

How accurate is the asymptotic approximation to the distribution of realised variance?

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

June 2002

www.levyprocess.org

Prepared for a volume on “Identification and Inference for Econometric Models,” edited by Donald W.F. Andrews, James L. Powell, Paul A. Ruud and James H. Stock (festschrift for Thomas J. Rothenberg). To be published by Cambridge University Press.

Abstract

In this paper we study the reliability of the mixed normal asymptotic distribution of realised variance error, which we have previously derived using the theory of realised power variation. Our experiments suggest that the asymptotics is reliable when we work with the logarithmic transform of the realised variance.

Keywords: Lévy process; Mixed Gaussian limit; OU process; Quadratic variation; Realised power variation; Realised variance; Realised volatility; Square root process; Stochastic volatility; Superposition.

1 Introduction

Tom Rothenberg’s outstanding teaching and research has raised the level of understanding econometricians have of the asymptotic properties of estimators and testing procedures used in economics. His frequent trips away from the United States, and his particular kindness to research students during his academic visits, has spread his influence changing the way we carry out theoretical econometric research. This paper touches on some of Tom’s research interests. It will look at the effectiveness of an asymptotic theory. His influential paper Rothenberg (1984) was devoted to issues of this type.

1.1 The model

This paper assesses the accuracy of the mixed normal asymptotic approximation to the distribution of realised variance (that is the sum of squares of financial returns) we recently derived

in Barndorff-Nielsen and Shephard (2002b) and extended in Barndorff-Nielsen and Shephard (2003) and Barndorff-Nielsen and Shephard (2002a). This theory assumes a flexible stochastic volatility (SV) model for log-prices.

In the SV model for log-prices a basic Brownian motion is generalised to allow the volatility term to vary over time. Then the log-price y^* follows

$$y^*(t) = \alpha(t) + \int_0^t \sigma(u)dw(u), \quad t \geq 0, \quad (1)$$

where σ and α is assumed to be stochastically independent of the standard Brownian motion w . We call σ the *instantaneous* or *spot volatility*, σ^2 the corresponding *variance* and α the *mean* process. A simple example of this is

$$\alpha(t) = \mu t + \beta \sigma^{2*}(t), \quad \text{where} \quad \sigma^{2*}(t) = \int_0^t \sigma^2(u)du,$$

in which case we might call β a risk premium. The process σ^{2*} is called the *integrated variance*. Throughout we will assume the following conditions hold with probability one:

(C) $\sigma^2 > 0$ is càdlàg on $[0, \infty)$ and α has the property

$$\hbar^{-3/4} \max_{1 \leq j \leq M} |\alpha(j\hbar) - \alpha((j-1)\hbar)| = o(1), \quad (2)$$

in \hbar .

Condition (C) implies that the α process is continuous and so is predictable. Hence y^* is a rather flexible special semimartingale. See, for example, Back (1991) for a discussion of the economic implications of this type of property. Assumption (C) also allows the volatility to have, for example, deterministic diurnal effects, jumps, long memory, no unconditional mean or be non-stationary. The mean process α is much more constrained (e.g. it cannot be a Lévy process). A rather flexible example of the process is

$$\alpha(t) = \int_0^t g(\sigma^2(s))ds.$$

where g is a smooth function. Note that condition (C) implies that σ^2 and α are bounded Riemann integrable functions on any finite interval $[0, t]$.

Over an interval of time of length $\hbar > 0$, which could represent a day or a month for example, returns are defined as

$$y_i = y^*(i\hbar) - y^*((i-1)\hbar), \quad i = 1, 2, \dots, T, \quad (3)$$

which implies that

$$y_i | \alpha_i, \sigma_i^2 \sim N(\alpha_i, \sigma_i^2) \quad \text{where} \quad \alpha_i = \alpha(i\hbar) - \alpha\{(i-1)\hbar\},$$

while

$$\sigma_i^2 = \sigma^{2*}(i\hbar) - \sigma^{2*}\{(i-1)\hbar\}.$$

Here σ_i^2 is called *actual variance* and α_i is the *actual mean*. Reviews of the literature on the SV 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 (2001).

