

Econometric analysis of realised volatility and its use in estimating Lévy based non-Gaussian OU type stochastic volatility models

OLE E. BARNDORFF-NIELSEN

*The Centre for Mathematical Physics and Stochastics (MaPhySto),
University of Aarhus, Ny Munkegade, DK-8000 Aarhus C, Denmark.*
oebn@mi.aau.dk

NEIL SHEPHARD

Nuffield College, Oxford OX1 1NF, UK
neil.shephard@nuf.ox.ac.uk

November 2000 (First circulated 26th October 2000)

Abstract

The availability of intra-day data on the prices of speculative assets means that we can use quadratic variation like measures of activity in financial markets, called realised volatility, to study the stochastic properties of returns. Here we provide a statistical basis for realised volatility and show how it can be used to estimate the parameters of stochastic volatility models. Models covered included those which are based on Lévy driven non-Gaussian OU volatility processes, as well as more traditional type models such as constant elasticity of variance processes or superpositions of such processes.

Keywords: Econometrics; Kalman filter; Lévy process; OU process; Realised volatility; Stochastic volatility; Subordination; Superposition.

1 Introduction

1.1 Stochastic volatility

In stochastic volatility (SV) models the basic Brownian motion model for log-prices is generalised to allow the volatility terms to vary over time. Then the log-price $x^*(t)$ follows the solution to the stochastic differential equation (SDE),

$$dx^*(t) = \{\mu + \beta\sigma^2(t)\} dt + \sigma(t)dw(t), \quad (1)$$

where $\sigma^2(t)$, the *instantaneous volatility*, is going to be assumed to be stationary, latent and stochastically independent of the standard Brownian motion $w(t)$. This model can be thought of as a time series version of a scale mixture of normals model. To see this directly recall the definition of a return over an interval of time of length $\Delta > 0$

$$y_n = x^*(\Delta n) - x^*((n-1)\Delta), \quad n = 1, 2, \dots \quad (2)$$

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

$$\sigma_n^2 = \sigma^{2*}(n\Delta) - \sigma^{2*} \{(n-1)\Delta\}, \quad \text{and} \quad \sigma^{2*}(t) = \int_0^t \sigma^2(u) du.$$

In econometrics $\sigma^{2*}(t)$ is called *integrated volatility*, while 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, Harvey, and Renault (1996), while 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

$$E(\sigma_n^2) = \xi\Delta, \quad \text{Var}(\sigma_n^2) = 2\omega^2 r^{**}(\Delta) \quad \text{and} \quad \text{Cov}\{\sigma_n^2, \sigma_{n+s}^2\} = \omega^2 \diamond r^{**}(\Delta s), \quad (3)$$

where

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

and

$$r^*(t) = \int_0^t r(u) du \quad \text{and} \quad r^{**}(t) = \int_0^t r^*(u) du. \quad (5)$$

That is 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 $x^*(t)$. In particular, for the above SV model the *quadratic variation* is $\sigma^{2*}(t)$, i.e. we have

$$[x^*](t) = \text{p-lim}_{r \rightarrow \infty} \sum \{x^*(t_{i+1}^r) - x^*(t_i^r)\}^2 = \sigma^{2*}(t) \quad (6)$$

for any sequence of partitions $t_0^r = 0 < t_1^r < \dots < t_{m_r}^r = t$ with $\sup_i \{t_{i+1}^r - t_i^r\} \rightarrow 0$ for $r \rightarrow \infty$. This is a powerful result for it does not depend upon the model for instantaneous volatility nor the drift terms in the SDE for log-prices given in (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 (2000a) in foreign exchange markets. Some of the effects of market microstructure on these estimates are studied by Bai, Russell, and Tiao (2000).

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.). As a result we should regard the above results as indicating that we can estimate actual volatility, for example over a day, reasonably accurately by sums of squared returns, say, using five, ten or thirty minute periods. Suppose we have M intra-day observations during each day, then the sum of squared intra-day changes over a day is

$$s_n^2 = \sum_{j=1}^M \left\{ x^* \left((n-1)\Delta + \frac{\Delta j}{M} \right) - x^* \left((n-1)\Delta + \frac{\Delta(j-1)}{M} \right) \right\}^2, \quad (7)$$

which is an estimate of σ_n^2 . It is a consistent estimate as $M \rightarrow \infty$, while it is unbiased when μ and β are zero. In econometrics s_n^2 has recently been labelled *realised volatility*, and we will follow that convention here. In a series of important papers Andersen, Bollerslev, Diebold, and Labys (2000a), Andersen, Bollerslev, Diebold, and Ebens (2000) and Andersen, Bollerslev, Diebold, and Labys (2000b) have empirically studied the properties of s_n^2 in foreign exchange and equity markets (earlier, less formal work on this topic includes Schwert (1989) and Taylor and Xu (1997)). In their econometric analysis they have regarded s_n^2 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/s_n is more or less Gaussian. So far no measure of error has been obtained which indicates the difference between s_n^2 and σ_n^2 . We will 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 prepared to use a model for $\sigma^2(t)$.

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

1.2 Empirical example

To illustrate some of these features we have used the same return data as employed by Andersen, Bollerslev, Diebold, and Labys (2000a), although we have made slightly different adjustments to deal with some missing data. Full details of this are given in Barndorff-Nielsen and Shephard (2001b, Ch. 1). This United States Dollar/ German Deutsche Mark series covers the ten year period from 1st December 1986 until 30th November 1996. It records every five minutes the most recent quote to appear on the Reuters screen. It has been kindly supplied to us by the Olsen group in Zurich and preprocessed by Tim Bollerslev. In the top left graph in Figure 1

we have drawn the correlogram of the squared five minute returns. It shows the well known very strong diurnal effect (the x-axis is drawn in days). This will be discussed in detail in Section 4 but for now will be ignored entirely. The graph on the top right of the Figure shows the correlogram of realised volatility, s_n^2 , computed using $M = 288$. Hence this second graph corresponds to volatility over a single day. The graph starts out at around 0.6, decays very quickly for a number of days and then decays at a slower rate. The graphs on the bottom show a cumulative version of the squared five minute returns drawn on a small scale, while on the right the same cumulative function is drawn over a larger time scale. It is the daily increments of this process which makes up realised volatility.

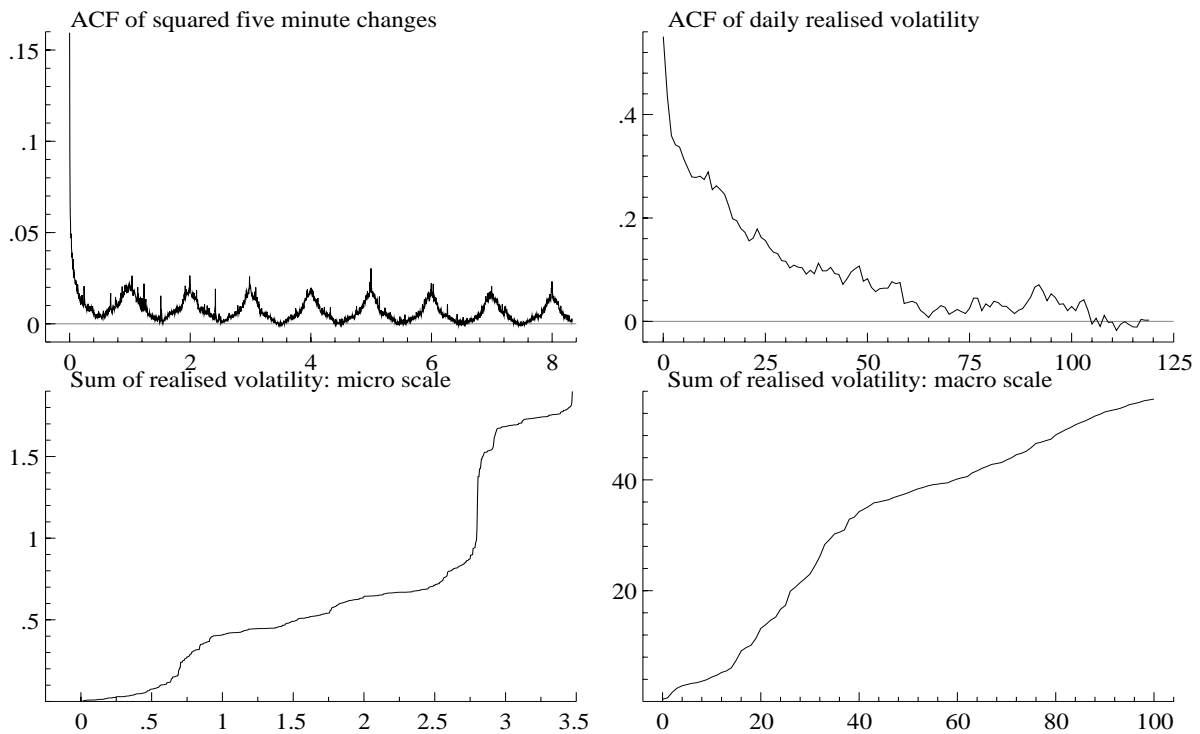


Figure 1: All graphs refer to the Olsen group's five minute changes data. Top left: ACF of five minute returns. Bottom left: cumulative sum of squared 5 minute changes over short interval. Top right: ACF of realised volatility measured over each day. Bottom right: cumulative sum of squared 5 minute changes over long interval.

