J. R. Statist. Soc. B (2002)
64, *Part* 2, *pp.* 253–280

Econometric analysis of realized volatility and its use in estimating stochastic volatility models

Ole E. Barndorff-Nielsen

University of Aarhus, Denmark

and Neil Shephard

Nuffield College, Oxford, UK

[Received November 2000. Final revision November 2001]

Summary. The availability of intraday data on the prices of speculative assets means that we can use quadratic variation-like measures of activity in financial markets, called realized volatility, to study the stochastic properties of returns. Here, under the assumption of a rather general stochastic volatility model, we derive the moments and the asymptotic distribution of the realized volatility error—the difference between realized volatility and the discretized integrated volatility (which we call actual volatility). These properties can be used to allow us to estimate the parameters of stochastic volatility models without recourse to the use of simulation-intensive methods.

Keywords: Kalman filter; Leverage; Lévy process; Power variation; Quadratic variation; Quarticity; Realized volatility; Stochastic volatility; Subordination; Superposition

1. Introduction

1.1. Stochastic volatility

In the stochastic volatility (SV) model for log-prices of stocks and for log-exchange-rates a basic Brownian motion is generalized to allow the volatility term to vary over time. Then the log-price $y^*(t)$ follows the solution to the stochastic differential equation (SDE)

$$\mathrm{d}y^*(t) = \{\mu + \beta\,\sigma^2(t)\}\,\mathrm{d}t + \sigma(t)\,\mathrm{d}w(t), \tag{1}$$

where $\sigma^2(t)$, the *instantaneous* or *spot volatility*, will be assumed (almost surely) to have locally square integrable sample paths, while being stationary and stochastically independent of the standard Brownian motion $w(t)$. We shall label μ the drift and β the risk premium. Over an interval of time of length $\Delta > 0$ returns are defined as

$$y_n = y^*(\Delta n) - y^*\{(n-1)\Delta\}, \qquad n = 1, 2, \ldots, \tag{2}$$

which implies that, whatever the model for σ^2, it follows that

$$y_n|\sigma_n^2 \sim N(\mu\Delta + \beta\sigma_n^2, \sigma_n^2),$$

where

$$\sigma_n^2 = \sigma^{2*}(n\Delta) - \sigma^{2*}\{(n-1)\Delta\}$$

Address for correspondence: Neil Shephard, Nuffield College, Oxford, OX1 1NF, UK.
E-mail: neil.shephard@nuf.ox.ac.uk

 1369–7412/02/64253

and

$$\sigma^{2*}(t) = \int_0^t \sigma^2(u)\,\mathrm{d}u.$$

In econometrics $\sigma^{2*}(t)$ is called *integrated volatility*, whereas we call σ_n^2 *actual volatility*. Both definitions play a central role in the probabilistic analysis of SV models. Reviews of the literature on this topic are given in Taylor (1994), Shephard (1996) and Ghysels *et al.* (1996), whereas statistical and probabilistic aspects are studied in detail in Barndorff-Nielsen and Shephard (2001a). One of the key results in this literature (Barndorff-Nielsen and Shephard, 2001a) is that if we write (when they exist) ξ, ω^2 and r respectively as the mean, variance and the autocorrelation function of the process $\sigma^2(t)$ then

$$\left.\begin{aligned} E(\sigma_n^2) &= \xi\Delta, \\ \operatorname{var}(\sigma_n^2) &= 2\omega^2\, r^{**}(\Delta), \\ \operatorname{cov}(\sigma_n^2, \sigma_{n+s}^2) &= \omega^2 \diamond r^{**}(\Delta s), \end{aligned}\right\} \tag{3}$$

where

$$\diamond r^{**}(s) = r^{**}(s+\Delta) - 2\,r^{**}(s) + r^{**}(s-\Delta), \tag{4}$$

and

$$\begin{aligned} r^{*}(t) &= \int_0^t r(u)\,\mathrm{d}u, \\ r^{**}(t) &= \int_0^t r^{*}(u)\,\mathrm{d}u, \end{aligned} \tag{5}$$

i.e. the second-order properties of $\sigma^2(t)$ completely determine the second-order properties of actual volatility.

One of the most important aspects of SV models is that $\sigma^{2*}(t)$ can be exactly recovered using the entire path of $y^*(t)$. In particular, for the above SV model the *quadratic variation* is $\sigma^{2*}(t)$, i.e. we have

$$[y^*](t) = \operatorname*{plim}_{q\to\infty} \left[\sum\{y^*(t_{i+1}^q) - y^*(t_i^q)\}^2\right] = \sigma^{2*}(t) \tag{6}$$

for any sequence of partitions $t_0^q = 0 < t_1^q < \ldots < t_{m_r}^q = t$ with $\sup_i(t_{i+1}^q - t_i^q) \to 0$ for $q \to \infty$. Here plim denotes the probability limit of the sum. This is a powerful result for it does not depend on the model for instantaneous volatility nor the drift terms in the SDE for log-prices given in equation (1). The quadratic variation estimation of integrated volatility has recently been highlighted, following the initial draft of Barndorff-Nielsen and Shephard (2001a) and the concurrent independent work of Andersen and Bollerslev (1998a), by Andersen, Bollerslev, Diebold and Labys (2001) and Maheu and McCurdy (2001) in foreign exchange markets and Andersen, Bollerslev, Diebold and Ebens (2001) and Areal and Taylor (2002) in equity markets. See also the contribution of Comte and Renault (1998).

In practice, although we often have a continuous record of quotes or transaction prices, at a very fine level the SV model is a poor fit to the data. This is due to market microstructure effects (e.g. discreteness of prices, bid–ask bounce, irregular trading etc.; see Bai *et al.* (2000)). As a result we should regard the above quadratic variation result as indicating that we can estimate actual volatility, e.g. over a day, reasonably accurately by sums of squared returns, say, by using periods of 30 min but keeping in mind that taking returns over increasingly finer time periods will lead to the introduction of important biases. Hence the limit argument in quadratic variation is interesting but of limited direct practical use on its own. Suppose instead that we have fixed M intraday observations during each day; then the sum of squared intraday changes over a day is

$$\{y\}_n = \sum_{j=1}^{M} \left[y^*\left\{(n-1)\Delta + \frac{\Delta j}{M}\right\} - y^*\left\{(n-1)\Delta + \frac{\Delta(j-1)}{M}\right\}\right]^2, \tag{7}$$

which is an estimate of σ_n^2. It is a consistent estimate as $M \to \infty$, while it is unbiased when μ and β are 0. In econometrics $\{y\}_n$ has recently been labelled *realized volatility*, and we shall follow that convention here. Andersen, Bollerslev, Diebold and Labys (2001), Andersen, Bollerslev, Diebold and Ebens (2001) and Andersen *et al.* (2000) have empirically studied the properties of $\{y\}_n$ in foreign exchange and equity markets (earlier, less formal work on this topic includes Poterba and Summers (1986), Schwert (1989), Taylor and Xu (1997) and Christensen and Prabhala (1998)). In their econometric analysis they have regarded $\{y\}_n$ as a very accurate estimate of σ_n^2. Indeed they often regard the estimate as basically revealing the true value of actual volatility so that $y_n/\sqrt{\{y\}_n}$ is virtually Gaussian. So far no measure of error has been obtained which indicates the difference between $\{y\}_n$ and σ_n^2. We shall show that this difference is approximately mixed Gaussian, can be substantial and that more accurate estimates of σ_n^2 are readily available if we are willing to use a model for $\sigma^2(t)$. Andreou and Ghysels (2001) have independently approximated the properties of realized volatility using the methods of Foster and Nelson (1996) in their study of rolling estimators of the spot volatility $\sigma^2(t)$.

In this paper we shall discuss a simple way of formally bridging the gap between realized and actual volatility, providing a discussion of the properties of $\{y\}_n$ which has so far been lacking in the literature. Inevitably for finite M these properties will depend on the dynamics of the instantaneous volatility as well as the drift term in the SDE for log-prices. This must be the case, for the shorthand of ignoring the small sample effects of estimating σ_n^2 with the consistent $\{y\}_n$ is only valid for infeasibly large values of M.

In particular the contribution of our paper will be to allow us

(a) to derive the asymptotic distribution of $(\{y\}_n - \sigma_n^2)\sqrt{M}$ for large M, showing that this does not depend on μ and β, and is statistically sensible,
(b) to analyse the properties of realized volatility by assuming that $\mu = \beta = 0$ as the corresponding error has been shown to be small,
(c) to understand the exact second-order properties of $\{y\}_n$ when $\mu = \beta = 0$,
(d) to use the models for instantaneous volatility to provide *model-based* estimates of actual volatility (rather than model-free estimates which assume $M \to \infty$) using the series of realized volatility measurements when $\mu = \beta = 0$ (these model-based estimates can be based on past, current or historical sequences of realized volatilities) and
(e) to estimate the parameters of SV models by using simple and rather accurate statistical procedures when $\mu = \beta = 0$.

1.2. Empirical example

To illustrate some of the empirical features of realized volatility we have used the same return data as employed by Andersen, Bollerslev, Diebold and Labys (2001), although Appendix A will describe the slightly different adjustments that we have made to deal with some missing data. This US dollar–German Deutschmark series covers the 10-year period from December 1st, 1986, until November, 30th, 1996. Every 5 min it records the most recent quote to appear on the Reuters screen. It has been kindly supplied to us by Olsen and Associates in Zurich and preprocessed by Tim Bollerslev. It will be used throughout our paper to illustrate the results that we shall develop. In Fig. 1(a) we have drawn the correlogram of the squared 5-min returns over the 10 years' sample. It shows the well-known very strong diurnal effect (the x-axis is drawn in days). This will be discussed in detail in Section 6 but for now will be ignored entirely. Fig. 1(b) shows the correlogram of realized volatility, $\{y\}_n$, computed using $M = 288$ (i.e. based on 5-min data) and again using the whole 10 years of data. The graph starts out at around 0.6, decays very quickly for a number of days and then decays at a slower rate. Fig. 1(c) shows a cumulative version of the squared 5-min returns drawn on a small scale, while in Fig. 1(d) the same cumulative function is drawn over a larger timescale. It is the daily increments of this process which make up realized volatility.

1.3. Outline of the paper

The outline of the rest of the paper is as follows. In Section 2 we discuss the basic approach in the most straightforward set-up where μ and β are 0, providing the second-order properties of realized volatility. These can be used in estimating the value of actual volatility from a time series of realized volatilities. This is discussed in Section 3, which also contains a discussion of using the realized volatilities to provide estimates of continuous SV models. Section 4 gives an

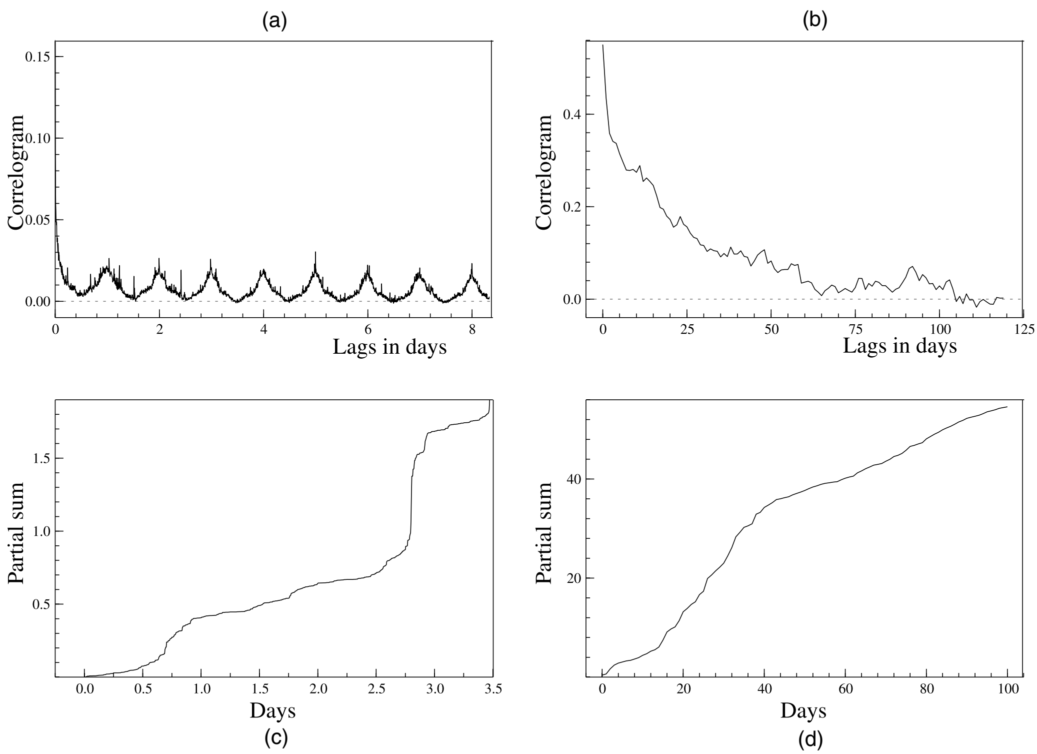

Fig. 1. Summary statistics for Olsen group's 5-min changes data: (a) autocorrelation function of 5-min returns; (b) autocorrelation function of realized volatility; (c) cumulative sum of squared 5-min changes over a short interval; (d) cumulative sum of squared 5-min changes over a long interval

empirical illustration of the methods developed in the previous two sections. Section 5 provides the asymptotic distribution of $(\{y\}_n - \sigma_n^2)\sqrt{M}$, covering the case where there is drift and a risk premium. This section shows that the effect on realized volatility of the drift and the risk premium is extremely small. Section 6 studies diurnal effects and leverage extensions. Section 7 concludes, whereas Appendix A contains a discussion of the data set that is used in this paper together with a proof of lemma 1 and theorem 1.