1.2 Realised variance

Suppose one was interested in estimating actual volatility σ_i^2 using M intra- \hbar observations (that is M equally spaced increments in the interval $(i-1)\hbar$ to $i\hbar$). A natural candidate is the *realised variance*¹

$$[y_M^*]_i = \sum_{j=1}^M y_{j,i}^2 \quad \text{where} \quad y_{j,i} = y^* \left((i-1)\hbar + \frac{\hbar j}{M} \right) - y^* \left((i-1)\hbar + \frac{\hbar(j-1)}{M} \right). \quad (4)$$

It, and its cousin *realised volatility*

$$\sqrt{\sum_{j=1}^M y_{j,i}^2},$$

have been used in financial economics for many years by, for example, Poterba and Summers (1986), Schwert (1989), Taylor and Xu (1997), Christensen and Prabhala (1998), Andersen, Bollerslev, Diebold, and Labys (2001) and Andersen, Bollerslev, Diebold, and Ebens (2001). However, until recently little theory was known about realised variance outside the Brownian motion case. See the review by Andersen, Bollerslev, and Diebold (2003).

In independent and concurrent work Barndorff-Nielsen and Shephard (2001) and Andersen and Bollerslev (1998) pointed out that the theory of quadratic variation (e.g. Jacod and Shiryaev (1987, p. 55), Protter (1990) and Back (1991)) implies $[y_M^*]_i$ is a consistent estimator of σ_i^2 as $M \rightarrow \infty$. This is an interesting result for it is semi-parametric — it does not depend upon the exact form of α or σ^2 . Unfortunately quadratic variation does not provide a theory of the magnitude of the *realised variance error*

$$u_i = [y_M^*]_i - \sigma_i^2 \quad \text{or} \quad e_i = [y_M^*]_i / \sigma_i^2$$

the *realised variance ratio*. This is important, for although modern econometricians routinely have transaction based data, continuous sample paths processes such as SV models are rather

¹Sums of squared returns are often called realised volatility in econometrics, while we use the name realised variance for that term and realised volatility for the corresponding square root. The use of volatility to denote standard deviations rather than variances is standard in financial economics. See, for example, the literature on volatility and variance swaps, which are derivatives written on realised volatility or variance, which includes Demeterfi, Derman, Kamal, and Zou (1999), Howison, Rafailidis, and Rasmussen (2000) and Chriss and Morokoff (1999). We have chosen to follow this nomenclature rather than the one more familiar in econometrics. Confidence intervals for the realised volatility follow by square rooting the confidence intervals for the realised variance.

poorly fitting at very short time horizons. There are a number of reasons for this, mostly due to market microstructure effects. In particular assets are usually quoted and traded on fixed meshes of points (e.g. decimals or eighths), while the quoting and trading process tends to occur at irregular points in time. See Bai, Russell, and Tiao (2000) who discusses the impact of these type of effects on realised variance. The implication of this is that it is dangerous to make inference based on extremely large values of M for the effect of model misspecification can swamp the effects we are trying to measure. Instead it seems sensible to use moderate values of M and properly account for the fact that the realised variance error is not negligible.

To see that realised variance error can be substantial, we have carried out a small simulation. This could have been based on the familiar constant elasticity of variance (CEV) process which is the solution to the SDE

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

where b is standard Brownian motion uncorrelated with w . 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 have been heavily favoured by Meddahi and Renault (2002) in this context. Instead of this we will mainly work with the non-Gaussian Ornstein-Uhlenbeck, or OU type for short, process which is the solution to the

$$d\sigma^2(t) = -\lambda \sigma^2(t) dt + dz(\lambda t), \tag{5}$$

where z is a subordinator (that is a Lévy process with non-negative increments). These models have been developed in this context by Barndorff-Nielsen and Shephard (2001). In Figure 1 we have taken $\alpha = 0$ and drawn a curve to represent a simulated sample path of σ_i^2 from an OU process where $\sigma^2(t)$ has a $\Gamma(4, 8)$ stationary distribution, $\lambda = -\log(0.99)$ and $\hbar = 1$, along with the associated realised variance (depicted using crosses) computed using a variety of values of M . In this case z is a simple compound Poisson process whose details are given in, for example, Barndorff-Nielsen and Shephard (2001). We see that as M increases the precision of realised variance increases, while Figure 1.d shows that the variance of the realised variance error increases with the actual variance. We will return to this important observation in a moment.