1.3 OU type processes

A main focus of our paper will be where the volatility is the solution to the SDE

$$d\sigma^2(t) = -\lambda\sigma^2(t)dt + dz(\lambda t), \quad (8)$$

where $z(t)$ is a Lévy process with non-negative increments (which rules out Brownian motion as a candidate model for $z(t)$). The unusual timing $dz(\lambda t)$ is deliberately chosen so that it will turn out that whatever the value of λ the marginal distribution of $\sigma^2(t)$ will be unchanged. Hence we separately parameterise the distribution of the volatility and the dynamic structure. These models, called non-Gaussian Ornstein-Uhlenbeck, or OU type for short, processes, have been developed in this context by Barndorff-Nielsen and Shephard (2001a) and can be generalised to allow for a superposition of such processes (that is a sum of independent OU processes with different decay rates). They are mathematically very tractable allowing us to derive, for example,

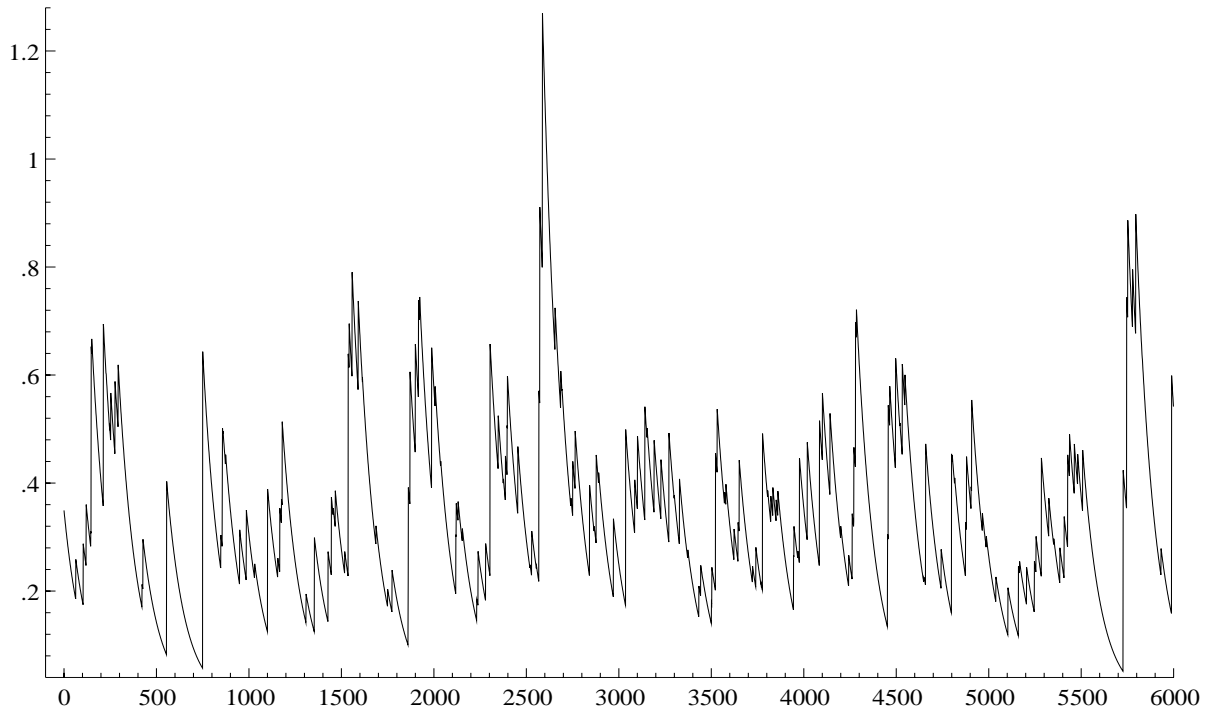


Figure 2: *OU process with $\Gamma(\nu, \alpha)$ marginals. Throughout, $\nu = 3$, $\alpha = 8.5$, $\lambda = 0.01$ and $\Delta = 1$. Plot of $\sigma^2(n\Delta)$ against n .*

analytic option prices for these processes for wide choices of the Lévy process (see Barndorff-Nielsen and Shephard (2001a) and subsequently Nicolato and Venardos (2000) and Tomkins and Hubalek (2000)). A major reason for this is that integrated volatility has a very simple form for these models with (Barndorff-Nielsen and Shephard (2001a))

$$\sigma^{2*}(t) = \lambda^{-1} \{z(\lambda t) - \sigma^2(t) + \sigma^2(0)\}. \quad (9)$$

A simulated example of the paths that the $\sigma^2(t)$ and $z(\lambda t)$ processes follow is given in Figure 2, showing their large upwards jumps and slow downward decays. In this example the process has

been designed to have a marginal law for $\sigma^2(t)$ which is gamma, although in practice many other marginal laws are attractive from an empirical viewpoint. These are developed in Barndorff-Nielsen and Shephard (2001a). Importantly, for these models $r(t) = \exp(-\lambda|t|)$, which imply

$$r^{**}(t) = \lambda^{-2} \left\{ e^{-\lambda|t|} - 1 + \lambda t \right\} \quad \text{and} \quad \diamond r^{**}(\Delta s) = \lambda^{-2} (1 - e^{-\lambda\Delta})^2 e^{-\lambda\Delta(s-1)}, \quad s > 0.$$

An alternative, which is also covered by our analysis, is where volatility follows a constant elasticity of variance process (or a superposition of such processes)

$$d\sigma^2(t) = -\lambda \{ \sigma^2(t) - \xi \} dt + \omega \sigma(t)^\eta db(\lambda t), \quad \eta \in [1, 2],$$

where $b(t)$ is standard Brownian motion uncorrelated with $w(t)$. Of course the special cases of $\eta = 1$ delivers the square root process, while when $\eta = 2$ we have Nelson's GARCH diffusion. These models, favoured heavily by Meddahi and Renault (2000) in this context, also have the property that $r(t) = \exp(-\lambda|t|)$, implying we can also compute the second order properties of $\{\sigma_n^2\}$ for this model. Our results will be based around these second order properties and so results for the OU process will carry over to this process. However, we will not explore this aspect in any detail.

Our analysis will allow us to

- have an understanding of the approximate distribution of $\sqrt{M} (s_n^2 - \sigma_n^2)$ for large M and the exact, generic second order properties of s_n^2 and σ_n^2 .
- use the models for instantaneous volatility to provide *model based* estimates of actual volatility (rather than model free estimates which assume $M \rightarrow \infty$) using the series of realised volatility measurements. These model based estimates can be forecasts, filtered or smoothed estimates of actual volatility depending upon how much information we can use. In particular, these will be based on past, current or historical sequences of realised volatilities.
- estimate the parameters of SV models using simple and rather accurate statistical procedures.

1.4 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 setup where μ and β are zero. Section 3 extends the results to allow for drift, while section 4 studies diurnal effects. Section 5 discusses working with multivariate observations. Section 6 concludes.