2. Relating actual to realized volatility

2.1. Generic results

Actual volatility σ_n^2 plays a crucial role in SV models. It can be estimated using realized volatility $\{y\}_n$, given in equation (7). Here we discuss this in the simplest context where $\mu = \beta = 0$, delaying our discussion of the effect on $\{y\}_n$ of the drift and risk premium until Section 5. In that section we shall show that the effect is minor and so the results that we develop here will still be important in that wider case.

In SV models we can always make the decomposition

$$\{y\}_n = \sigma_n^2 + u_n, \qquad \text{where } u_n = \{y\}_n - \sigma_n^2. \tag{8}$$

Here we call u_n the *realized volatility error*, which has the property that $E(u_n|\sigma_n^2) = 0$. Hence realized volatility is an unbiased estimator of actual volatility. We know that as $M \to \infty$ then $\{y\}_n \to \sigma_n^2$, almost surely, so it is also consistent. However, the purpose of this section is to discuss the properties of $\{y\}_n$ for finite M. We can see that

$$E(\{y\}_n) = \Delta\xi, \qquad \operatorname{var}(\{y\}_n) = \operatorname{var}(u_n) + \operatorname{var}(\sigma_n^2), \qquad \operatorname{cov}(\{y\}_n, \{y\}_{n+s}) = \operatorname{cov}(\sigma_n^2, \sigma_{n+s}^2).$$

Further, writing

$$\sigma_{j,n}^2 = \sigma^{2*}\left\{(n-1)\Delta + \frac{\Delta j}{M}\right\} - \sigma^{2*}\left\{(n-1)\Delta + \frac{\Delta(j-1)}{M}\right\}$$

we have that

$$u_n \overset{\mathcal{L}}{=} \sum_{j=1}^{M} \sigma_{j,n}^2(\varepsilon_{j,n}^2 - 1),$$

where $\varepsilon_{j,n} \overset{\text{IID}}{\sim} N(0,1)$ and independent of $\{\sigma_{j,n}^2\}$. It is clear that $\{u_n\}$ is a weak white noise sequence which is uncorrelated with the actual volatility series $\{\sigma_n^2\}$.

Now, unconditionally,

$$\begin{aligned} \operatorname{var}(u_n) &= 2M\, E\{(\sigma_{1,n}^2)^2\} \\ &= 2\, M\{\operatorname{var}(\sigma_{1,n}^2) + E(\sigma_{1,n}^2)^2\}, \end{aligned} \tag{9}$$

for $\sigma_{1,n}^2$ has the same marginal distribution as each element of $\{\sigma_{j,n}^2\}$. In general we have from equations (3) that

$$\begin{aligned} E(\sigma_{1,n}^2) &= \Delta M^{-1}\xi, \\ \operatorname{var}(\sigma_{1,n}^2) &= 2\omega^2\, r^{**}(\Delta M^{-1}). \end{aligned} \tag{10}$$

Hence we can compute $\text{var}(u_n)$ for all SV models when $\mu = \beta = 0$. In turn, having established the second-order properties of σ_n^2 and u_n, we can immediately use the results in Whittle (1983) to provide best linear prediction and smoothing results for the unobserved actual volatilities σ_n^2 from the time series of realized volatilities $\{y\}_n$. The only issues which remain on this front are computational. Otherwise this covers all covariance stationary models for $\sigma^2(t)$—including long memory processes.

One of the implications of the result given above is that

$$\begin{aligned}\text{corr}(\{y\}_n, \{y\}_{n+s}) &= \frac{\text{cov}(\sigma_n^2, \sigma_{n+s}^2)}{\text{var}(u_n) + \text{var}(\sigma_n^2)} \\ &= \frac{\omega^2 \diamond r^{**}(\Delta s)}{2M^{-1}\{2\omega^2 M^2 r^{**}(\Delta M^{-1}) + (\Delta\xi)^2\} + 2\omega^2 r^{**}(\Delta)}.\end{aligned}$$

Notice that

$$\text{corr}(y_n^2, y_{n+s}^2) = \frac{\omega^2 \diamond r^{**}(\Delta s)}{2\{2\omega^2 r^{**}(\Delta) + (\Delta\xi)^2\} + 2\omega^2 r^{**}(\Delta)}$$

can be derived from this result, for $\{y\}_n = y_n^2$ when $M = 1$. Hence the decay rates in the autocorrelation function of $\{y\}_n$, σ_n^2 and y_n^2 are the same in general but the correlation varies considerably, being the highest for σ_n^2, followed by $\{y\}_n$ and lowest for y_n^2.

In practice we tend to use realized volatility measures with M being moderately large. Hence it is of interest to think of a central limit approximation to the distribution of u_n. This will depend on the limit of $t^{-2} r^{**}(t)$ as $t \to 0$ from above. Now, by Taylor series expansion

$$r^{**}(t) = r^{**}(0+) + t\, r^{*}(0+) + \tfrac{1}{2}t^2\, r(0+) + o(t^2) = \tfrac{1}{2}t^2\, r(0+) + o(t^2).$$

This means that the limit of $t^{-2} r^{**}(t)$ is $\frac{1}{2}$. A consequence of this is that

$$\lim_{M\to\infty} \{M^2\, \text{var}(\sigma_{1,n}^2)\} = \Delta^2\omega^2, \tag{11}$$

implying that, as $M \to \infty$,

$$\text{var}(u_n\sqrt{M}) = \text{var}\{(\{y\}_n - \sigma_n^2)\sqrt{M}\} \to 2\Delta^2(\omega^2 + \xi^2).$$

This is an important result. We have moved away from the standard consistency result of $\{y\}_n \to \sigma_n^2$ in probability as $M \to \infty$, which follows from familiar quadratic variation results. Now we have the more refined measure of the uncertainty of this error term.

2.2. *Simple example*

Suppose that the volatility process has the autocorrelation function $r(t) = \exp(-\lambda\,|t|)$. Here we recall two classes of processes which have this property. The first is the constant elasticity of variance (CEV) process which is the solution to the SDE

$$\text{d}\sigma^2(t) = -\lambda\{\sigma^2(t) - \xi\}\,\text{d}t + \omega\,\sigma(t)^\eta\,\text{d}b(\lambda t), \qquad \eta \in [1, 2],$$

where $b(t)$ is standard Brownian motion uncorrelated with $w(t)$. Of course the special case of $\eta = 1$ delivers the square-root process, whereas when $\eta = 2$ we have Nelson's generalized

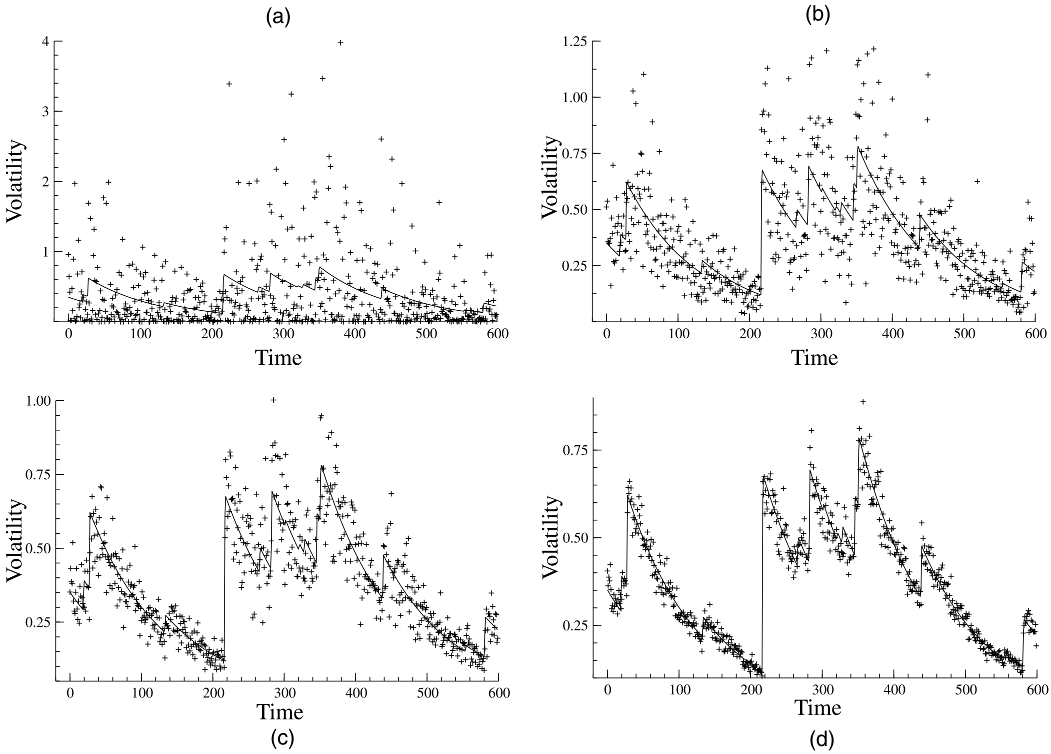

Fig. 2. Actual σ_n^2 (———) and realized $\{y\}_n$ (+) (with M varying) volatility based on a $\Gamma(4, 8)$ OU process with $\lambda = -\log(0.98)$ and $\Delta = 1$ (this implies $\xi = 0.5$ and $\xi\omega^{-2} = 8$): (a) $M = 1$; (b) $M = 12$; (c) $M = 48$; (d) $M = 288$

autoregressive conditional heteroscedastic diffusion. These models have been heavily favoured by Meddahi and Renault (2002) in this context. The second process is the non-Gaussian Ornstein–Uhlenbeck (OU) process which is the solution to the SDE

$$\mathrm{d}\sigma^2(t) = -\lambda\,\sigma^2(t)\,\mathrm{d}t + \mathrm{d}z(\lambda t), \tag{12}$$

where $z(t)$ is a Lévy process with non-negative increments. These models have been developed in this context by Barndorff-Nielsen and Shephard (2001a). In Fig. 2 we have drawn a curve to represent a simulated sample path of σ_n^2 from an OU process where $\sigma^2(t)$ has a $\Gamma(4, 8)$ stationary distribution, $\lambda = -\log(0.98)$ and $\Delta = 1$, along with the associated realized volatility (depicted by using crosses) computed using a variety of values of M. We see that as M increases the precision of realized volatility increases, whereas Fig. 2(d) shows that the variance of the realized volatility error increases with the volatility, a result which we shall come back to in Section 5 where the asymptotic distribution that we develop for $u_n\sqrt{M}$ will reflect this feature.