1.3 Mixed normal asymptotic theory

In a recent paper Barndorff-Nielsen and Shephard (2002b), consequently extended in Barndorff-Nielsen and Shephard (2003) and Barndorff-Nielsen and Shephard (2002a), have strengthened the above consistency result $[y_M^*]_i \xrightarrow{P} \sigma_i^2$ considerably. The results have two parts. The first gives the asymptotic distribution of realised variance.

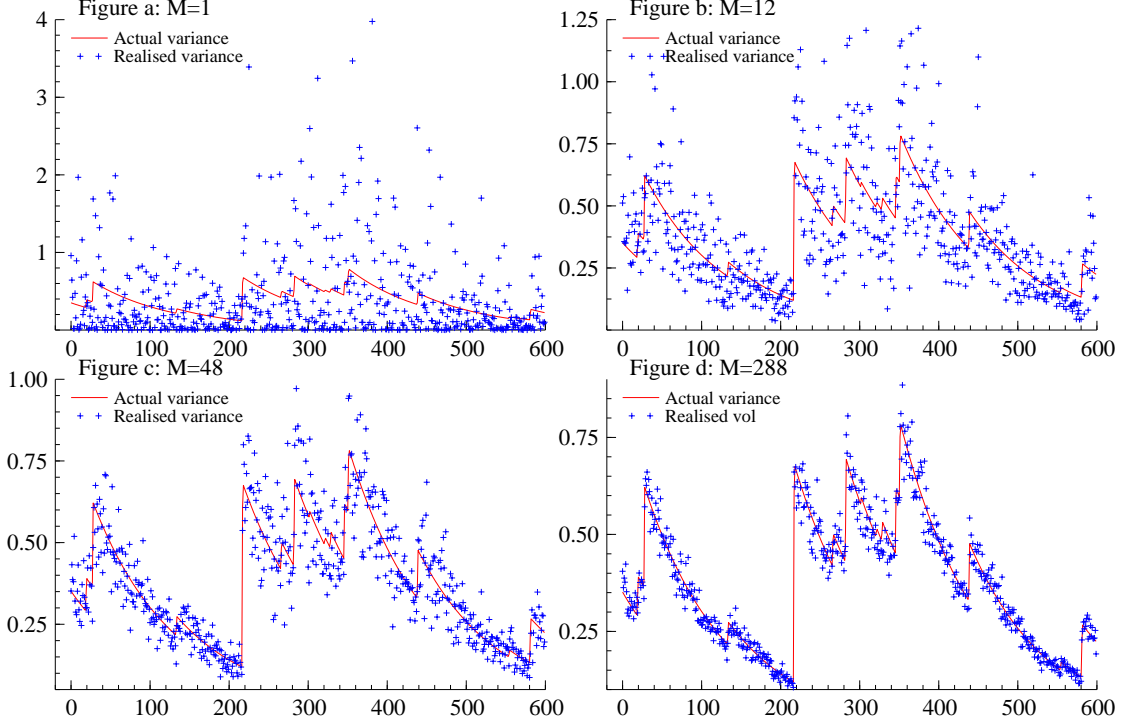


Figure 1: *Actual σ_i^2 and realised $[y_M^*]_i$ (with M varying) variance based upon a $\Gamma(4, 8)$ -OU process with $\lambda = -\log(0.98)$ and $\hbar = 1$. This implies $\xi = 0.5$ and $\xi\omega^{-2} = 8$. The file containing the code used to carry out these calculations is called `ssf_mse.oa`.*

Theorem 1 *For the SV model in (1), for any positive \hbar and $M \rightarrow \infty$*

$$\frac{\sqrt{\frac{M}{\hbar}} ([y_M^*]_i - \sigma_i^2)}{\sqrt{2\sigma_i^{[4]}}} \xrightarrow{\mathcal{L}} N(0, 1), \quad \text{where } \sigma_i^{[4]} = \int_{(i-1)\hbar}^{i\hbar} \sigma^4(u) du. \quad (6)$$

□

We call σ^4 and $\sigma_i^{[4]}$ the *spot* and *actual quarticity*, respectively. This theorem implies the mixed normal limit

$$\sqrt{\frac{M}{\hbar}} ([y_M^*]_i - \sigma_i^2) | \sigma_i^{[4]} \xrightarrow{\mathcal{L}} N\left(0, 2\sigma_i^{[4]}\right). \quad (7)$$

Of course the problem with this theory is that $\sigma_i^{[4]}$ is unknown. This is tackled by the following theorem on realised power variation which is due to Barndorff-Nielsen and Shephard (2003).