2 Realised volatility

2.1 Relating actual to realised volatility

2.1.1 Generic results

Actual volatility, σ_n^2 , plays a crucial role in SV models. It can be estimated using realised volatility, s_n^2 , given in (7). Here we discuss this in the simplest context where $\mu = \beta = 0$, extending to the more complicated case in Section 3. In SV models we can always decompose

$$s_n^2 = \sigma_n^2 + u_n, \quad \text{where} \quad u_n = s_n^2 - \sigma_n^2. \quad (10)$$

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

$$E(s_n^2) = \Delta\xi, \quad \text{Var}(s_n^2) = \text{Var}(u_n) + \text{Var}(\sigma_n^2), \quad \text{Cov}(s_n^2, s_{n+s}^2) = \text{Cov}(\sigma_n^2, \sigma_{n+s}^2).$$

Further, writing $\sigma_{j,n}^2$ as the volatility of the j -th intra-day return, so that,

$$x^* \left((n-1)\Delta + \frac{\Delta j}{M} \right) - x^* \left((n-1)\Delta + \frac{\Delta(j-1)}{M} \right) | \sigma_{j,n}^2 \sim N(0, \sigma_{j,n}^2),$$

we have that

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

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

Now unconditionally,

$$\begin{aligned} \text{Var}(u_n) &= 2ME(\sigma_{1,n}^4) \\ &= 2M \left\{ \text{Var}(\sigma_{1,n}^2) + E(\sigma_{1,n}^2)^2 \right\}, \end{aligned} \quad (11)$$

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

$$E(\sigma_{1,n}^2) = \Delta M^{-1} \xi, \quad \text{Var}(\sigma_{1,n}^2) = 2\omega^2 r^{**} (\Delta M^{-1}). \quad (12)$$

Hence we can compute $\text{Var}(u_n)$ for all SV models when $\mu = \beta = 0$. One of the implications of this results is that

$$\begin{aligned} \text{Cor}(s_n^2, s_{n+s}^2) &= \frac{\text{Cov}(\sigma_n^2, \sigma_{n+s}^2)}{\text{Var}(u_n) + \text{Var}(\sigma_n^2)} \\ &= \frac{\omega^2 \diamond r^{**}(\Delta s)}{2M^{-1} \left\{ 2\omega^2 M^2 r^{**} (\Delta M^{-1}) + (\Delta\xi)^2 \right\} + 2\omega^2 r^{**}(\Delta)}. \end{aligned}$$

Notice the autocorrelation function of $\{y_n^2\}$ is also given by this result, for $s_n^2 = y_n^2$ when $M = 1$. Hence the decay rates in the acf of $\{s_n^2\}$, $\{\sigma_n^2\}$ and $\{y_n^2\}$ are, in general, the same but the degree of correlation varies considerably with the correlation being the highest for $\{\sigma_n^2\}$, followed by $\{s_n^2\}$ and ending with the lowest correlation in $\{y_n^2\}$.

In practice we tend to use realised 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 upon the limit of $t^{-2}r^{**}(t)$ as $t \rightarrow 0$ from above. Now, by Taylor expansion

$$\begin{aligned} r^{**}(t) &= r^{**}(0+) + tr^*(0+) + \frac{1}{2}t^2r(0+) + o(t^2) \\ &= \frac{1}{2}t^2r(0+) + o(t^2). \end{aligned}$$

This means the limit of $t^{-2}r^{**}(t)$ is c where $c = r(0+)$. A consequence of this is that

$$\lim_{M \rightarrow \infty} M^2 \text{Var}(\sigma_{1,n}^2) = c\Delta^2\omega^2 \quad (13)$$

implying, as M goes to infinity,

$$\text{Var}(\sqrt{M}u_n) = \text{Var}\left\{\sqrt{M}(s_n^2 - \sigma_n^2)\right\} \rightarrow 2\Delta^2(c\omega^2 + \xi^2).$$

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

The above limiting result can be strengthened to calculate the asymptotic distribution of this error. In particular we show in the Appendix that

$$\frac{\sqrt{M}(s_n^2 - \sigma_n^2)}{\sqrt{2M \sum_{j=1}^M \sigma_{j,n}^4}} \xrightarrow{\mathcal{L}} N(0, 1), \quad (14)$$

where

$$\sigma_{j,n}^2 = \sigma^{2*} (jM^{-1}\Delta + (n-1)\Delta) - \sigma^{2*} \{(j-1)M^{-1}\Delta + (n-1)\Delta\}, \quad j = 1, 2, \dots, M.$$

Importantly

$$M \sum_{j=1}^M \sigma_{j,n}^4$$

has a stochastic limit, and so $\sqrt{M}(s_n^2 - \sigma_n^2)$ has a mixed normal asymptotic distribution.

2.2 Lévy based volatility models

2.2.1 Recalling some properties

In this paper we will be primarily interested in estimating non-Gaussian OU based SV models of the type (8). The OU structure allows us to explicitly compute the mean, variance and

autocorrelation function of the actual volatility. In particular

$$E(\sigma_n^2) = \Delta\xi, \quad \text{Var}(\sigma_n^2) = \frac{2\omega^2}{\lambda^2} \left\{ e^{-\lambda\Delta} - 1 + \lambda\Delta \right\}$$

and

$$\text{Cor}\{\sigma_n^2, \sigma_{n+s}^2\} = de^{-\lambda\Delta(s-1)}, \quad s = 1, 2, \dots, \quad (15)$$

where

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

Importantly this implies actual volatility has the autocorrelation function of an autoregressive

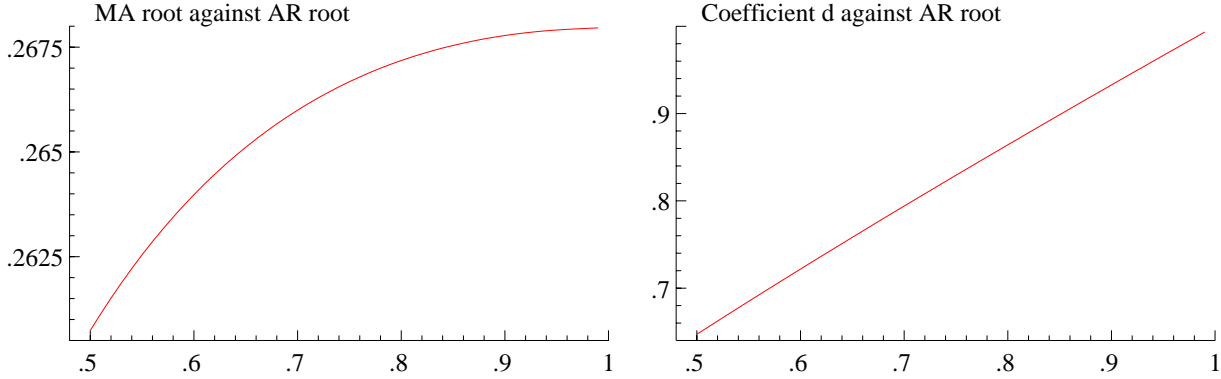


Figure 3: *Left graph shows plot of moving average root against autoregressive root $e^{-\Delta\lambda}$ for ARMA(1,1) representation. Right graph shows d in expression for $\text{Cor}\{\sigma_n^2, \sigma_{n+s}^2\}$ against autoregressive root $e^{-\Delta\lambda}$.*

moving average (ARMA) model of order (1,1). Its autoregressive root is $e^{-\lambda\Delta}$ (which will be typically close to one unless Δ is very large), while the moving average root is also determined by $e^{-\lambda\Delta}$ but has to be found numerically. A graph of the moving average root against $e^{-\lambda\Delta}$ is given in the left hand side of Figure 3 and shows that for a wide range of the autoregressive root the moving average root is around 0.265. Likewise Figure 3 shows a plot of d against $e^{-\lambda\Delta}$ and indicates a rapid decline in this coefficient as the autoregressive root falls. In particular, in financial econometrics the literature suggests volatility is quite persistent, which would imply d should be close to one. Thus if t is recorded in days and Δ is set to one day, then empirically reasonable values of λ will imply d should be close to one.

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

$$\text{Cor}\{y_n^2, y_{n+s}^2\} = c'e^{-\lambda\Delta(s-1)}, \quad (16)$$

where

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

Importantly this means $\{y_n^2\}$ also has a linear ARMA(1,1) representation. Further, it has the same autocorrelation function as the familiar generalised autoregressive conditional heteroskedastic (GARCH) model used extensively in econometrics (see, for example, Bollerslev, Engle, and Nelson (1994)). Finally, the autoregressive root of the ARMA representation is the same for $\{y_n^2\}$ as for $\{\sigma_n^2\}$, 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 *much* less clear than if we had observed the correlograms of the latent $\{\sigma_n^2\}$. This can be most easily seen by noting that for small λ ,

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

which is much smaller than d which is approximately $1 - \lambda\Delta$. For example if the $\xi = \omega^2$, then c' will be approximately 0.2 for daily data.

2.2.2 Implied error in realised volatility

As σ_n^2 has an ARMA(1,1) representation, we can think of (10) as a very compact linear state space representation (see, for example, the textbook expositions in Harvey (1989, Ch. 3) and Hamilton (1994, Ch. 13)). This is briefly outlined in the Appendix. In order to fully characterise it we just have to calculate the exact unconditional variance of $\{u_n\}$. We gave a generic expression for it in (11) and (12), now we become more specialised. Having carried that out we could use a Kalman filter to unbiasedly and efficiently (in a linear sense) estimate and predict future actual volatilities, refining the realised volatility estimate. As a bi-product the Kalman filter also provides a quasi-likelihood function for the realised volatility which can be used to effectively estimate the parameters of the model. We will briefly study the properties of these estimators in a moment.