For both CEV and OU models

$$r^{**}(t) = \lambda^{-2}\{\exp(-\lambda\,|t|) - 1 + \lambda t\}$$

and

$$\diamond r^{**}(\Delta s) = \lambda^{-2}\{1 - \exp(-\lambda\Delta)\}^2 \exp\{-\lambda\Delta(s-1)\}, \qquad s > 0.$$

This implies that

$$E(\sigma_n^2) = \Delta\xi,$$

$$\mathrm{var}(\sigma_n^2) = \frac{2\,\omega^2}{\lambda^2}\{\exp(-\lambda\Delta) - 1 + \lambda\Delta\}$$

and

$$\mathrm{corr}(\sigma_n^2, \sigma_{n+s}^2) = d\exp t\{-\lambda\Delta(s-1)\}, \qquad s = 1, 2, \ldots, \tag{13}$$

where

$$d = \frac{\{1 - \exp(-\lambda\Delta)\}^2}{2\{\exp(-\lambda\Delta) - 1 + \lambda\Delta\}} \leqslant 1.$$

Finally

$$\begin{aligned}\mathrm{var}(u_n) &= 2M\{\mathrm{var}(\sigma_{1,n}^2) + E(\sigma_{1,n}^2)^2\} \\ &= 2M[2\omega^2\,\lambda^{-2}\{\exp(-\lambda\Delta/M) - 1 + \lambda\Delta M^{-1}\} + (\Delta M^{-1})^2\xi^2].\end{aligned} \tag{14}$$

Importantly this analysis implies that actual volatility has the autocorrelation function of an autoregressive moving average (ARMA) model of order (1, 1). Its autoregressive root is $\exp(-\lambda\Delta)$, which will be typically close to 1 unless Δ is very large, whereas the moving average root is also determined by $\exp(-\lambda\Delta)$ but must be found numerically. A graph of the moving average root against $\exp(-\lambda\Delta)$ is given in Fig. 3(a) and shows that for a wide range of the autoregressive root the moving average root is around 0.265. Likewise Fig. 3(b) shows a plot of d against $\exp(-\lambda\Delta)$ and indicates a rapid decline in this coefficient as the autoregressive root falls. In particular, in financial econometrics the literature suggests that volatility is quite persistent, which would imply that d should be close to 1. Thus, if t is recorded in days and Δ is set to 1 day, then empirically reasonable values of λ will imply that d should be close to 1.

In turn the autocorrelation function for σ_n^2 implies that the squares of returns have autocorrelations of the form

$$\mathrm{corr}(y_n^2, y_{n+s}^2) = c'\exp\{-\lambda\Delta(s-1)\}, \tag{15}$$

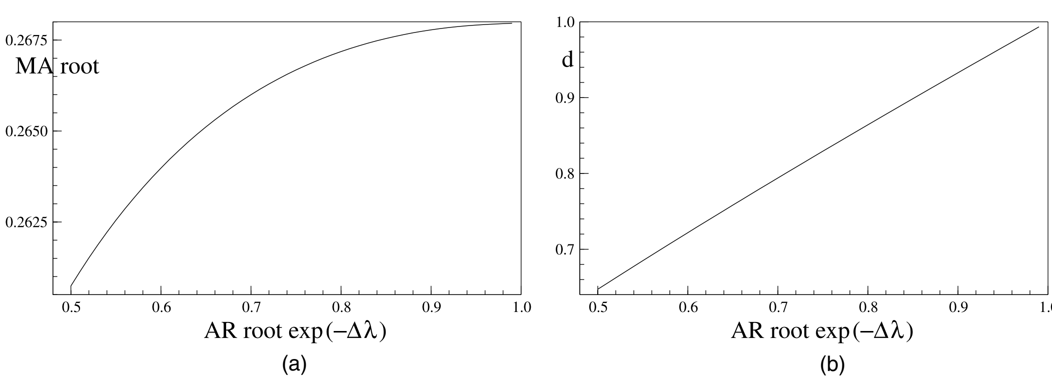

Fig. 3. (a) Moving average root plotted against autoregressive root $\exp(-\Delta\lambda)$ for the ARMA(1, 1) representation and (b) d in the expression for $\mathrm{corr}(\sigma_n^2, \sigma_{n+s}^2)$ against autoregressive root $\exp(-\Delta\lambda)$

where

$$\frac{1}{3} \geqslant \frac{1}{3}d \geqslant c' = \frac{\{1 - \exp(-\lambda\Delta)^2\}}{6\,\{\exp(-\lambda\Delta) - 1 + \lambda\Delta\} + 2(\lambda\Delta)^2\,(\xi/\omega)^2} \geqslant 0.$$

This means that y_n^2 also has a linear ARMA(1, 1) representation. Further, it has the same autocorrelation function as the familiar generalized autoregressive conditional heteroscedastic model that is used extensively in econometrics (see, for example, Bollerslev *et al.* (1994)). Finally, the autoregressive root of the ARMA representation is the same for y_n^2 as for σ_n^2; however, the moving average root of the square changes is much larger in absolute value. The implication is that the correlograms for y_n^2 will be less clear than if we had observed the correlograms of the σ_n^2. This can be most easily seen by noting that for small λ

$$c' \simeq \frac{1 - \lambda\Delta}{3 + 2\,(\xi/\omega)^2},$$

which is much smaller than d which is approximately $1 - \lambda\Delta$. For example, if $\xi = \omega$, then c' will be approximately 0.2 for daily data.

2.3. Extension of the example: superpositions

The OU or CEV volatility models are often too simple to fit accurately the types of dependence structures that we observe in financial economics. This can be seen Fig. 1(b) which displays the autocorrelation function of realized volatility for the Olsen group's 5-min data. This graph shows a relatively quick initial decline in the autocorrelation function, followed by a slower decay. This single observation is sufficient to dismiss the OU and CEV models.

One mathematically tractable way of improving the flexibility of the volatility model is to let the instantaneous volatility be the sum, or superposition, of independent OU or CEV processes. As the processes do not need to be identically distributed, this offers plenty of flexibility while still being mathematically tractable. Superpositions of such processes also have potential for modelling long-range dependence and self-similarity in volatility. This is discussed in the OU case in Barndorff-Nielsen and Shephard (2001a) and in more detail by Barndorff-Nielsen (2001) who formalizes the use of superpositions as a way of modelling long-range dependence. This follows earlier related work by Granger (1980), Barndorff-Nielsen *et al.* (1990), Cox (1991), Ding and Granger (1996), Engle and Lee (1999), Shephard (1996), pages 36–37, Andersen and Bollerslev (1997a), Barndorff-Nielsen (1998) and Comte and Renault (1998).

Consider volatility based on the sum of J independent OU or CEV processes,

$$\sigma^2(t) = \sum_{i=1}^{J} \tau^{(i)}(t),$$

where the $\tau^{(i)}(t)$ process has the memory parameter λ_i. We assume that

$$E\{\tau^{(i)}(t)\} = w_i\xi,$$
$$\operatorname{var}\{\tau^{(i)}(t)\} = w_i\omega^2,$$

where $\{w_i \geqslant 0\}$ and $\Sigma_{i=1}^{J} w_i = 1$, implying that $E\{\sigma^2(t)\} = \xi$ and $\operatorname{var}\{\sigma^2(t)\} = \omega^2$. The implication is that

$$\operatorname{cov}\{\sigma^2(t), \sigma^2(t+s)\} = \sum_{i=1}^{J} \operatorname{cov}\{\tau^{(i)}(t), \tau^{(i)}(t+s)\} = \omega^2 \sum_{i=1}^{J} w_i \exp(-\lambda_i\,|s|)\,.$$

Hence the autocorrelation function of instantaneous volatility can have components which are a mix of quickly and slowly decaying components. For fixed J the statistical identification of this model (imposing a constraint like $\lambda_1 < \ldots < \lambda_J$) is a consequence of the form of the autocorrelation function and the uniqueness of the Laplace transformation.

The linearity of the superposition of OU processes means that actual volatility has the form $\sigma_n^2 = \Sigma_{i=1}^J \tau_n^{(i)}$ where

$$\tau_n^{(i)} = \tau^{(i)*}(n\Delta) - \tau^{(i)*}\{(n-1)\Delta\}$$

and

$$\tau^{(i)*}(t) = \int_0^t \tau^{(i)}(u)\,\mathrm{d}u.$$

The key feature is that each $\tau_n^{(i)}$ has an ARMA(1, 1) representation of the type discussed earlier. As the autocovariance function of a sum of independent components is the sum of the autocovariances of the terms in the sum, we can compute the autocorrelation function of σ_n^2 without any new work. Computationally it is helpful to realize that the sum of uncorrelated ARMA(1, 1) processes can be fed into a linear state space representation when combined with decomposition (8). The only new issue is computing

$$\mathrm{var}(u_t) = 2M\{\mathrm{var}(\sigma_{1,n}^2) + E(\sigma_{1,n}^2)^2\}.$$

Clearly $E(\sigma_{1,n}^2) = \xi\Delta M^{-1}$ whereas

$$\begin{aligned}
\mathrm{var}(\sigma_{1,n}^2) &= \sum_{i=1}^J \mathrm{var}(\tau_{1,n}^{(i)}) = 2\omega^2 \sum_{i=1}^J w_i\, r_i^{**}(\Delta M^{-1}) \\
&= 2\omega^2 \sum_{i=1}^J \frac{w_i}{\lambda_i^2}\{\exp(-\lambda_i\Delta M^{-1}) - 1 + \lambda_i\Delta M^{-1}\} \\
&= 2\omega^2 \sum_{i=1}^J \frac{w_i}{2\lambda_i^2}(\lambda_i\Delta M^{-1})^2 + o(M^{-2}).
\end{aligned}$$

Importantly, for large M this expression simplifies and so we again obtain

$$\mathrm{var}(u_n\sqrt{M}) \to 2\Delta^2(\omega^2 + \xi^2), \qquad \text{as } M \to \infty.$$

3. Efficiency gains: model-based estimators of volatility

3.1. State space representation

If $\sigma^2(t)$ is an OU or CEV process then σ_n^2 has an ARMA(1, 1) representation and so it is computationally convenient to place decomposition (8) into a linear state space representation (see, for example, Harvey (1989), chapter 3, and Hamilton (1994), chapter 13). In particular we write $\alpha_{1n} = \sigma_n^2 - \Delta\xi$ and $u_n = \sigma_u v_{1n}$; then the state space representation is explicitly

$$\begin{aligned}
\{y\}_n &= \Delta\xi + (1 \quad 0)\,\alpha_n + \sigma_u v_{1n}, \\
\alpha_{n+1} &= \begin{pmatrix} \phi & 1 \\ 0 & 0 \end{pmatrix} \alpha_n + \begin{pmatrix} \sigma_\sigma \\ \sigma_\sigma\theta \end{pmatrix} v_{2n},
\end{aligned} \tag{16}$$

where v_n is a zero-mean, white noise sequence with an identity covariance matrix. The parameters ϕ, θ and σ_σ^2 represent the autoregressive root, the moving average root and the variance of the innovation to this process, whereas σ_u^2 is found from equations (9) and (10). Software for handling linear state space models is available in Koopman *et al.* (1999). Having constructed this representation we can use a Kalman filter to estimate unbiasedly and efficiently (in a linear sense) σ_n^2 by prediction (i.e. the estimate of σ_n^2, using $\{y\}_1, \ldots, \{y\}_{n-1}$) and smoothing (i.e. the estimate of σ_n^2, using $\{y\}_1, \ldots, \{y\}_T$ where T is the sample size). By-products of the Kalman filter are the mean-square errors of these *model-based* (i.e. they depend on the assumption that σ_n^2 has an ARMA(1, 1) representation) estimators.

Table 1 reports the mean-square error of the model-based predictor and smoother of actual volatility, as well as the corresponding result for the model-free raw realized volatility (the mean-square errors of the model-based estimators will be above the figures quoted towards the very start and end of the sample, for we have quoted steady state quantities). The results in the left-hand block of Table 1 correspond to the model which was simulated in Fig. 2, whereas the other blocks vary the ratio of ξ to ω^2. The exercise is repeated for two values of λ.