Theorem 2 *Let y^* be a stochastic process of the form (1). Define the realised power variation (of order $2q$)*

$$[y_M^*]_i^{[2q]} = \sum_{j=1}^M y_{j,i}^{2q}.$$

Then, for $M \rightarrow \infty$,

$$(\hbar/M)^{-q+1} c_q [y_M^*]_i^{[2q]} \xrightarrow{p} \int_{\hbar(i-1)}^{i\hbar} \sigma^{2q}(u) du,$$

where q denotes a positive integer and $c_q = \{1 \cdot 3 \cdot \dots \cdot (2q - 1)\}^{-1}$.

□

The case of $q = 1$ is the standard quadratic variation result. When $q = 2$ we obtain

$$\frac{M}{\hbar} [y_M^*]_i^{[4]} = \frac{M}{\hbar} \sum_{j=1}^M y_{j,i}^4 \xrightarrow{p} 3\sigma_i^{[4]}.$$

An implication of this is that we can use the feasible limit theory

$$\frac{[y_M^*]_i - \sigma_i^2}{\sqrt{\frac{2}{3} \sum_{j=1}^M y_{j,i}^4}} \xrightarrow{\mathcal{L}} N(0, 1). \quad (8)$$

Of course in practice it may make sense to transform the above limit theorem to impose, a priori, positivity on the approximating distribution (see, for example, the discussion of transformations in Rothenberg (1984, pp. 887–889)). In particular it seems natural to work with the logarithmic transformation of the realised variance ratio so that

$$\frac{\sqrt{\frac{M}{\hbar}} \{\log [y_M^*]_i - \log \sigma_i^2\}}{\sqrt{2\sigma_i^{[4]} / (\sigma_i^2)^2}} \xrightarrow{\mathcal{L}} N(0, 1)$$

and so, using the realised power variation theory,

$$\frac{\log [y_M^*]_i - \log \sigma_i^2}{\sqrt{\frac{2}{3[y_M^*]_i^2} \sum_{j=1}^M y_{j,i}^4}} \xrightarrow{\mathcal{L}} N(0, 1). \quad (9)$$

It is not clear, without further study, whether the log-based asymptotic theory (9) is more or less accurate in finite samples than the non-transformed version (8).

The following remarks can be made about these results.

- $\sum_{j=1}^M y_{j,i}^2$ converges to $\int_{\hbar(i-1)}^{\hbar i} \sigma^2(u) du$ at rate \sqrt{M} . This considerably strengthens the quadratic variation result, for now we know the rate of convergence, not just that it converges.
- The limit theorem is unaffected by the form of the drift process α , smoothness assumption (C) is sufficient that its effect becomes negligible. Again this considerably strengthens the quadratic variation result which says the p-lim is unaffected by the drift. Now we know this result extends to the next order term as well.
- Knowledge of the form of the volatility dynamics is not required in order to use this theory. In a sense this is a semi-parametric result.

- The fourth moment of returns need not exist for the asymptotic normality to hold. In such heavy tailed situations, the stochastic denominator $\int_{(i-1)\hbar}^{i\hbar} \sigma^4(u)du$ loses its unconditional mean. However, this property is irrelevant to the workings of the theory.
- The volatility process $\sigma^2(t)$ can be non-stationary, exhibit long-memory or include intra-day effects.
- $\sum_{j=1}^M y_{j,i}^2 - \int_{\hbar(i-1)}^{\hbar i} \sigma^2(u)du$ has a mixed Gaussian limit implying that marginally it will have heavier tails than a normal.
- The magnitude of the error $\sum_{j=1}^M y_{j,i}^2 - \int_{\hbar(i-1)}^{\hbar i} \sigma^2(u)du$ is likely to be large in times of high volatility. This can be seen in the simulation results in Figure 1.
- Conditionally on $\int_{\hbar(i-1)}^{\hbar i} \sigma^4(u)du$ and $\int_{\hbar(k-1)}^{\hbar k} \sigma^4(u)du$, the errors

$$\sum_{j=1}^M y_{j,i}^2 - \int_{\hbar(i-1)}^{\hbar i} \sigma^2(u)du \quad \text{and} \quad \sum_{j=1}^M y_{j,k}^2 - \int_{\hbar(k-1)}^{\hbar k} \sigma^2(u)du$$

are asymptotically independent and jointly normal for $i \neq k$.