Now unconditionally

$$\begin{aligned} \text{Var}(u_n) &= 2ME(\sigma_{1,n}^4) \\ &= 2M \left\{ \text{Var}(\sigma_{1,n}^2) + E(\sigma_{1,n}^2)^2 \right\} \\ &= 2M \left\{ 2\omega^2\lambda^{-2} \left(e^{-\lambda\Delta/M} - 1 + \lambda\Delta M^{-1} \right) + (\Delta M^{-1})^2 \xi^2 \right\}. \end{aligned} \quad (17)$$

Of course, for large M

$$2 \left(e^{-\lambda\Delta/M} - 1 + \lambda\Delta M^{-1} \right) \simeq (\lambda\Delta M^{-1})^2,$$

implying, as M goes to infinity,

$$\text{Var}(\sqrt{M}u_n) \rightarrow 2\Delta^2(\omega^2 + \xi^2).$$

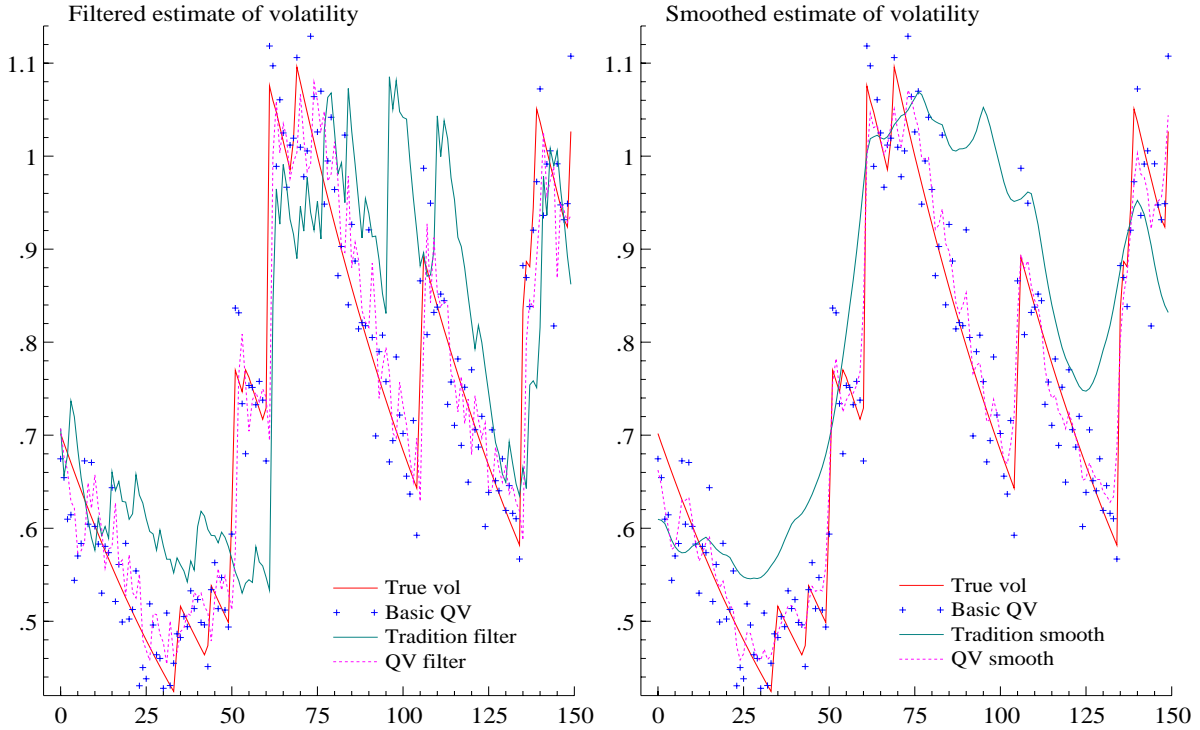


Figure 4: *Payoff of using intra-day information in estimating volatility in simulated SV models. OU process $\sigma^2(t) \sim \Gamma(2, 4)$ distribution and $\lambda = -\log(0.98)$, with $\Delta = 1$ representing a day. Throughout $M = 128$. Left shows filtered quantities, right smoothed ones. Throughout, basic QV denotes s_n^2 , tradition smooth is the smoothed estimate based on squared daily returns ($M = 1$) and QV smooth denotes the model based smoothed estimate based on all the intra-day information.*

This is the same result as in (14) with $r(0+) = c = 1$.

To illustrate this method we have constructed a simulation experiment. We simulated an OU process for $\{\sigma_n^2\}$ with a marginal $\Gamma(2, 4)$ distribution and $\lambda = -\log(0.98)$, with $\Delta = 1$ representing a day. We assumed that $M = 128$ and that we can observe the corresponding intra-day returns at that frequency as well as the aggregated return at the daily level. Figure 4 shows the plot of true daily volatility, drawn as $\{\sqrt{\sigma_n^2}\}$, over time, together with the corresponding realised volatility for each day $\{\sqrt{s_n^2}\}$. The realised volatility is a pretty good estimate of the true volatility, but is quite noisy. However, it is far more precise than the corresponding best linear estimator of the volatility based on the above model (10) using squared daily returns. On the left hand graph we have shown the one-step ahead prediction of the volatility σ_n^2 , using y_1^2, \dots, y_{n-1}^2 . This prediction is quite close to the one which would result from a discrete time GARCH model and so should be regarded as being reasonable by the academic literature. However, it does not exploit the intra-day data and so is much less precise than the realised volatility estimate. This carries over to the smoothed estimate of the volatility, which is given

on the right hand side of the graph. This uses all the information in the daily returns sample, y_1^2, \dots, y_{n-1}^2 , to estimate σ_n^2 . This is more precise, but is still inferior to the realised volatility estimate.

M	Mean	Standard error	0.1 quantile	0.9 quantile
λ , truth is 0.03046				
1	0.206	0.612	0.0157	0.561
10	0.0447	0.233	0.0223	0.0425
100	0.0315	0.00507	0.0255	0.0384
ν , truth is 2				
1	1.93	1.04	0.800	3.20
10	2.06	0.449	1.53	2.65
100	2.05	0.304	1.67	2.45
a , truth is 4				
1	4.00	2.11	1.60	6.80
10	4.14	0.852	3.16	5.17
100	4.13	0.591	3.38	4.94

Table 1: *Monte Carlo study of the quasi-likelihood estimator of SV model with OU volatility. Volatility model has $\sigma^2(t) \sim \Gamma(\nu, a)$ with 500 daily observations. The true values of ν and a are 2 and 4 respectively. M denotes the number of intra-day observations used. 1,000 replications are used in the study. Code is available to carry out these calculations in `ssf.ox`.*

Figure 4 also gives the filtered (that is the estimate of σ_n^2 , using s_1^2, \dots, s_{n-1}^2) and smoothed (that is the estimate of σ_n^2 , using s_1^2, \dots, s_T^2) estimates based on the model based use of intra-day data. The parameters in this analysis are set at the true values. The filtered estimate is less spread out than the realised version and is in fact more precise. This is remarkable as it is based on a predictive information, while realised volatility uses current information. The smoothed estimate is far more precise than the realised volatility estimate. Indeed the error in the estimates are very minor.

Table 1 shows the result of a small simulation experiment which investigates the effectiveness of the quasi-likelihood estimation methods based on the time series of realised 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, however 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, Jeantheau, and Larédo (2000) who show if volatility is strong mixing then squared returns are strong mixing). This means we can immediately apply the asymptotic theory results of Francq and Zakoian (2000) in this context so long as $\sigma^2(t)$ is strong mixing. Further the estimation takes only around 5 seconds on a modern notebook computer.

The setup of the simulation study uses 500 daily observations where the volatility is an OU process with a gamma marginal distribution. The Table varies the value of M , the number of intra-day observations available. When $M = 1$ this corresponds to using the classical approach of squared daily returns. When M is higher we are using intra-day data. The results suggest that the intra-day data allows 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 be jointly Gaussian. The results are quite encouraging for they are based on only two years of data but suggests we can construct quite precise estimates of these models with this.

The above method can be contrasted with the large literature on the estimation of (partially observed) stochastic differential equations based on discrete data. A very incomplete list of references include Gouriéroux, Monfort, and Renault (1993), Gallant and Long (1997), Elerian, Chib, and Shephard (2001) and Sørensen (2000). The closest paper to ours is a recent one by Bollerslev and Zhou (2000) who use a method of moment approach based on an assuming that the actual volatility process $\{\sigma_n^2\}$ is observed via the quadratic variation estimator.