The main conclusion from the results in Table 1 is that model-based approaches can potentially lead to very significant reductions in mean-square error, with the reductions being highest for persistent (low value of λ) volatility processes with high values of $\xi\omega^{-2}$. Even for moderately large values of M the model-based predictor can be more accurate than realized volatility, sometimes by a considerable amount. This is an important result from a forecasting viewpoint. However, when there is not much persistence and M is very large, this result is reversed and realized volatility can be moderately more accurate. The smoother is always substantially more accurate than realized volatility, even when M is very large and there is not much memory in volatility. This suggests that model-based methods may be particularly helpful in estimating historical records of actual volatility. Finally, we should place some *caveats* on these conclusions. These results represent a somewhat favourable set-up for the model-based approach. In these calculations we have assumed knowledge of the second-order properties of volatility whereas in

Table 1. Exact mean-square error (steady state) of the estimators of actual volatility†

M	*Results for $\xi = 0.5$ and $\xi\omega^{-2} = 8$*			*Results for $\xi = 0.5$ and $\xi\omega^{-2} = 4$*			*Results for $\xi = 0.5$ and $\xi\omega^{-2} = 2$*		
	Smoother	*Predictor*	*$\{y\}_n$*	*Smoother*	*Predictor*	*$\{y\}_n$*	*Smoother*	*Predictor*	*$\{y\}_n$*
$\exp(-\Delta\lambda) = 0.99$									
1	0.0134	0.0226	0.624	0.0209	0.0369	0.749	0.0342	0.0625	0.998
12	0.00383	0.00792	0.0520	0.00586	0.0126	0.0624	0.00945	0.0211	0.0833
48	0.00183	0.00430	0.0130	0.00276	0.00692	0.0156	0.00440	0.0116	0.0208
288	0.000660	0.00206	0.00217	0.000967	0.00343	0.00260	0.00149	0.00600	0.00347
$\exp(-\Delta\lambda) = 0.9$									
1	0.0345	0.0456	0.620	0.0569	0.0820	0.741	0.0954	0.148	0.982
12	0.0109	0.0233	0.0520	0.0164	0.0396	0.0624	0.0259	0.0697	0.0832
48	0.00488	0.0150	0.0130	0.00707	0.0260	0.0156	0.0108	0.0467	0.208
288	0.00144	0.00966	0.00217	0.00195	0.0178	0.00260	0.00280	0.0338	0.00347

†The first two estimators are model based (smoother and predictor) and the third is model free (realized volatility $\{y\}_n$). These measures are calculated for different values of $\omega^2 = \mathrm{var}\{\sigma^2(t)\}$ and λ, keeping $\xi = E\{\sigma^2(t)\}$ fixed at 0.5.

practice we shall have to build such a model and then to estimate it, inducing additional biases that we have not reported on.

3.2. Estimating parameters: a numerical illustration

Estimating the parameters of continuous time SV models is known to be difficult owing to our inability to compute the appropriate likelihood function. This has prompted the development of a sizable collection of methods to deal with this problem. A very incomplete list of references includes Gourieroux *et al.* (1993), Gallant and Long (1997), Kim *et al.* (1998), Elerian *et al.* (2001) and Sørensen (2000). Here we study a simple approach based on the realized volatilities. The closest reference to ours in this respect is Bollerslev and Zhou (2001) who use a method-of-moments approach based on assuming that the actual volatility process $\{\sigma_n^2\}$ is observed via the quadratic variation estimator, i.e. they assume that there is no realized volatility error.

Table 2 shows the result of a small simulation experiment which investigates the effectiveness of the quasi-likelihood estimation methods based on the time series of realized volatility. The quasi-likelihood is constructed using the output of the Kalman filter. It is suboptimal for it does not exploit the non-Gaussian nature of the volatility dynamics, but it provides a consistent and asymptotically normal set of estimators. This follows from the fact that the Kalman filter builds the Gaussian quasi-likelihood function for the ARMA representation of the process, where the noise in the representation is both white and strong mixing (strong mixing follows from Sørensen (2000) and Genon-Catalot *et al.* (2000) who showed that if the volatility is strong mixing then squared returns are strong mixing). This means that we can immediately apply the asymptotic theory results of Francq and Zakoïan (2000) in this context so long as $\sigma^2(t)$ is strong mixing. Further the estimation takes only around 5 s on a Pentium III notebook computer.

The set-up of the simulation study uses 500 daily observations where the volatility is an OU process with a gamma marginal distribution. Table 2 varies the value of M. When $M=1$ this corresponds to using the classical approach of squared daily returns. When M is higher

Table 2. Monte Carlo estimates of the 0.1- and 0.9-quantiles of the maximum quasi-likelihood estimator of an SV model with OU volatility†

M	*Quantiles for the following parameter values:*					
	$\lambda = 0.01$	$\xi = 0.5$	$\omega^2 = \xi/8 = 0.0625$	$\lambda = 0.01$	$\xi = 0.5$	$\omega^2 = \xi/4 = 0.125$
1	0.00897, 1.76	0.318, 0.659	0.00751, 0.152	0.00750, 0.400	0.272, 0.752	0.0172, 0.225
12	0.00891, 0.0409	0.341, 0.669	0.0130, 0.0759	0.00789, 0.0367	0.265, 0.751	0.0197, 0.168
48	0.00920, 0.0348	0.339, 0.672	0.0134, 0.0715	0.00920, 0.0320	0.266, 0.727	0.0199, 0.149
288	0.00928, 0.0336	0.334, 0.674	0.0130, 0.0755	0.00906, 0.0299	0.269, 0.731	0.0207, 0.152
	$\lambda = 0.1$	$\xi = 0.5$	$\omega^2 = \xi/8 = 0.0625$	$\lambda = 0.1$	$\xi = 0.5$	$\omega^2 = \xi/4 = 0.125$
1	0.0451, 1.57	0.400, 0.573	0.0271, 0.151	0.0505, 0.312	0.374, 0.599	0.0548, 0.226
12	0.0725, 0.165	0.420, 0.572	0.0383, 0.0847	0.0713, 0.158	0.397, 0.593	0.0717, 0.170
48	0.0748, 0.152	0.421, 0.566	0.0397, 0.0829	0.0754, 0.148	0.398, 0.592	0.0763, 0.163
288	0.0792, 0.141	0.425, 0.572	0.0410, 0.0788	0.0755, 0.136	0.403, 0.619	0.0774, 0.176

†The volatility model has $\sigma^2(t) \sim \Gamma(\nu, a)$ with 500 daily observations, which implies $\xi = \nu a^{-1}$ and $\omega^2 = \nu a^{-2}$. The true value of ξ is always fixed at 0.5, while ω^2 and λ vary. M denotes the number of intraday observations used. 1000 replications are used in the study.

we are using intraday data. The results suggest that the intraday data allow us to estimate the parameters much more efficiently. Indeed when M is large the estimators have very little bias and turn out to be quite close to being jointly Gaussian. The experiment also suggests that when λ is larger, which corresponds to the process having less memory, then the estimates of ξ and ω^2 are sharper. Taken together the results are quite encouraging for they are based on only 2 years of data but suggest that we can construct quite precise estimates of these models with this.

4. Empirical illustration

To illustrate some of these results we have fitted a set of superposition-based OU or CEV SV models to the realized volatility time series constructed from the 5-min exchange rate return data discussed in Section 1. Here we use the quasi-likelihood method to estimate the parameters of the model: ξ, ω^2, $\lambda_1, \ldots, \lambda_J$ and $w_1, \ldots, w_J$. We do this for a variety of values of M, starting with $M = 6$, which corresponds to working with 4-h returns. The resulting parameter estimates are given in Table 3. For the moment we shall focus on this case.

The fitted parameters suggests a dramatic shift in the fitted model as we go from $J = 1$ to $J = 2$ or $J = 3$. The more flexible models allow for a factor which has quite a large degree of memory, as well as a more rapidly decaying component or two. A simple measure of fit of the model is the Box–Pierce statistic, which shows a large jump from a massive 302 when $J = 1$, down to an acceptable number for a superposition model.

To provide a more detailed assessment of the fit of the model we have drawn a series of graphs in Fig. 4. Except where explicitly noted we have computed the graphs using the $J = 3$ fitted model, although there would be very little difference if we had taken $J = 2$. Fig. 4(a) draws the computed realized volatility $\{y\}_n$, together with the corresponding smoothed estimate of actual volatility using the fitted SV model. We can see that realized volatility is much more jagged than the smoothed quantity. In Fig. 4(b) we have drawn a kernel-based estimate of the log-density of log(realized volatility). The bandwidths were taken to be $1.06\hat{\sigma}T^{-1/5}$, where T is the sample size and $\hat{\sigma}$ is the empirical standard deviation of log(realized volatility) (this is an optimal choice against a mean-square error loss for Gaussian data, e.g. Silverman (1986)) while we have chosen the range of the display to match the upper and lower 0.05% of the data—so trimming very little of the data. Andersen, Bollerslev, Diebold and Labys (2001) have suggested that the marginal

Table 3. Fit of the superposition of J volatility processes for an SV model based on realized volatility computed using $M = 6$, $M = 18$ and $M = 144$†

M	J	ξ	ω^2	λ_1	λ_2	λ_3	w_1	w_2	*Quasi-likelihood*	*BP*
6	3	0.4783	0.376	0.0370	1.61	246	0.212	0.180	−113258	11.2
6	2	0.4785	0.310	0.0383	3.76	—	0.262	—	−113261	11.3
6	1	0.4907	0.358	1.37	—	—	—	—	−117397	302
18	3	0.460	0.373	0.0145	0.0587	3.27	0.0560	0.190	−101864	26.4
18	2	0.460	0.533	0.0448	4.17	—	0.170	—	−101876	26.5
18	1	0.465	0.497	1.83	—	—	—	—	−107076	443
144	3	0.508	4.79	0.0331	0.973	268	0.0183	0.0180	−68377	15.3
144	2	0.509	0.461	0.0429	3.74	—	0.212	—	−68586	23.3
144	1	0.513	0.374	1.44	—	—	—	—	−76953	765

†We do not record w_J as this is 1 minus the sum of the other weights. The estimation method is quasi-likelihood using output from a Kalman filter. BP denotes the Box–Pierce statistic, based on 20 lags, which is a test of serial dependence in the scaled residuals.

