance, Volume 14, Handbook of Statistics*. Edited by G.S. Maddala and C.R. Rao. Elsevier: Amsterdam, The Netherlands, pp. 119–91.

Heber, Gerd, Asger Lunde, Neil Shephard, and Kevin Sheppard. 2009. *Oxford-Man Institute's Realized Library*. Oxford: Oxford-Man Institute, University of Oxford.

Jacquier, Eric, Nicholas G. Polson, and Peter E. Rossi. 1994. Bayesian analysis of stochastic volatility models. *Journal of Business & Economic Statistics* 12: 371–89.

Kastner, Gregor, and Sylvia Frühwirth-Schnatter. 2014. Ancillarity-sufficiency interweaving strategy (ASIS) for boosting MCMC estimation of stochastic volatility models. *Computational Statistics and Data Analysis* 76: 408–23.

Kastner, Gregor. 2016. Dealing with stochastic volatility in time series using the R package stochvol. *Journal of Statistical Software* 69: 1–30.

Kim, Sangjoon, Neil Shephard, and Siddhartha Chib. 1998. Stochastic volatility: Likelihood inference and comparison with ARCH models. *Review of Economic Studies* 65: 361–93.

Koenker, Roger, and Kevin F. Hallock. 2001. Quantile regression. *Journal of Economic Perspectives* 15: 143–56.

McAleer, Michael. 2005. Automated inference and learning in modeling financial volatility. *Econometric Theory* 21: 232–61.

McCausland, William J., Shirley Miller, and Denis Pelletier. 2011. Simulation smoothing for state-space models: A computational efficiency analysis. *Computational Statistics and Data Analysis* 55: 199–212.

Poon, Ser-Huang, and Clive W. J. Granger. 2003. Forecasting volatility in financial markets: A review. *Journal of Economic Literature* 41: 478–539.

Poon, Ser-Huang, and Clive Granger. 2005. Practical issues in forecasting volatility. *Financial Analysts Journal* 61: 45–56.

Poterba, James M., and Lawrence H. Summers. 1986. The persistence of volatility and stock market fluctuations. *American Economic Review* 76: 1124–41.

Rue, Håvard. 2001. Fast sampling of Gaussian Markov random fields. *Journal of the Royal Statistical Society Series B* 63: 325–38.

Schwert, G. William. 1989. Why does stock market volatility change over time? *Journal of Finance* 44: 1115–53.

Shephard, Neil, and Kevin Sheppard. 2010. Realising the future: Forecasting with high-frequency-based volatility (HEAVY) models. *Journal of Applied Econometrics* 25: 197–231.

Taylor, Stephen John. 1982. Financial returns modelled by the product of two stochastic processes: A study of daily sugar prices 1691–79. In *Time Series Analysis: Theory and Practice 1*. Edited by O. D. Anderson. Amsterdam, The Netherlands, pp. 203–26.

Taylor, Stephen J. 1994. Modeling stochastic volatility: A review and comparative study. *Mathematical Finance* 4: 183–204.

Taylor, Stephen J., and Xinzhong Xu. 1997. The incremental volatility information in one million foreign exchange quotations. *Journal of Empirical Finance* 4: 317–40.

Tauchen, George E., and Mark Pitts. 1983. The price variability-volume relationship on speculative markets. *Econometrica* 5: 485–505.

Tsay, Ruey S. 1987. Conditional heteroscedastic time series models. *Journal of the American Statistical Association* 82: 590–604.

Yu, Yaming, and Xiao-Li Meng. 2011. To center or not to center: That is not the question—An ancillarity-suffiency interweaving strategy (ASIS) for boosting MCMC efficiency. *Journal of Computational and Graphical Statistics* 20: 531–70.