2.3 Superpositions of OU processes

The OU model (8) for the instantaneous volatility is often too simple in practice to fit the types of dependence structures we observe in financial economics. This can be seen in the top right piece of Figure 1 which displays the autocorrelation function of realised volatility for the Olsen group's five minute data. This graph shows a relatively quick initial decline in the acf, followed by a slower decay. The autocorrelation at lag one is around 0.6, which is a long way from 1. This can only be reconciled with an OU based SV model if the decay rate of the OU process is very fast. However, this would not allow us to fit the rest of the correlogram which displays considerable memory. This single observation is sufficient to dismiss the OU (and constant elasticity of variance) model.

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 processes. As the processes do not need to be identically distributed, this offers a great deal of flexibility while still being mathematically tractable (for example, see Nicolato and Venardos (2000) for analytic derivative pricing based on these models). Superpositions of OU processes also have potential for modelling long-range dependence in volatility. This is discussed briefly in Barndorff-Nielsen and Shephard (2001a) and at more depth by Barndorff-Nielsen (2000) who formalises the use of superpositions as a way of modelling long-range dependence. This follows earlier related work by Granger (1980), Cox (1991), Ding and Granger (1996), Engle and Lee (1999) and Comte and Renault (1998).

Consider volatility based up of the sum of J independent OU processes

$$\sigma^2(t) = \sum_{j=1}^J \tau^j(t), \quad \text{where} \quad d\tau^j(t) = -\lambda_j \tau^j(t) dt + dz_j(\lambda_j t),$$

where the $z_j(t)$ are independent Lévy processes with non-negative increments. It is not necessary for the Lévy processes to have the same marginal distributions at time one. In practice in this paper we have used the following structure to parameterise the model.

We assume the non-negative weights $\{w_j\}$ are such that

$$E(\tau^j(t)) = w_j \xi \quad \text{Var}(\tau^j(t)) = w_j \omega^2 \quad \text{where} \quad \sum_{j=1}^J w_j = 1,$$

implying

$$E(\sigma^2(t)) = \xi, \quad \text{Var}(\sigma^2(t)) = \omega^2,$$

and

$$\begin{aligned} \text{Cov}(\sigma^2(t), \sigma^2(t+s)) &= \sum_{j=1}^J \text{Cov}(\tau^j(t), \tau^j(t+s)) \\ &= \omega^2 \sum_{j=1}^J w_j \exp(-\lambda_j |s|). \end{aligned}$$

Hence the acf of instantaneous volatility can have components which are a mix of quickly and slowly decaying components. A simple parametric example of this construction is where $\tau^j(t) \sim \Gamma(w_j \nu, \alpha)$ -OU, implying

$$\sigma^2(t) \sim \Gamma(\nu, \alpha)\text{-OU}_J,$$

a superposition of J independent OU processes with a marginal distribution which is gamma. A similar type of process can be constructed with inverse Gaussian marginal distributions.

The linearity of the superposition of OU processes means that actual volatility has the form

$$\sigma_n^2 = \sum_{j=1}^J \tau_n^j,$$

where

$$\tau_n^j = \tau^{j*}(n\Delta) - \tau^{j*}\{(n-1)\Delta\}, \quad \text{and} \quad \tau^{j*}(t) = \int_0^t \tau^j(u) du.$$

The key feature is that each τ_n^j has an ARMA(1,1) representation of the type discussed above. 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 acf of σ_n^2 without any new work.

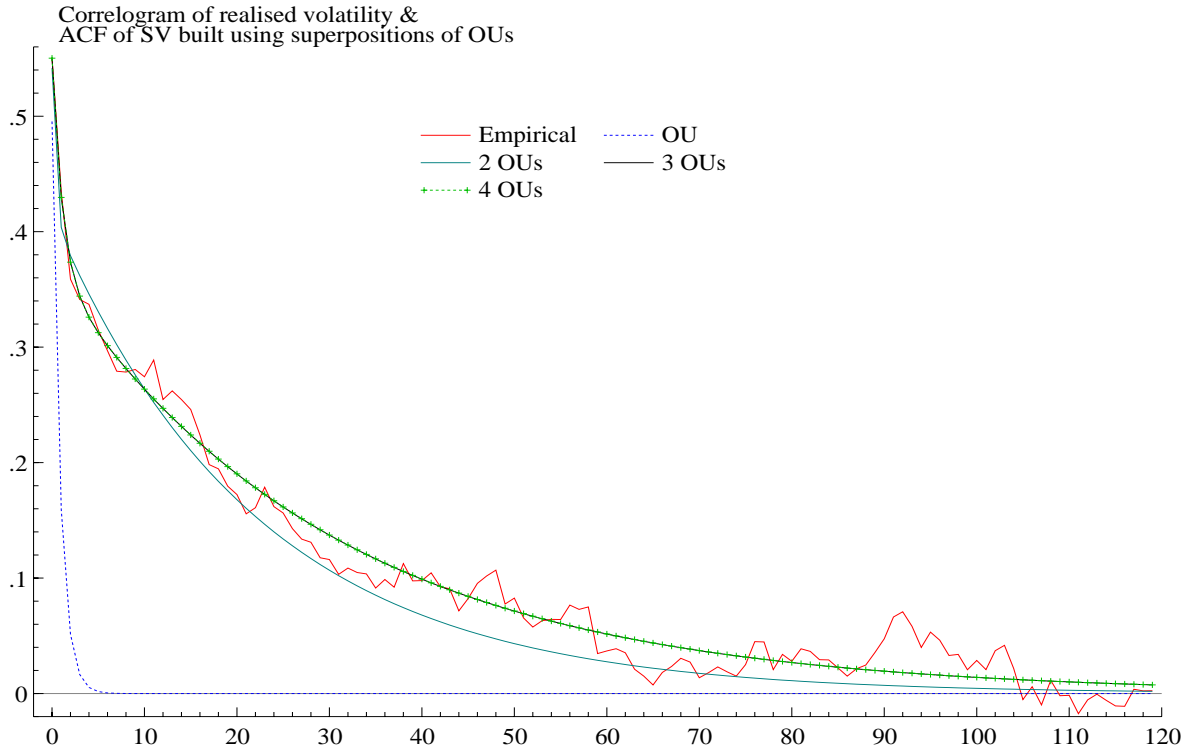


Figure 5: *Autocorrelations for realised volatility. Graphed are the correlogram for realised volatility and the fitted autocorrelation functions for various SV models based on the superposition of up to four OU processes.*

Computationally it is helpful to realise that the sum of uncorrelated ARMA(1,1) processes can be fed into a linear state space representation when combined with (10). The only new issue is computing

$$\text{Var}(u_t) = 2M \left\{ \text{Var}(\sigma_{1,n}^2) + \text{E}(\sigma_{1,n}^2)^2 \right\}.$$

Clearly

$$\text{E}(\sigma_{1,n}^2) = \xi \frac{\Delta}{M},$$

while

$$\begin{aligned} \text{Var}(\sigma_{1,n}^2) &= \sum_{j=1}^J \text{Var}(\tau_{1,n}^j) = 2\omega^2 \sum_{j=1}^J w_j r_j^{**} (\Delta M^{-1}) \\ &= 2\omega^2 \sum_{j=1}^J \frac{w_j}{\lambda_j^2} \left\{ e^{-\lambda_j \Delta M^{-1}} - 1 + \lambda_j \Delta M^{-1} \right\}. \end{aligned}$$

This implies we can also estimate the parameters of this model via the Kalman filter which moreover gives us filtered, smoothed and forecasted estimates of actual volatility. No new issues arise. Further it allows us to compute the acf of realised volatility straightforwardly as the

autocovariance function is just the autocovariance function of actual volatility plus the white noise measurement error.

To illustrate these points we have drawn in Figure 6 the fitted acf (drawn in lags of days) of realised volatility for a variety of values of M , taking the parameter values from Table 3 which fits this model to high frequency exchange rate returns. This data will be discussed at some length in the next subsection. For now it is sufficient to note that as M decreases the correlation in realised volatility tends to fall. This is most marked at the low lag values.