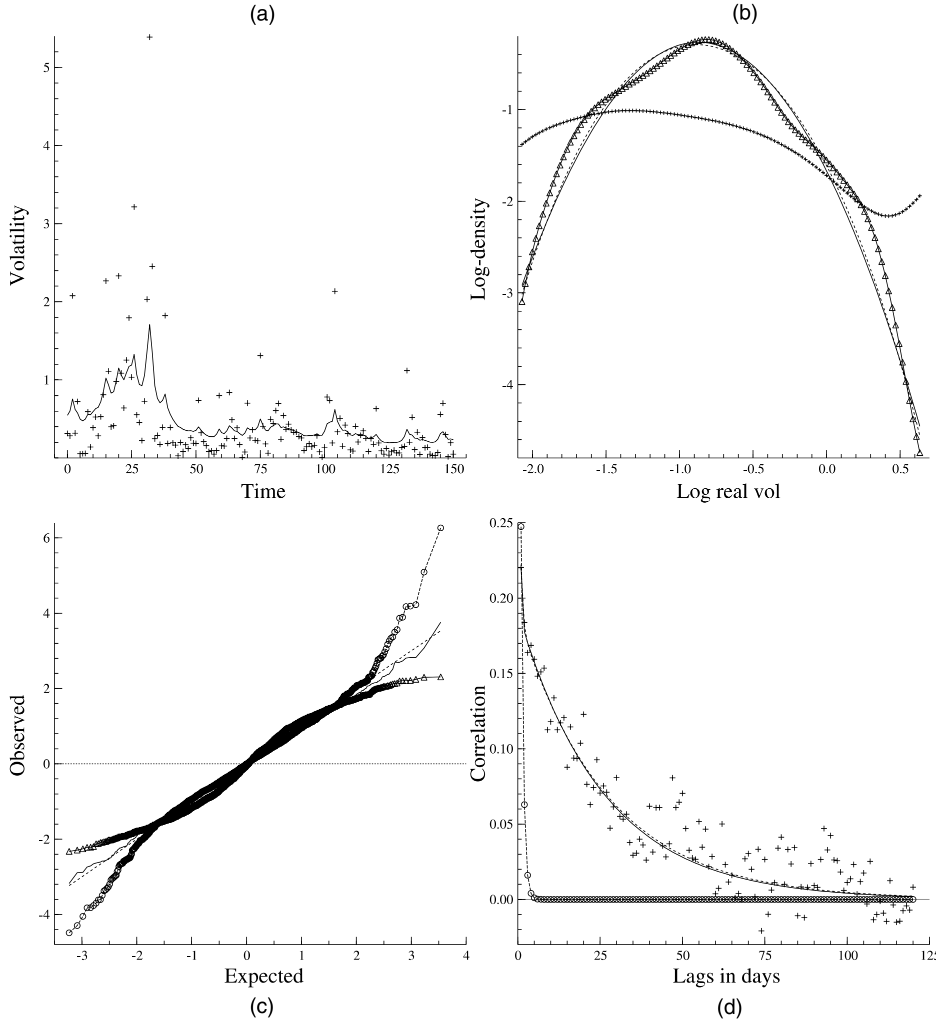

Fig. 4. Results from the SV model using $M=6$ (4-h returns): (a) first 150 observations of $\{y\}_n$ (+) and smoothed estimates of σ_n^2 (———); (b) kernel-based estimates of the density of $\log(\{y\}_n)$ (+) and log-smoothed $\log(\sigma_n^2)$ (Δ), and the log-normal (———) and inverse Gaussian (·······) fits; (c) *QQ*-plot for y_n scaled by estimated σ_n, using $\surd\{y_n\}$ (Δ), predicted (∘) and smoothed (———) volatility (·······, 45° line); (d) autocorrelation function of $\{y_n\}$ (+) and the fit of OU SV models with $J=1$ (∘), $J=2$ (———) and $J=3$ (·······)

distribution of realized volatility is closely approximated by a log-normal distribution when M is high, and that this would support a model for actual volatility which is log-normal. Such models go back to Clark (1973) and Taylor (1986). However, when we draw the corresponding fitted log-normal log-density, choosing the parameters by using maximum likelihood based on the smoothed realized volatilities as data, we see that the fit is poor. The same holds for the inverse Gaussian log-density. This is also drawn on Fig. 4(b), but is so close to the fit of the log-normal curve that it is extremely difficult to tell the difference between the two curves. Inverse Gaussian models for volatility were suggested by Barndorff-Nielsen and Shephard (2001a). The rejection of the log-normal and inverse Gaussian marginal distributions for realized volatility itself seems conclusive here. However, when we carry out the same action on the smoothed realized volatilities this rejection no longer holds, implying that the realized volatility error matters greatly here. The kernel-based estimate of the log-density of the logarithmic smoothed estimates is very much in line with the log-normal or inverse Gaussian hypothesis. This seems to extend the observations of Andersen, Bollerslev, Diebold and Labys (2001) in at least two directions:

(a) our model-based estimated actual volatility is fitted well, not just by the log-normal distribution, but equally well by the inverse Gaussian distribution;
(b) by using a model-based smoother the above stylized fact can be deduced using quite a low value of M.

Of course we have yet to see whether these results continue to hold as M increases.

Fig. 4(c) draws a QQ-plot of returns y_n divided by a number of estimates of σ_n. If the SV model holds correctly and there is no measurement error then these variables should be Gaussian and the QQ-plot should appear on a 45° line. Fig. 4(c) indicates that when we scale returns by realized volatility the returns are highly non-Gaussian, whereas when we use the smoothed estimate then the model seems to fit extremely well. If we replace the smoothed estimate by the predictor of actual volatility, then we see that the fit is as poor as the plot based on the realized volatility. Overall, Fig. 4(c) again confirms the fit of the model, which suggests that when $M = 6$ the difference between realized and smoothed volatility is important.

Fig. 4(d) shows the corresponding autocorrelation function for the realized volatility series together with the corresponding empirical correlogram. We see from Fig. 4(d) that when $J = 1$ we are entirely unable to fit the data, as its autocorrelation function starts at around 0.6 and then decays to 0 in a couple of days. A superposition of two processes is much better, picking up the longer-range dependence in the data. The superpositions of two and three processes give very similar fits; indeed in the graph they are hardly distinguishable.

We next ask how these results vary as M increases. We reanalyse the situation when $M = 144$, which corresponds to working with 10-min returns. Table 3 contains the estimated parameters for this problem. They suggest that moving to a superposition of three processes has an important effect on the fit of the model. Again the fitted models indicate that the volatility has elements which have a substantial memory, whereas other components are much more transitory. An important feature of Table 3 is the jump in the value of the estimated ω^2 when we move to having $J = 3$. This is caused by the third component which has a very high value of λ, which does not overly change the variance of the actual volatility.

The fit of the model can also be seen from Fig. 5. This broadly shows the same results as Fig. 4 except for the following. Realized volatility is now less jagged, and so the smoothed estimator of actual volatility and realized volatility are much more in line. The plots of the estimated log-densities show that realized and smoothed volatilities are again closer, with both being quite well fitted by the log-normal and inverse Gaussian distributions. The smoothed estimators are still more closely approximated than the realized volatilities, however. The QQ-plots for realized and smoothed volatility are roughly similar, whereas the plot for prediction is still not satisfactory. This indicates that the uncertainty of predicting volatility 1 day ahead is substantial. Finally, the autocorrelation functions show an improvement in fit as we go from $J = 2$ to $J = 3$ in the SV model.

We finish this section by briefly repeating this exercise with an intermediate value of M, taking $M = 18$, which corresponds to working with returns calculated over 80-min periods. The results are given in Table 3. They, and the corresponding plots not reproduced here, are very much in

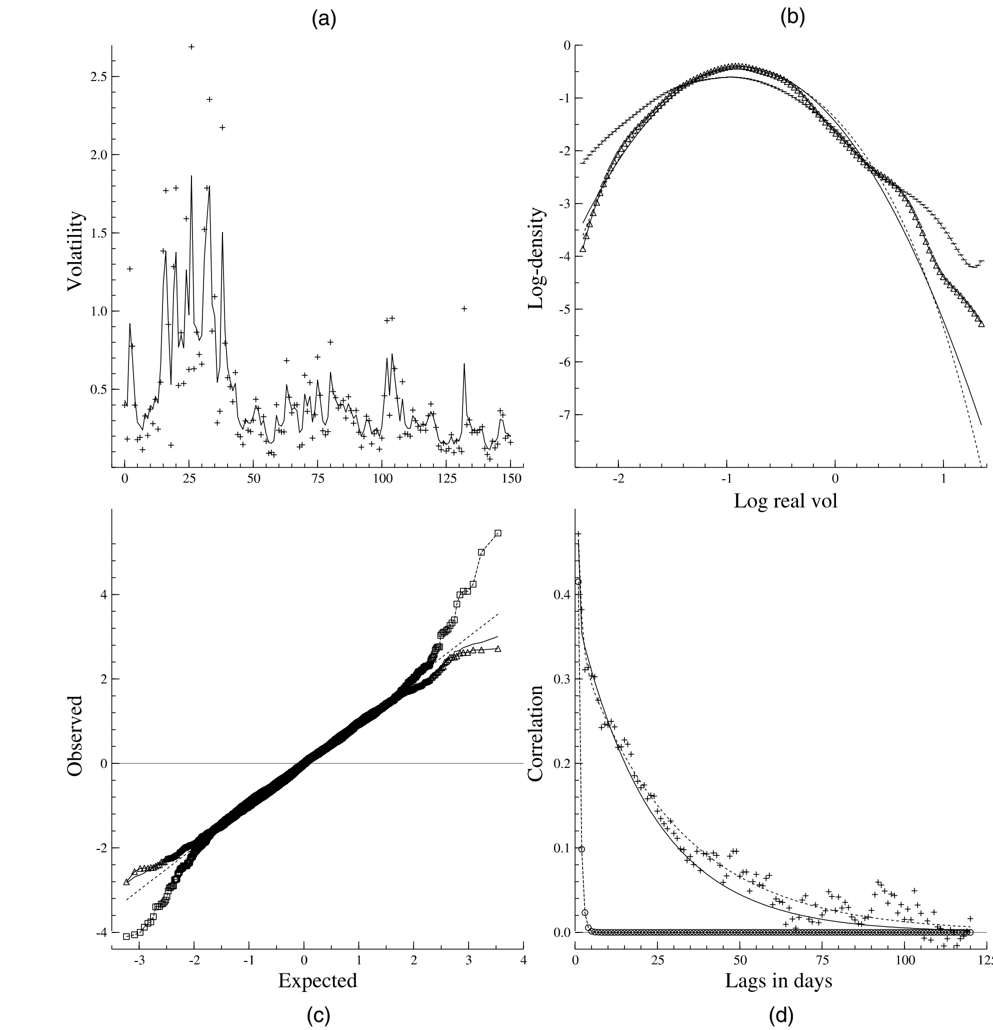

Fig. 5. Results from the fit of the SV model using $M=144$ (10-min returns): (a) $\{y_n\}$ (+) together with the smoothed estimator of σ_n^2 (———); (b) estimates of the log-density $\log(\{y\}_n)$ (_–‾) and smoothed estimates of $\log(\sigma_n^2)$ (Δ), and the log-normal (———) and inverse Gaussian (·······) fits; (c) *QQ*-plot for the standardized returns (———, smoothed; $\square$, prediction; Δ, $\{y\}_n$; ·······, 45° line); (d) autocorrelation function of $\{y\}_n$ (+) and the fit of OU SV models with $J=1$ (∘), $J=2$ (———) and $J=3$ (·······)

line with the previous graphs with the smoothed estimates of actual volatility performing well, although the *QQ*-plot is not as good as it was when we used 4-h data.

5. Asymptotic distribution of realized volatility error

5.1. The theory

In Section 2 we derived the mean and variance of the realized volatility error for a continuous time SV model when $\mu=\beta=0$. Although it is possible to derive the corresponding result when $\mu \neq 0$ but $\beta=0$, adapting to the risk premium case seems difficult. Instead we take an asymptotic route. In this section we obtain a limit theory for

$$(\{y\}_n - \sigma_n^2)\sqrt{M},$$

which covers the case of a drift and risk premium.

Theorem 1. For the SV model (1) suppose that the volatility process σ^2 is of locally bounded variation (i.e. with probability 1 the paths of σ^2 are of bounded variation on any compact subinterval of $[0, \infty)$). Then, for any positive Δ and $M \to \infty$

$$(\{y\}_n - \sigma_n^2) \Big/ \sqrt{\left\{2\sum_{j=1}^{M}(\sigma_{j,n}^2)^2\right\}} \xrightarrow{\mathcal{L}} N(0, 1), \tag{17}$$

where

$$\sigma_{j,n}^2 = \sigma^{2*}\left\{(n-1)\Delta + \frac{\Delta j}{M}\right\} - \sigma^{2*}\left\{(n-1)\Delta + \frac{\Delta(j-1)}{M}\right\}.$$

Furthermore,

$$\Delta^{-1} M \sum_{j=1}^{M}(\sigma_{j,n}^2)^2 \to \sigma_n^{[4]} \qquad \text{almost surely} \tag{18}$$

where

$$\sigma_n^{[4]} = \int_{(n-1)\Delta}^{n\Delta} \sigma^4(s)\,\mathrm{d}s$$

In particular, then, the limiting law of $\sqrt{M}(\{y\}_n - \sigma_n^2)$ is a normal variance mixture.

Proof. The proof of theorem 1 is given in Appendix A.

This theorem implies that

$$\sqrt{M}(\{y\}_n - \sigma_n^2)|\sigma_n^{[4]} \xrightarrow{\mathcal{L}} N(0, 2\Delta\sigma_n^{[4]}), \tag{19}$$

which has the important implication that we can strengthen the usual quadratic variation result that the drift and risk premium has no effect on the limit of $\{y\}_n$ to the result that the asymptotic distribution of $(\{y\}_n - \sigma_n^2)\sqrt{M}$ does not depend on μ or β. Thus the effect on realized volatility of the drift and risk premium is of only *third order*, which suggests that it may be safe to ignore it in many cases.