2.4 Empirical illustration

To illustrate some of these results we have fitted a set of superposition based OU type SV models to the realised volatility series discussed in the introduction to this paper. There we computed realised volatility using 5 minute changes to the exchange rate. Here we use the quasi-likelihood method to estimate the parameters of the superposition — ξ , ω^2 , $\{\lambda_j\}$ and $\{w_j\}$. Our empirical work will look at fitting a series of these models increasing J until the fit of the model does not improve very much. The results, given in Table 2, are striking. The third OU process has a value of w_3 which is between 0.96 and 0.94, while the damping factor is very high with λ_3 being over 200. This means the vast majority of the instantaneous volatility has very little predictability even over a five minute period. This means that σ_n^2 will be far less jagged than $\sigma^2(n\Delta)$, which is the reason why ω^2 rose so much as J went from two to three. This is an important result. (We note that our SV models have continuous sample paths and so are, in principle, different from SV models with added jumps which are the usual way of modelling this type of occurrence in SDEs.) The other two OU processes have similar weights, but the first process has a daily damping factor of around 0.97 which means that around four percent of a shock to this process is left after 100 trading days. The second factor lasts only a couple of days.

J	ξ	ω^2	λ_1	λ_2	λ_3	w_1	w_2	Quasi-L	BP
3	0.5321	4.712	0.03263	0.8865	289.6	0.01870	0.02081	-1153.1	17.331
2	0.53206	0.37044	0.045212	2.8396	—	0.27294	—	-1160.2	33.848
1	0.53206	0.32107	1.1344	—	—	—	—	-1325.3	700.28

Table 2: *Fit of the superposition of J OU type volatility process for a SV model. The data is realised volatility computed using $M = 288$, that is five minute returns. We do not record w_J as this is 1 minus the sum of the other weights. Estimation method: quasi-likelihood using output from a Kalman filter. BP denotes Box-Pierce statistic, based on 20 lags, which is a test of serial dependence in the scaled residuals.*

Figure 5 shows the corresponding autocorrelation function for the realised volatility together with the corresponding empirical correlogram. We see from this figure that the simple OU

process is entirely unable to fit the data, as it starts at around 0.6 and then almost instantly decays to zero. A superposition of two processes is much better, but damps too slowly initially and then too quickly. In particular it is poor at picking up the longer-range dependence in the data. The superposition of three and four processes give very similar fits, indeed in the graph they are hardly distinguishable, suggesting a model with three OU processes is sufficiently flexible.

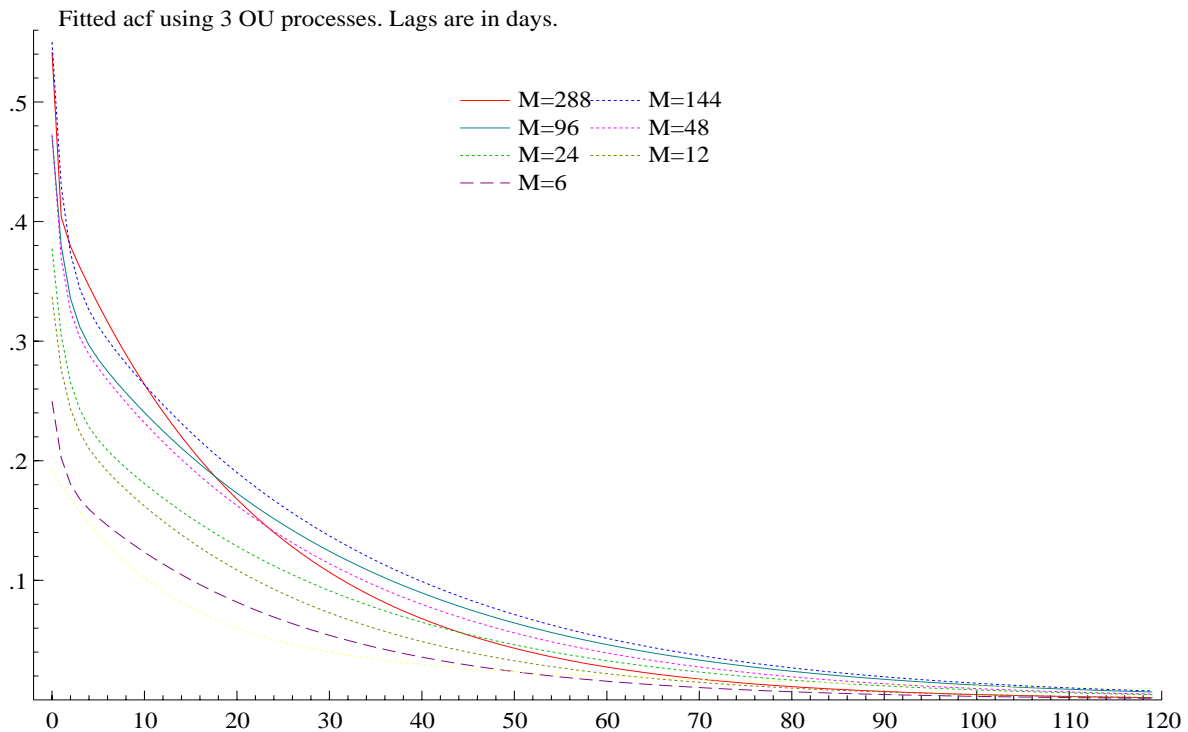


Figure 6: *Fitted acf for realised volatility using a superposition of 3 OU processes. Different curves are given for different values of M , reestimating the model's parameters at each value of M . The fitted parameters are given in the next subsection.*

To assess the influence of M on the estimates of the parameters of the model we have reestimated the model using a variety of values of M . The results are reported in Table 3. There are a number of interesting features of these results. The estimator of parameter ξ falls very slightly as M increases, while ω^2 does change quite considerably. The reason for this is that as w_3 changes, which corresponds to a very short memory OU process, so ω^2 must vary considerably in order for the unconditional variance of the return over a day to remain constant. One of the most encouraging features of the results is that the estimators of $\{\lambda_j\}$ do not vary very much with M while there is not much sign of any residual serial dependence after the fitting of the model. Finally, one of the features of Figure 6 is that as M goes from 288 to 144 it crosses the other curves. This is not the expected result if the SV model fits perfectly in continuous

M	ξ	ω^2	λ_1	λ_2	λ_3	w_1	w_2	Quasi-L	BP
288	0.5321	4.712	0.03263	0.886	289.6	0.01870	0.02081	-1153.1	17.3
144	0.5112	5.175	0.03292	0.890	288.8	0.01693	0.01536	-1371.7	16.9
96	0.4936	3.536	0.03547	1.046	291.5	0.02450	0.02649	-1375.6	17.9
48	0.4727	4.398	0.03413	0.740	289.9	0.01785	0.01670	-1704.1	21.9
24	0.4630	2.479	0.03997	0.775	293.7	0.03384	0.02652	-1885.2	23.7
12	0.4714	1.8146	0.04131	1.000	294.29	0.05180	0.04197	-2426.5	25.5
6	0.4813	0.4733	0.01172	0.0812	294.30	0.04498	0.19678	-2633.4	30.4

Table 3: *Fit of the superposition of 3 OU type volatility process for a SV model with different values of M. Estimation method: quasi-likelihood using output from a Kalman filter built on realised volatility. BP denotes Box-Pierce statistic, based on 20 lags, which is a test of serial dependence in the scaled residuals.*

time and so we can regard this as a sign of misspecification at the level of 5 minute returns.

3 Drift

Suppose we generalise the standard model to allow for a drift effect

$$dx^*(t) = \mu dt + \sigma(t)dw(t). \quad (18)$$

We saw in equation (6) that the addition of the drift into the model does not change the quadratic variation at all and so realised volatility is still a consistent estimator of actual volatility. However, realised volatility will now be biased.

To study the dynamic properties of realised volatility we note that

$$\begin{pmatrix} y_n \\ s_n^2 \end{pmatrix} = \begin{pmatrix} \mu\Delta + v_n \\ \mu^2\Delta^2M^{-1} + \sigma_n^2 + u_n \end{pmatrix}, \quad (19)$$

where $\{u_n, v_n\}$ is a zero mean, vector white noise process. Of course the $\mu^2\Delta^2M^{-1}$ term, which we call an offset, is likely to be very small in practice if M is large for μ is typically very small.

In particular, using the same notation as before

$$\begin{pmatrix} v_n \\ u_n \end{pmatrix} \stackrel{\mathcal{L}}{=} \left\{ \begin{array}{c} \sum_{j=1}^M \varepsilon_{j,n} \sigma_{j,n} \\ 2\mu\Delta M^{-1} \sum_{j=1}^M \varepsilon_{j,n} \sigma_{j,n} + \sum_{j=1}^M \sigma_{j,n}^2 (\varepsilon_{j,n}^2 - 1) \end{array} \right\}$$

The implication is that

$$\begin{aligned} \text{Var}(v_n) &= ME(\sigma_{1,n}^2) = \Delta^{-1}\xi, \\ \text{Cov}(v_n, u_n) &= 2\mu\Delta E(\sigma_{1,n}^2) = 2M^{-1}\mu\xi, \end{aligned}$$

and

$$\begin{aligned} \text{Var}(u_n) &= 4\mu^2\Delta^2M^{-1}E(\sigma_{1,n}^2) + 2ME(\sigma_{1,n}^4) \\ &= 4\mu^2\Delta M^{-2}\xi + 2M \{ \Delta^2M^{-2}\xi^2 + 2\omega^2r^{**}(\Delta M^{-1}) \}. \end{aligned}$$

For short-memory volatility models we have seen in (13) that $M^2\mathbf{E}(\sigma_{1,n}^4)$ converges to a constant as M goes to infinity. Hence

$$\text{Cov}\left(\begin{array}{c} v_n \\ \sqrt{M}u_n \end{array}\right) \rightarrow \left(\begin{array}{cc} \Delta^{-1}\xi & 0 \\ 0 & 2\Delta^2(c\omega^2 + \xi^2) \end{array}\right), \quad \text{as } M \rightarrow \infty.$$

Of course v_n does not converge to being normally distributed as M goes to infinity, although $\sqrt{M}u_n$ will again converge to a mixed normal. This implies that for large M the return y_n has barely any information in it about the volatility process. Indeed it seems sensible to entirely ignore the effect of the drift term, only correcting for it in the offset term in the representation of s_n^2 in (19). However, even this correction seems relatively unimportant.

4 Diurnal affects and actual volatility

An important aspect of the realised volatility series is that it is not very sensitive to the substantial and complicated intra-day diurnal pattern in volatility found in many empirical studies (e.g. Andersen and Bollerslev (1997) and Andersen and Bollerslev (1998b)) as well as being clear from the top left of Figure 1. To understand this it is helpful to think of the instantaneous volatility as the sum of a deterministic diurnal component, $\sigma_\psi^2\{\text{mod}(t, \Delta)\}$ where $t = \Delta$ represents a year, plus a stochastic process, $\sigma_\lambda^2(t)$, then we have

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

As a result

$$\sigma_n^2 = c + \sigma_{n,\lambda}^2, \quad \text{where } c = \int_0^\Delta \sigma_\psi^2\{\text{mod}(u, \Delta)\} du$$

and

$$\sigma_{n,\lambda}^2 = \sigma_\lambda^{2*}(n) - \sigma_\lambda^{2*}\{(n-1)\}, \quad \text{and } \sigma_\lambda^{2*}(t) = \int_0^t \sigma_\lambda^2(u) du.$$

Hence in this structure the dynamics of realised volatility is unaffected by the presence of 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 will neglect this deficiency.

5 Extensions

5.1 Basics

There are at least three important extensions to these results. Each of these results are more intricate than the ones we report here and so we have discussed them at length in Barndorff-Nielsen and Shephard (2001b). Here we just outline the problems. The first deals with the case where β , in (1), is not zero. This does not raise too many new problems.

The second extension allows for the introduction of a leverage term into the model. This can be carried out in a number of ways. In Barndorff-Nielsen and Shephard (2001a) we parameterise the effect as

$$dx^*(t) = \{\mu + \beta\sigma^2(t)\} dt + \sigma(t)dw(t) + \rho dz(\lambda t),$$

where we assume the volatility process is of OU type (8). Here the Lévy process which drives the volatility also appears in the price equation. This is quite a significant change to the process for now the log-prices do not have continuous sample paths, but will jump at the same time that the volatility jumps. A major advantage of this approach is that it is still possible to produce analytic option pricing formulae based on these types of models — see Nicolato and Venardos (2000). The corresponding quadratic variation for this process is

$$[x^*](t) = \sigma^{2*}(t) + \rho^2[z](\lambda t).$$

Hence realised volatility will not consistently estimate integrated volatility.

5.2 Multivariate case

The final extension we could think of is to the multivariate case. A simple multivariate structure for a log-price vector can be generated off a $N \times 1$ multivariate SV model. In particular a simple extension of the univariate setup is to write

$$dx^*(t) = \{\mu + \Sigma(t)\beta\} dt + \Sigma(t)^{1/2}dw(t),$$

where $w(t)$ is a vector of independent standard Brownian motions. Then the return vector

$$y_n = x^*(n\Delta) - x^*((n-1)\Delta),$$

is a multivariate mixture of normals. In particular

$$y_n | \Sigma_n \sim N(\mu\Delta + \Sigma_n\beta, \Sigma_n),$$

where

$$\Sigma_n = \Sigma^*(n\Delta) - \Sigma^*((n-1)\Delta) \quad \text{and} \quad \Sigma^*(t) = \int_0^t \Sigma(t)dt.$$

We call $\Sigma^*(t)$ *integrated covolatility* and Σ_n *actual covolatility*. For the above SV model the *quadratic covariation* is $\Sigma^*(t)$, i.e. we have

$$[x^*](t) = \text{p-lim}_{r \rightarrow \infty} \sum \{x^*(t_{i+1}^r) - x^*(t_i^r)\} \{x^*(t_{i+1}^r) - x^*(t_i^r)\}' = \Sigma^*(t) \quad (20)$$

for any sequence of partitions $t_0^r = 0 < t_1^r < \dots < t_{m_r}^r = t$ with $\sup_i \{t_{i+1}^r - t_i^r\} \rightarrow 0$ for $r \rightarrow \infty$. Again this is a robust measure as it produces the integrated covolatility even if μ and β are

non-zero. However, it is an entirely asymptotic concept and so is not directly applicable in practice.

The actual covolatility can be estimated using intra-day observations

$$S_n = \sum_{j=1}^M y_{i,n} y'_{i,n}, \quad \text{where} \quad y_{i,n} = x^* \left((n-1) \Delta + \frac{\Delta i}{M} \right) - x^* \left((n-1) \Delta + \frac{\Delta (i-1)}{M} \right).$$

This is consistent as M goes to infinity (see Barndorff-Nielsen and Shephard (2001a)) and is unbiased if μ and β is zero. We call S_n *realised covolatility*. Some of the empirical properties of $\{S_n\}$ are studied in Andersen, Bollerslev, Diebold, and Labys (2000a) in the bivariate case.

When μ and β are zero then

$$S_n = \Sigma_n + U_n, \quad \text{where} \quad U_n = S_n - \Sigma_n$$

where U_n is a zero mean, white noise error which is uncorrelated with all of the elements of actual covolatility. Its properties are studied at length in Barndorff-Nielsen and Shephard (2001b). Although the properties of the elements of S_n are intricate, they raise no new issues compared to the univariate case.

6 Conclusion

In this paper we have studied the statistical properties of realised volatility in the context of SV models. Our results are entirely general, providing both a central limit theory approximation to their distribution as well as an exact second order analysis. These results can be used within a linear state space representation, 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.

The results extend to allow us to deal with drift, skewness, leverage and the multivariate cases. 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.

7 Acknowledgments

Ole E. Barndorff-Nielsen's work is supported by CAF (www.caf.dk), which is funded by the Danish Social Science Research Council, and by MaPhySto (www.maphysto.dk), which is funded by the Danish National Research Foundation. Neil Shephard's research is supported by the UK's ESRC through the grant "Econometrics of trade-by-trade price dynamics," which is coded R00023839. All the calculations made in this paper are based on software written by the authors using the Ox language of Doornik (1998) combined with the state space software library

described by Koopman, Shephard, and Doornik (1999). It is contained within the file `ssf.ox`. We owe a large debt to Michel M. Dacorogna at the Olsen Group for allowing us to use Olsen's high frequency exchange rate data in our study and Tim Bollerslev for supplying us with a semi-cleaned version of this data. The comments of Frank Gerhard help the exposition. Finally, we thank Nour Meddahi and Enrique Sentana who pointed out to us the work of Francq and Zakoian on ARMA models with mixing white noise error terms.

8 Appendix

In the appendix we derive two results. First we discuss the asymptotic theory of realised volatility as $M \rightarrow \infty$. Second we present the details of the linear state space form for realised volatility.