Theorem 1 also implies that we would expect the variance of the realized volatility error to depend positively on the level of volatility. We have already seen an example of this in the simulated data in Fig. 2. Further, the marginal distribution of the realized volatility error should be thicker than normal owing to the normal variance mixture (19) averaged over the random $\sigma_n^{[4]}$. We call $\sigma_n^{[4]}$ the *actual quarticity*, whereas the associated $\sigma^4(t)$ is the *spot quarticity*.

We should note that

$$\sum_{j=1}^{M}(\sigma_{j,n}^2)^2$$

is the realized volatility of $\sigma^{2*}(t)$. Of course the limit, as $M \to \infty$, of this realized volatility is 0; however, theorem 1 shows that the scaled

$$M \sum_{j=1}^{M} (\sigma_{j,n}^2)^2$$

has a stochastic limit. This is a special case of a more general lemma that we prove in Appendix A on what we call *power variation*.

Lemma 1. Assume that $\tau(t)$ is of locally bounded variation. Then, for $M \to \infty$ and r a positive integer,

$$\Delta^{-r+1} M^{r-1} \sum_{j=1}^{M} \tau_j^r \to \tau^{r*}(\Delta) \qquad \text{almost surely} \tag{20}$$

where

$$\tau_j = \tau^*(jM^{-1}\Delta) - \tau^*\{(j-1)M^{-1}\Delta\}$$

and

$$\tau^{r*}(t) = \int_0^t \tau^r(s)\,\mathrm{d}s.$$

The proof of lemma 1 is given in Appendix A.

We report fully on the implications of this result, and various possible extensions, in Barndorff-Nielsen and Shephard (2001c). One of these extensions is that, again writing

$$y_{j,n} = y^*\{(n-1)\Delta + j\Delta M^{-1}\} - y^*\{(n-1)\Delta + (j-1)\,\Delta M^{-1}\},$$

it can be shown that the *realized quarticity*

$$\frac{1}{3} M\Delta^{-1} \sum_{j=1}^{M} y_{j,n}^4 \to \sigma_n^{[4]} \qquad \text{almost surely,}$$

which is also the limit for $M \to \infty$ of $M\Delta^{-1} \Sigma_{j=1}^{M} (\sigma_{j,n}^2)^2$. Consequently, the former—known—sum can be used instead of the latter—unknown—sum in the denominator on the left-hand side of the key limiting result (17). In particular

$$(\{y\}_n - \sigma_n^2) \Big/ \sqrt{\left(\frac{2}{3} \sum_{j=1}^{M} y_{j,n}^4\right)} \xrightarrow{\mathcal{L}} N(0, 1).$$

Following the first draft of this paper we have used Monte Carlo methods to study the finite sample behaviour of this asymptotic approximation. Results are reported in Barndorff-Nielsen and Shephard (2001b). These experiments suggest that we need quite large values of M for the result to be reliable; however, a better performance is obtained by transforming the approximation on to the log-scale. Then the approximation becomes

$$\{\log(\{y\}_n) - \log(\sigma_n^2)\} \Bigg/ \sqrt{\left\{ \frac{2}{3} \frac{\sum_{j=1}^{M} y_{j,n}^4}{\left(\sum_{j=1}^{M} y_{j,n}^2\right)^2} \right\}} \xrightarrow{\mathcal{L}} N(0, 1).$$

This seems to be quite accurate even for moderate values of M. Following the developments of this paper, our further work on power variation has recently been reported in Barndorff-Nielsen and Shephard (2001c).

Finally we note that theorem 1 requires that τ is of locally bounded variation. In the OU case this is easily checked for we know that

$$\tau(t) = \exp(-\lambda t)\,\tau(0) + \int_0^t \exp\{-\lambda(t-s)\}\;\mathrm{d}z(\lambda s).$$

5.2. Application

Suppose that our interest is in estimating μ and β, knowing that

$$y_n|\sigma_n^2 \sim N(\mu + \beta\sigma_n^2, \sigma_n^2).$$

A naïve approach would be to regress returns on a constant and the sequence of feasible realized volatilities to produce a simple regression-based estimator. Such an estimator will be both biased and inconsistent owing to an *errors-in-variables* effect of mismeasuring actual volatility by using realized volatility (see Hendry (1995), chapter 12, for a discussion of this in a historical context). The bias is determined by the variance of $\{y\}_n - \sigma_n^2$, which we have seen is $O(M^{-1})$ even in the presence of drift and risk premium. A smaller bias would result if we use a model-based estimator of actual volatility, instead of the simpler realized volatility. We saw in Section 2 that this can substantially reduce the variance, and this will carry over to the bias reduction of the regression-based estimator.

An alternative strategy is to employ an instrumental variable approach. This requires us to find an estimator of σ_n^2 which does not rely on data at time n but is correlated with σ_n^2. A model-free candidate for this task is

$$\widehat{\sigma_n^2} = \tfrac{1}{2}(\{y\}_{n-1} + \{y\}_{n+1}),$$

the average of contiguous realized volatilities. If actual volatility is temporally dependent, the theory developed in Section 2 shows that $\widehat{\sigma_n^2}$ will be correlated with σ_n^2 and so is a valid instrument. Of course, within the context of a model, the best instrument will be the jack-knife estimator, which is the best linear estimator of σ_n^2 given

$$\{y\}_1\,,\ \{y\}_2\,,\ \ldots,\ \{y\}_{n-1}\,,\ \{y\}_{n+1}\,,\ \ldots,\ \{y\}_T\,.$$

This is readily computed for models which can be placed within a linear state space form. A final approach to dealing with this issue is to append an extra measurement equation to the linear state space form (16) and to estimate μ and β at the same time as other parameters in a fully specified model.

6. Extensions

6.1. Diurnal effects and actual volatility

An important aspect of the realized volatility series is that it is not very sensitive to the substantial and complicated intraday diurnal pattern in volatility that is found in many empirical studies (e.g. Andersen and Bollerslev (1997b, 1998b)) as well as being clear from Fig. 1(a). To understand this it is helpful to think of the spot volatility as the sum of a (potentially unknown) deterministic diurnal component $\sigma^2_\psi\{\mathrm{mod}(t,\Delta)\}$ where Δ represents a day, plus a stochastic process $\sigma^2_\lambda(t)$; then we have

$$\sigma^2(t) = \sigma^2_\psi\{\mathrm{mod}(t,\Delta)\} + \sigma^2_\lambda(t).$$

Hence in this model the spot volatility has a repeating intraday (i.e. diurnal) component, but does not have a day of the week or monthly seasonal component. As a result

$$\sigma^2_n = c + \sigma^2_{n,\,\lambda},$$

where

$$c = \int_0^\Delta \sigma^2_\psi\{\mathrm{mod}(u,\Delta)\}\,\mathrm{d}u$$

and

$$\sigma^2_{n,\,\lambda} = \sigma^{2*}_\lambda(n) - \sigma^{2*}_\lambda(n-1),$$

and

$$\sigma^{2*}_\lambda(t) = \int_0^t \sigma^2_\lambda(u)\,\mathrm{d}u.$$

In this structure the dynamics of realized volatility are unaffected by a diurnal effect. Of course, in practice this additive structure should be regarded as holding only approximately, in which case the diurnal effect may not be completely ignorable. However, in this paper we shall neglect this deficiency.

6.2. *Leverage*

This analysis has not included a leverage term in the model. This can be added in various ways. We follow Barndorff-Nielsen and Shephard (2001a) in parameterizing the effect as

$$\mathrm{d}y^*(t) = \{\mu + \beta\,\sigma^2(t)\}\,\mathrm{d}t + \sigma(t)\,\mathrm{d}w(t) + \rho\,\mathrm{d}\bar{z}(\lambda t),$$

where we assume that $\bar{z}(t) = z(t) - E\{z(t)\}$ and $z(t)$ is a Lévy process potentially correlated with $\sigma^2(t)$. The corresponding quadratic variation for this process is

$$[y^*](t) = \sigma^{2*}(t) + \rho^2[\bar{z}](\lambda t),$$

whereas the realized volatility error

$$u_n = (\{y^0\}_n - \sigma^2_n) + \rho^2(\{\bar{z}\}_n - [\bar{z}]_n) + 2\rho c_n,$$

where

$$y^0(t) = \int_0^t \sigma(t)\,\mathrm{d}w(t),$$

$$c_n = \sum_{j=1}^M \bar{z}_{j,n}\varepsilon_{j,n}\sigma_{j,n}$$

$$\{y^0\}_n - \sigma_n^2 = \sum_{j=1}^{M} \sigma_{j,n}^2 (\varepsilon_{j,n}^2 - 1),$$

using the generic realized volatility notation that was developed in Section 2. Here the three terms which make up u_n are zero meaned and uncorrelated when we assume that $\mu = \beta = 0$. The only task is to calculate the variances of each of the terms.

The new terms are straightforward to study once we have the following lemma which relates a Lévy process to its quadratic variation and realized volatility.

Lemma 2. Let t be fixed and let $z(t)$ be a Lévy process with finite fourth cumulant. Then, defining

$$\{z\}(t) = \sum_{j=1}^{M} [z(jtM^{-1}) - z\{(j-1)tM^{-1}\}]^2,$$

we have that

$$E\begin{Bmatrix} z(t) \\ \{z\}(t) \\ [z](t) \end{Bmatrix} = t\begin{pmatrix} \kappa_1 \\ \kappa_2 \\ \kappa_2 \end{pmatrix},$$

$$\operatorname{cov}\begin{Bmatrix} z(t) \\ \{z\}(t) \\ [z](t) \end{Bmatrix} = t\begin{pmatrix} \kappa_2 & \kappa_3 & \kappa_3 \\ \kappa_3 & \kappa_4 + 3\kappa_2^2 tM^{-1} & \kappa_4 \\ \kappa_3 & \kappa_4 & \kappa_4 \end{pmatrix},$$

where κ_r denotes the rth cumulant of $z(1)$. An implication is that $\{z\}(t) - [z](t)$ has zero mean, whereas

$$\operatorname{var}\left\langle \{z\}(t) - [z](t) \right\rangle = 3\kappa_2^2 t^2 M^{-1}.$$

Proof. Most of the results follow immediately, recognizing that the rth cumulant of $z(t)$ is $t\kappa_r$. This is a consequence of the Lévy–Khintchine representation. The only piece of this result which is not trivial is

$$\begin{aligned} \operatorname{var}[\{z\}(t)] &= M\,\mu_4\{z(tM^{-1})\} \\ &= M[\kappa_4\{z(tM^{-1})\} + 3\,\kappa_2\{z(tM^{-1})\}^2] \\ &= M(tM^{-1}\kappa_4 + 3t^2M^{-2}\kappa_2^2). \end{aligned}$$

Here $\mu_4\{\cdot\}$ and $\kappa_4\{\cdot\}$ denote the fourth centred moment and cumulant respectively. □

We achieve our desired result immediately for

$$M\operatorname{cov}\begin{pmatrix} \{\bar{z}\}_n - [\bar{z}]_n \\ \{y^0\}_n - \sigma_n^2 \\ c_n \end{pmatrix} \to \Delta^2 \begin{pmatrix} 3\kappa_2^2 & 0 & 0 \\ 0 & 2(\omega^2 + \xi^2) & 0 \\ 0 & 0 & \kappa_2\xi \end{pmatrix},$$

as $M \to \infty$. Repeating the pattern we had before, no new issues arise when we allow for a drift or a risk premium for their effect will be small compared with the other terms. Of course, the central limit theory that we developed in Section 5 will apply to $(\{y^0\}_n - \sigma_n^2)\sqrt{M}$ not to $(\{y\}_n - [y]_n)\sqrt{M}$.

Trivially this analysis also deals with the situation of a model which is an SV process plus jumps, where the volatility is not correlated with the jumps.