8.1 Asymptotic distribution of realised volatility

For this subsection we introduce the notation

$$\tau_j = \sigma^{2*} (jM^{-1}\Delta) - \sigma^{2*} \{(j-1)M^{-1}\Delta\}, \quad j = 1, 2, \dots, M,$$

and then note the error term for realised volatility for day one is

$$\begin{aligned} u_1 &= s_1^2 - \sigma^{2*}(\Delta) \\ &= \sum_{j=1}^M [x^*(jM^{-1}\Delta) - x^*\{(j-1)M^{-1}\Delta\}]^2 - \sigma^{2*}(\Delta). \end{aligned}$$

Then the conditional cumulant function of u_1 is

$$\log \phi(\zeta; u_1 | \tau_1, \dots, \tau_M) = -\frac{1}{2} \sum_{j=1}^M \{\log(1 - 2i\zeta\tau_j) - i\zeta\tau_j\}.$$

For $|\zeta| \max_{1 \leq j \leq M} \tau_j < 1$ we find, by Taylor's formula,

$$\begin{aligned} \log \phi(\zeta; u_1 | \tau_1, \dots, \tau_M) &= -\frac{1}{2} \sum_{j=1}^M \{\log(1 - 2i\zeta\tau_j) - i\zeta\tau_j\} \\ &= -\zeta^2 M \sum_{j=1}^M \tau_j^2 \int_0^1 \frac{1-u}{1-i\zeta\tau_j u} du. \end{aligned}$$

Since $E(\tau_j) = \Delta\xi M^{-1}$ and $\text{Var}(\tau_j) = 2\omega^2 r^{**}(\Delta M^{-1}) \sim r(0+)\omega^2 \Delta^2 M^{-2}$ we have, under mild regularity conditions, that

$$|\sqrt{M}| \max_{1 \leq j \leq M} \tau_j \xrightarrow{p} 0$$

for $M \rightarrow \infty$. It follows that the conditional characteristic function of $\sqrt{M}u_1$ satisfies

$$\lim_{M \rightarrow \infty} \log \phi(\zeta; \sqrt{M}u_1 | \tau_1, \dots, \tau_M) = -\zeta^2 \lim_{M \rightarrow \infty} M \sum_{j=1}^M \tau_j^2.$$

Hence

$$\frac{s_1^2 - \sigma_1^2}{\sqrt{2 \sum_{j=1}^M \tau_j^2}} \xrightarrow{\mathcal{L}} N(0, 1).$$

8.2 State space representation of s_n^2 for single OU process

Recall actual volatility σ_n^2 has an ARMA(1,1) representation while

$$s_n^2 = \sigma_n^2 + u_n, \quad \text{where} \quad u_n = s_n^2 - \sigma_n^2,$$

where u_n is a zero mean white noise process with variance (17) which we write as σ_u^2 . This noise process is uncorrelated with actual volatility. Hence the process has a linear state space representation

$$\begin{aligned} s_n^2 &= \Delta\xi + x_n\alpha_n + u_n, & u_n &\sim WN(0, \sigma_u^2) \\ \alpha_{n+1} &= T_n\alpha_n + G_nv_n, & v_n &\sim WN(0, 1). \end{aligned}$$

Here $WN(.,.)$ denotes white noise errors. The error terms u_n and v_n are uncorrelated while

$$x_n = (1 \quad 0), \quad T_n = \begin{pmatrix} \phi & 1 \\ 0 & 0 \end{pmatrix}, \quad G_n = \begin{pmatrix} \sigma_\sigma \\ \sigma_\sigma\theta \end{pmatrix},$$

where ϕ , θ and σ_σ represent the autoregressive root, the moving average root and the variance of the innovation to this process. Software, based on the Kalman filter, for carrying out best linear estimates of $\{\sigma_n^2\}$ using the realised volatility is available in Koopman, Shephard, and Doornik (1999).

References

- Andersen, T. G. and T. Bollerslev (1997). Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance* 4, 115–158.
- Andersen, T. G. and T. Bollerslev (1998a). Answering the skeptics: yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39, 885–905.
- Andersen, T. G. and T. Bollerslev (1998b). Deutsche mark-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies. *Journal of Finance* 53, 219–265.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and H. Ebens (2000). The distribution of stock return volatility. Unpublished paper: Economics Department, Duke University.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys (2000a). The distribution of exchange rate volatility. *Journal of the American Statistical Association* 95. Forthcoming.

- Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys (2000b). Exchange rate returns standardised by realised volatility are (nearly) Gaussian. Unpublished paper: Economics Department, Duke University.
- Bai, X., J. R. Russell, and G. C. Tiao (2000). Beyond Merton's utopia: effects of non-normality and dependence on the precision of variance estimates using high-frequency financial data. Unpublished paper: Graduate School of Business, University of Chicago.
- Barndorff-Nielsen, O. E. (2000). Superposition of Ornstein-Uhlenbeck type processes. *Theory of Probability and Its Applications*. Forthcoming.
- Barndorff-Nielsen, O. E. and N. Shephard (2001a). Non-Gaussian OU based models and some of their uses in financial economics (with discussion). *Journal of the Royal Statistical Society, Series B 63*. Forthcoming, read to the RSS, 18 October 2000.
- Barndorff-Nielsen, O. E. and N. Shephard (2001b). *Lévy Based Dynamic Models for Financial Economics*. Oxford: Draft book, Nuffield College.
- Bollerslev, T., R. F. Engle, and D. B. Nelson (1994). ARCH models. In R. F. Engle and D. McFadden (Eds.), *The Handbook of Econometrics, Volume 4*, pp. 2959–3038. Amsterdam: North-Holland.
- Bollerslev, T. and H. Zhou (2000). Estimating stochastic volatility diffusion using conditional moments of integrated volatility. Unpublished paper: Dept. of Economics, Duke University.
- Comte, F. and E. Renault (1998). Long memory in continuous-time stochastic volatility models. *Mathematical Finance 8*, 291–323.
- Cox, D. R. (1991). Long-range dependence, non-linearity and time irreversibility. *Journal of Time Series Analysis 12*, 329–335.
- Ding, Z. and C. W. J. Granger (1996). Modeling volatility persistence of speculative returns: A new approach. *Journal of Econometrics 73*, 185–215.
- Doornik, J. A. (1998). *Ox: Object Oriented Matrix Programming, 2.0*. London: Timberlake Consultants Press.
- Elerian, O., S. Chib, and N. Shephard (2001). Likelihood inference for discretely observed non-linear diffusions. *Econometrica 69*. Forthcoming.
- Engle, R. F. and G. G. J. Lee (1999). A permanent and transitory component model of stock return volatility. In R. F. Engle and H. White (Eds.), *Cointegration, Causality, and Forecasting. A Festschrift in Honour of Clive W.J. Granger*, Chapter 20, pp. 475–497. Oxford: Oxford University Press.

- Francoq, C. and J.-M. Zakoïan (2000). Covariance matrix estimation for estimators of mixing weak ARMA models. *Journal of Statistical Planning and Inference* 83, 369–94.
- Gallant, A. R. and J. R. Long (1997). Estimating stochastic differential equations efficiently by minimum chi-square. *Biometrika* 84, 125–41.
- Genon-Catalot, V., T. Jeantheau, and C. Larédo (2000). Stochastic volatility as hidden Markov models and statistical applications. *Bernoulli* 6. Forthcoming.
- Ghysels, E., A. C. Harvey, and E. Renault (1996). Stochastic volatility. In C. R. Rao and G. S. Maddala (Eds.), *Statistical Methods in Finance*, pp. 119–191. Amsterdam: North-Holland.
- Gourieroux, C., A. Monfort, and E. Renault (1993). Indirect inference. *Journal of Applied Econometrics* 8, S85–S118.
- Granger, C. W. J. (1980). Long memory relationships and the aggregation of dynamic models. *Journal of Econometrics* 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.
- Koopman, S. J., N. Shephard, and J. A. Doornik (1999). Statistical algorithms for models in state space using SsfPack 2.2. *Econometrics Journal* 2, 107–66.
- Meddahi, N. and E. Renault (2000). Temporal aggregation of volatility models. Unpublished paper: CIRANO, Montreal.
- Nicolato, E. and E. Venardos (2000). Derivative pricing in Barndorff-Nielsen and Shephard’s OU type stochastic volatility models. Unpublished paper: Dept. of Mathematics Sciences, Aarhus University.
- Schwert, G. W. (1989). Why does stock market volatility change over time. *Journal of Finance* 44, 1115–1153.
- Shephard, N. (1996). Statistical aspects of ARCH and stochastic volatility. In D. R. Cox, D. V. Hinkley, and O. E. Barndorff-Nielsen (Eds.), *Time Series Models in Econometrics, Finance and Other Fields*, pp. 1–67. London: Chapman & Hall.
- Sørensen, M. (2000). Prediction based estimating equations. *Econometrics Journal* 3. Forthcoming.
- Taylor, S. J. (1994). Modelling stochastic volatility. *Mathematical Finance* 4, 183–204.
- Taylor, S. J. and X. Xu (1997). The incremental volatility information in one million foreign exchange quotations. *Journal of Empirical Finance* 4, 317–340.

Tomkins, R. and F. Hubalek (2000). On closed form solutions for pricing options with jumping volatility. Unpublished paper: Technical University, Vienna.