7. Conclusion

In this paper we have studied the statistical properties of realized volatility in the context of SV models. Our results are entirely general, providing both a central limit theory approximation as well as an exact second-order analysis. These results can be used, in conjunction with a model for the dynamics of volatility, to produce a more accurate estimate of actual volatility. Further, a simple quasi-likelihood results which could be used to perform computationally quite simple estimation. Potentially they allow us to exploit the availability of high frequency data in financial economics, giving us relatively simple and efficient ways of estimating these stochastic processes.

Finally, in our empirical work we have taken Δ to be 1 day. This choice is entirely *ad hoc*. Another possibility is to look simultaneously at several different Δ-values. This may have virtue as a way of checking the fit of the model, as well as allowing potentially more efficient estimation. However, we have yet to explore this issue. To do so it would be convenient to have a functional limit theorem for $(\{y\}_n - \sigma_n^2)$.

Acknowledgements

Barndorff-Nielsen's work is supported by the Centre for Analytical Finance, which is funded by the Danish Social Science Research Council, and by the Centre for Mathematical Physics and Stochastics, which is funded by the Danish National Research Foundation. Neil Shephard's research is supported by the UK's Economic and Social Research Council through grant R00023839. All the calculations made in this paper are based on software written by the second author using the `Ox` language of Doornik (2001) combined with the state space software library described by Koopman *et al.* (1999). We thank Michel M. Dacorogna for allowing us to use Olsen's high frequency exchange rate data in our study and Tim Bollerslev for supplying us with a semicleaned version of these data. The comments of Tim Bollerslev, Frank Gerhard, Nour Meddahi, Enrique Sentana and the two excellent referees on an initial draft were particularly helpful.

Appendix A

This appendix has three subsections. First we discuss some of the aspects of the data that we use in the paper. Second we give a proof of lemma 1, whereas in the third subsection we prove theorem 1.

A.1. Data appendix

The Olsen group have kindly made available to us a data set which records every 5 min the most recent quote to appear on the Reuters screen from December 1st, 1986, until November 30th, 1996. When prices are missing they have interpolated them. Details of this processing are given in Dacorogna *et al.* (2001). The same data set was analysed by Andersen, Bollerslev, Diebold and Labys (2001). We follow the extensive work of Torben Andersen and Tim Bollerslev on this data set, who removed much of the times when the market is basically closed. This includes almost all of the week-ends, and they have taken out most US holidays. The result is what we shall regard as a single time series of length 705 313 observations. Although many of the breaks in the series have been removed, sometimes there are sequences of very small price changes caused by, for example, unmodelled non-US holidays or data feed breakdowns. We deal with this by adding a Brownian bridge simulation to sequences of data where at each time point the absolute change in a 5-min period is below 0.01%, i.e., when this happens, we interpolate prices stochastically, adding a Brownian bridge with a standard deviation of 0.01 for each time period. By using a bridge process we are not affecting the long run trajectory of prices, and the effect on realized volatility is very small indeed. We have used this stochastic method here to be consistent with our other work on this topic where this effect is important. It is illustrated in Fig. 6, which shows the first 500 observations in the US dollar–

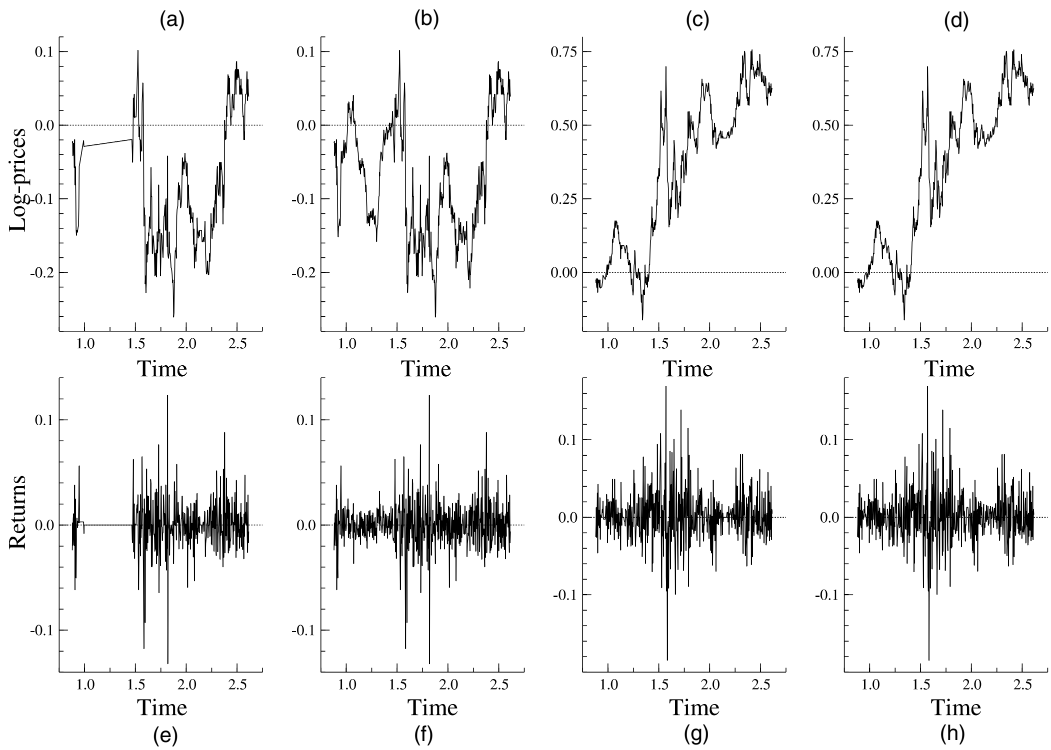

Fig. 6. (a) Raw yen–dollar data; (b) interpolated yen–dollar data; (c) raw Deutschmark–dollar data; (d) interpolated Deutschmark–dollar data; (e) returns for the raw yen–dollar data; (f) returns for the interpolated yen–dollar data; (g) returns for the raw Deutschmark–dollar data; (h) returns for the interpolated Deutschmark–dollar data (interpolation is by a Brownian bridge interpolator; the *x*-axes are in days)

Deutschmark series that we have used in this paper and another series which is for the yen–dollar comparison. Later stretches of the data have fewer breaks in them, but this graph illustrates the effects of our intervention. Clearly our approach is *ad hoc*. However, a proper statistical modelling of these breaks is very complicated because of their many causes and the fact that our data set is enormous.

A.2. Proof of lemma 1

Recall for the process τ that we use the notation

$$\tau^*(t) = \int_0^t \tau(s)\,\mathrm{d}s$$

and

$$\tau_j = \tau^*(jM^{-1}\Delta) - \tau^*\{(j-1)M^{-1}\Delta\}.$$

Proof. By the definition of τ_j, for every j there is a c_j such that

$$\inf_{(j-1)M^{-1}\Delta \leqslant s \leqslant jM^{-1}\Delta}\{\tau(s)\} \leqslant c_j \leqslant \sup_{(j-1)M^{-1}\Delta \leqslant s \leqslant jM^{-1}\Delta}\{\tau(s)\}$$

and

$$\tau_j = c_j \Delta M^{-1}. \tag{21}$$

The local bounded variation of τ implies that τ^r is locally bounded and Riemann integrable. Consequently

$$\Delta^{-r+1} M^{r-1} \sum_{j=1}^{M} \tau_j^r = \sum_{j=1}^{M} c_j^r \Delta M^{-1} \to \int_0^{\Delta} \tau^r(s)\,\mathrm{d}s = \tau^{r*}(\Delta). \qquad \square$$

The fact that τ^r is Riemann integrable is perhaps not immediately obvious. However, we recall that a bounded function f is Riemann integrable on an interval $[0, t]$ if and only if the set of discontinuity points of f has Lebesgue measure 0 (see Hobson (1927), pages 465–466, Munroe (1953), page 174, theorem 24.4, or Lebesgue (1902)). In our case the latter property follows immediately from the bounded variation of τ (any function of bounded variation is the difference between an increasing and a decreasing function and any monotone function has at most countably many discontinuities).

A.3. Proof of theorem 1

We first recall some definitions. Consider the SV model

$$y^*(t) = \mu t + \beta\,\tau^*(t) + \int_0^t \tau^{1/2}(s)\,\mathrm{d}w(t),$$

with τ positive, stationary and independent of w (we have switched our notation for the volatility as it simplifies our later derivation). Now writing u and $\{y\}$ for u_1 and $\{y\}_1$ we have

$$u = \{y\} - \tau^*(\Delta) = \sum_{j=1}^{M} y_j^2 - \tau^*(\Delta)$$

where

$$y_j = y^*(jM^{-1}\Delta) - y^*\{(j-1)M^{-1}\Delta\}.$$

Conditionally on $\tau_1, \ldots, \tau_M$, the increments $y_1, \ldots, y_M$ are independent, and

$$y_j \overset{\mathcal{L}}{=} N(\mu M^{-1}\Delta + \beta\tau_j, \tau_j).$$

Thus, conditionally, y_j^2 is non-central χ^2 distributed with cumulant function

$$C\{\zeta \ddagger y_j^2 | \tau_j\} = -\tfrac{1}{2}\log(1 - 2\mathrm{i}\tau_j\zeta) + \mathrm{i}\nu_j\zeta(1 - 2\mathrm{i}\tau_j\zeta)^{-1}$$

where $\mathrm{i} = \sqrt{-1}$ and

$$\nu_j = (\mu M^{-1}\Delta + \beta\tau_j)^2. \tag{22}$$

Consequently

$$C\{\zeta \ddagger u | \tau_1, \ldots, \tau_M\} = -\sum_{j=1}^{M} \{\tfrac{1}{2}\log(1 - 2\mathrm{i}\tau_j\zeta) - \mathrm{i}\nu_j\zeta(1 - 2\mathrm{i}\tau_j\zeta)^{-1} + \mathrm{i}\tau_j\zeta\}.$$

By Taylor's formula with remainder (see, for instance, Barndorff-Nielsen and Cox (1989), formula 6.122) we find, provided that

$$2|\zeta| \max_{1 \leqslant j \leqslant M}(\tau_j) < 1,$$

that

$$\tfrac{1}{2}\log(1 - 2\mathrm{i}\tau_j\zeta) - \mathrm{i}\nu_j\zeta(1 - 2\mathrm{i}\tau_j\zeta)^{-1} + \mathrm{i}\zeta\tau_j = \zeta^2\{\tau_j^2\, Q_{0j}(\zeta) + 2\nu_j\tau_j\, Q_{1j}(\zeta)\} - \mathrm{i}\nu_j\zeta,$$

where

$$Q_{0j}(\zeta) = 2\int_0^1 \frac{1-s}{(1 - 2\mathrm{i}\tau_j\zeta s)^2}\,\mathrm{d}s$$

and

$$Q_{1j}(\zeta) = 2\int_0^1 \frac{1-s}{(1-2\mathrm{i}\tau_j\zeta s)^3}\,\mathrm{d}s.$$

Hence

$$C\{\zeta\ddagger u|\tau_1,\ldots,\tau_M\} = \mathrm{i}\zeta\sum_{j=1}^M \nu_j - \zeta^2\sum_{j=1}^M\{\tau_j^2\,Q_{0j}(\zeta) + 2\nu_j\tau_j\,Q_{1j}(\zeta)\}. \tag{23}$$

Proof. Note first that expression (18) follows from lemma 1. Next, rewrite equation (23) as

$$\begin{aligned}C\{\zeta\ddagger u|\tau_1,\ldots,\tau_M\} &= \mathrm{i}\zeta\sum_{j=1}^M \nu_j - \zeta^2\sum_{j=1}^M(\tau_j^2+2\nu_j\tau_j) - \zeta^2\sum_{j=1}^M[\tau_j^2\{Q_{0j}(\zeta)-1\}+2\nu_j\tau_j\{Q_{1j}(\zeta)-1\}]\\ &= \frac{1}{2}\zeta^2\times 2\sum_{j=1}^M\tau_j + R(\zeta),\end{aligned}$$

where

$$R(\zeta) = \mathrm{i}\zeta\sum_{j=1}^M\nu_j - 2\zeta^2\sum_{j=1}^M\nu_j\tau_j - \zeta^2\sum_{j=1}^M[\tau_j^2\{Q_{0j}(\zeta)-1\}+2\nu_j\tau_j\{Q_{1j}(\zeta)-1\}].$$

Thus, to verify expression (17) we must show that

$$\sum_{j=1}^M\nu_j\bigg/\sqrt{\sum_{j=1}^M\tau_j^2}\to 0,$$

$$\sum_{j=1}^M\nu_j\tau_j\bigg/\sum_{j=1}^M\tau_j^2\to 0,$$

$$\sum_{j=1}^M\tau_j^2\left[Q_{0j}\left\{\zeta\bigg/\sqrt{\left(2\sum_{j=1}^M\tau_j^2\right)}\right\}-1\right]\bigg/\sum_{j=1}^M\tau_j^2\to 0$$

and

$$\sum_{j=1}^M\nu_j\tau_j\left[Q_{1j}\left\{\zeta\bigg/\sqrt{\left(2\sum_{j=1}^M\tau_j^2\right)}\right\}-1\right]\bigg/\sum_{j=1}^M\tau_j^2\to 0$$

or, equivalently, by expression (20), that

$$\surd M\sum_{j=1}^M\nu_j\to 0,$$

$$M\sum_{j=1}^M\nu_j\tau_j\to 0,$$

$$M\sum_{j=1}^M\tau_j^2\left[Q_{0j}\left\{\zeta\bigg/\sqrt{\left(2\sum_{j=1}^M\tau_j^2\right)}\right\}-1\right]\to 0 \tag{24}$$

and

$$M\sum_{j=1}^{M}\nu_j\tau_j\left[Q_{1j}\left\{\zeta\Big/\sqrt{\left(2\sum_{j=1}^{M}\tau_j^2\right)}\right\}-1\right]\to 0. \tag{25}$$

We have

$$\sqrt{M}\sum_{j=1}^{M}\nu_j = M^{-1/2}\left(\Delta^2\mu^2+2\mu\Delta\beta\sum_{j=1}^{M}\tau_j+\beta^2 M\sum_{j=1}^{M}\tau_j^2\right),$$

which tends to 0 on account of expression (20). Furthermore, also by expression (20) we find that

$$M\sum_{j=1}^{M}\nu_j\tau_j = M^{-1}\mu^2\Delta^2\sum_{j=1}^{M}\tau_j+2\mu\Delta\beta\sum_{j=1}^{M}\tau_j^2+\beta^2 M\sum_{j=1}^{M}\tau_j^3\to 0.$$

Finally, to show expressions (24) and (25) we first note that, by equation (21), the local boundedness of τ and expression (20),

$$\tau_j\Big/\sqrt{\sum_{j=1}^{M}\tau_j^2}=\sqrt{M}\tau_j\Big/\sqrt{\left(M\sum_{j=1}^{M}\tau_j^2\right)}=M^{-1/2}\Delta c_j\Big/\sqrt{\left(M\sum_{j=1}^{M}\tau_j^2\right)}=O(M^{-1/2})$$

uniformly in j. Hence

$$Q_{0j}\left\{\zeta\Big/\sqrt{\left(2\sum_{j=1}^{M}\tau_j^2\right)}\right\}-1\to 0 \tag{26}$$

and

$$Q_{1j}\left\{\zeta\Big/\sqrt{\left(2\sum_{j=1}^{M}\tau_j^2\right)}\right\}-1\to 0 \tag{27}$$

uniformly in j. Moreover, again using expression (20), we have

$$M\sum_{j=1}^{M}(\tau_j^2+\nu_j\tau_j)=(1+\Delta^2\mu^2)M\sum_{j=1}^{M}\tau_j^2+2\Delta\mu\beta M\sum_{j=1}^{M}\tau_j^3+\beta^2 M\sum_{j=1}^{M}\tau_j^4=O(1)$$

and expressions (24) and (25) follow from this and expressions (26) and (27).

References

Andersen, T. G. and Bollerslev, T. (1997a) Heterogeneous information arrivals and return volatility dynamics: uncovering the long-run in high frequency returns. *J. Finan.*, **52**, 975–1005.

——— (1997b) Intraday periodicity and volatility persistence in financial markets. *J. Emp. Finan.*, **4**, 115–158.

——— (1998a) Answering the skeptics: yes, standard volatility models do provide accurate forecasts. *Int. Econ. Rev.*, **39**, 885–905.

——— (1998b) Deutsche mark-dollar volatility: intraday activity patterns, macroeconomic announcements, and longer run dependencies. *J. Finan.*, **53**, 219–265.

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Ebens, H. (2001) The distribution of realised stock return volatility. *J. Finan. Econ.*, **61**, 43–76.

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P. (2000) Exchange rate returns standardised by realised volatility are (nearly) Gaussian. *Multinatn. Finan. J.*, **4**, 159–179.

——— (2001) The distribution of exchange rate volatility. *J. Am. Statist. Ass.*, **92**, 42–55.

Andreou, E. and Ghysels, E. (2001) Rolling-sampling volatility estimators: some new theoretical, simulation and empirical results. *J. Bus. Econ. Statist.*, **19**, in the press.

Areal, N. M. P. C. and Taylor, S. J. (2002) The realised volatility of FTSE-100 futures prices. *J. Fut. Markts*, **22**, in the press.

Bai, X., Russell, J. R. and Tiao, G. C. (2000) Beyond Merton's utopia: effects of non-normality and dependence on the precision of variance estimates using high-frequency financial data. Unpublished. Graduate School of Business, University of Chicago, Chicago.

Barndorff-Nielsen, O. E. (1998) Processes of normal inverse Gaussian type. *Finan. Stochast.*, **2**, 41–68.

——(2001) Superposition of Ornstein-Uhlenbeck type processes. *Theory Probab. Applic.*, **45**, 175–194.
Barndorff-Nielsen, O. E. and Cox, D. R. (1989) *Asymptotic Techniques for Use in Statistics.* London: Chapman and Hall.
Barndorff-Nielsen, O. E., Jensen, J. L. and Sørensen, M. (1990) Parametric modelling of turbulence. *Phil. Trans. R. Soc. Lond.*, **332**, 439–455.
Barndorff-Nielsen, O. E. and Shephard, N. (2001a) Non-Gaussian Ornstein–Uhlenbeck-based models and some of their uses in financial economics (with discussion). *J. R. Statist. Soc.* B, **63**, 167–241.
——(2001b) How accurate is the asymptotic approximation to the distribution of realised volatility? To be published.
——(2001c) Realised power variation and stochastic volatility. Unpublished. Nuffield College, Oxford.
Bollerslev, T., Engle, R. F. and Nelson, D. B. (1994) ARCH models. In *The Handbook of Econometrics*, vol. 4 (eds R. F. Engle and D. McFadden), pp. 2959–3038. Amsterdam: North-Holland.
Bollerslev, T. and Zhou, H. (2001) Estimating stochastic volatility diffusion using conditional moments of integrated volatility. *J. Econometr.*, to be published.
Christensen, B. J. and Prabhala, N. R. (1998) The relation between implied and realized volatility. *J. Finan. Econ.*, **37**, 125–150.
Clark, P. K. (1973) A subordinated stochastic process model with fixed variance for speculative prices. *Econometrica*, **41**, 135–156.
Comte, F. and Renault, E. (1998) Long memory in continuous-time stochastic volatility models. *Math. Finan.*, **8**, 291–323.
Cox, D. R. (1991) Long-range dependence, non-linearity and time irreversibility. *J. Time Ser. Anal.*, **12**, 329–335.
Dacorogna, M. M., Gencay, R., Muller, U. A., Olsen, R. B.and Pictet, O. V. (2001) *An Introduction to High-frequency Finance*. San Diego: Academic Press.
Ding, Z. and Granger, C. W. J. (1996) Modeling volatility persistence of speculative returns: a new approach. *J. Econometr.*, **73**, 185–215.
Doornik, J. A. (2001) *Ox: Object Oriented Matrix Programming, 3.0*. London: Timberlake.
Elerian, O., Chib, S. and Shephard, N. (2001) Likelihood inference for discretely observed non-linear diffusions. *Econometrica*, **69**, 959–993.
Engle, R. F. and Lee, G. G. J. (1999) A permanent and transitory component model of stock return volatility. In *Cointegration, Causality, and Forecasting: a Festschrift in Honour of Clive W. J. Granger* (eds R. F. Engle and H. White), ch. 20, pp. 475–497. Oxford: Oxford University Press.
Foster, D. P. and Nelson, D. B. (1996) Continuous record asymptotics for rolling sample variance estimators. *Econometrica*, **64**, 139–174.
Francq, C. and Zakoïan, J.-M. (2000) Covariance matrix estimation for estimators of mixing weak ARMA models. *J. Statist. Planng Inf.*, **83**, 369–394.
Gallant, A. R. and Long, J. R. (1997) Estimating stochastic differential equations efficiently by minimum chi-square. *Biometrika*, **84**, 125–141.
Genon-Catalot, V., Jeantheau, T. and Larédo, C. (2000) Stochastic volatility as hidden Markov models and statistical applications. *Bernoulli*, **6**, 1051–1079.
Ghysels, E., Harvey, A. C. and Renault, E. (1996) Stochastic volatility. In *Statistical Methods in Finance* (eds C. R. Rao and G. S. Maddala), pp. 119–191. Amsterdam: North-Holland.
Gourieroux, C., Monfort, A. and Renault, E. (1993) Indirect inference. *J. Appl. Econometr.*, **8**, S85–S118.
Granger, C. W. J. (1980) Long memory relationships and the aggregation of dynamic models. *J. Econometr.*, **14**, 227–238.
Hamilton, J. (1994) *Time Series Analysis*. Princeton: Princeton University Press.
Harvey, A. C. (1989) *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge: Cambridge University Press.
Hendry, D. F. (1995) *Dynamic Econometrics*. Oxford: Oxford University Press.
Hobson, E. W. (1927) *The Theory of Functions of a Real Variable and the Theory of Fourier's Series*, 3rd edn. Cambridge: Cambridge University Press.
Kim, S., Shephard, N. and Chib, S. (1998) Stochastic volatility: likelihood inference and comparison with ARCH models. *Rev. Econ. Stud.*, **65**, 361–393.
Koopman, S. J., Shephard, N. and Doornik, J. A. (1999) Statistical algorithms for models in state space using Ssf-Pack 2.2. *Econometr. J.*, **2**, 107–166.
Lebesgue, H. (1902) Integrale, longuer, aire. *Ann. Math. Pura Applic.*, **7**, 231–359.
Maheu, J. M. and McCurdy, T. H. (2001) Nonlinear features of realised FX volatility. *Rev. Econ. Statist.*, **83**, to be published.
Meddahi, N. and Renault, E. (2002) Temporal aggregation of volatility models. *J. Econometr.*, to be published.
Munroe, M. E. (1953) *Introduction to Measure and Integration*. Cambridge: Addison-Wesley.
Poterba, J. and Summers, L. (1986) The persistence of volatility and stock market fluctuations. *Am. Econ. Rev.*, **76**, 1124–1141.
Schwert, G. W. (1989) Why does stock market volatility change over time. *J. Finan.*, **44**, 1115–1153.

Shephard, N. (1996) Statistical aspects of ARCH and stochastic volatility. In *Time Series Models in Econometrics, Finance and Other Fields* (eds D. R. Cox, D. V. Hinkley and O. E. Barndorff-Nielsen), pp. 1–67. London: Chapman and Hall.
Silverman, B. W. (1986) *Density Estimation for Statistical and Data Analysis*. London: Chapman and Hall.
Sørensen, M. (2000) Prediction based estimating equations. *Econometr. J.*, **3**, 123–147.
Taylor, S. J. (1986) *Modelling Financial Time Series*. Chichester: Wiley.
———(1994) Modelling stochastic volatility. *Math. Finan.*, **4**, 183–204.
Taylor, S. J. and Xu, X. (1997) The incremental volatility information in one million foreign exchange quotations. *J. Emp. Finan.*, **4**, 317–340.
Whittle, P. (1983) *Prediction and Regulation*, 2nd edn. Oxford: Blackwell.