- Some of the features of (6) appear in the usual cross-section asymptotic theory of the estimation of σ^2 when $z_i \sim NID(0, \sigma^2)$. Then

$$\frac{\sqrt{M} \left\{ \frac{1}{M} \sum_{j=1}^M z_i^2 - \sigma^2 \right\}}{\sqrt{2\sigma^4}} \xrightarrow{L} N(0, 1),$$

whose natural feasible version is

$$\frac{\sqrt{M} \left\{ \frac{1}{M} \sum_{j=1}^M z_i^2 - \sigma^2 \right\}}{\sqrt{\frac{2}{3M} \sum_{j=1}^M z_i^4}} \xrightarrow{L} N(0, 1).$$

This has quite a few differences from (8). In particular the denominator divides by M rather than multiplies by M , while in the numerator $\sum_{j=1}^M z_i^2$ is divided by M while in the theory for realised variance it is left unscaled. Bartlett and Kendall (1946) have studied the asymptotic and finite sample behaviour of $\log \sum_{j=1}^M z_i^2$ in this case.

- These results are also quite closely related to the work of Foster and Nelson (1996) (note also the work of Genon-Catalot, Laredo, and Picard (1992), Florens-Zmirou (1993) and Hansen (1995)). In the case where the volatility follows a scalar diffusion, they provided an asymptotic distribution theory for an estimator of $\sigma^2(t)$. Their idea was to compute a local variance from the lagged data, e.g.,

$$\widehat{\sigma^2}(t) = \hbar^{-1} \sum_{j=1}^M \left\{ y^*(t - \hbar j M^{-1}) - y^*(t - \hbar(j-1)M^{-1}) \right\}^2. \quad (10)$$

They then studied its behaviour as $M \rightarrow \infty$ and $\hbar \downarrow 0$ under some assumptions. This “double asymptotics” yields a Gaussian limit theory so long as $\hbar \downarrow 0$ and $M \rightarrow \infty$ at the right, related rates. Of course this type of argument is familiar also in nonparametric econometrics (e.g. Pagan and Ullah (1999)). The double asymptotics makes it harder to use in practice than our own simpler analysis, which just needs $M \rightarrow \infty$. It is possible because our goal is to estimate the easier integrated variation rather than the harder spot variance.

In this paper we use simulation to study the finite sample behaviour of (8) and (9). We will show that (9) works well even for moderately small values of M , while (8) will typically require too large a value of M to be empirically very reliable, although it is a helpful guide.

This paper has three other sections. In Section 2 we run Monte Carlo experiments to assess the finite sample behaviour of (8) and (9). In Section 3 some new theory is introduced which improves the small sample performance of the statistics. Section 4 draws conclusions from our paper.

2 How good is the feasible asymptotic distribution?

2.1 Simple model

In this Section we will use simulation to assess the accuracy of our asymptotic approximations (8) and (9). Throughout this subsection the simulations will be based upon the type of OU based variance process described in the introduction. In particular we will work with a process where σ^2 has a $\Gamma(4, 8)$ stationary distribution, $\lambda = -\log(0.99)$ and $\hbar = 1$.

Figure 2 and Table 1 shows the results from the use of the two asymptotic results (8) and (9) in cases where $M = 12$ and $M = 48$. In the Figure the dots in the left hand graphs show a short sequence of realised variance errors $[y_M^*]_i - \sigma_i^2$, together with plus and minus twice their associated standard errors $\sqrt{\frac{2}{3} \sum_{j=1}^M y_{j,i}^4}$. The graphs show dramatic increases and decreases in the error bands. With $M = 12$ the bands range from 0.1 to around 0.75. These fluctuations correspond to increases and decreases in the overall level of the variance process, with wide bands occurring at periods of high levels of variance. The right hand graphs in Figure 2 give a corresponding normal QQ-plot for (8), which should lay on the 45 degree line if the asymptotics were to hold. The plot is calculated off a larger simulation run than that used to draw the plots in the middle and left hand side of the figure. It uses $T = 1,500$. The graph suggests the asymptotic theory provides a very poor guide to the finite sample behaviour of the standardised realised variance.

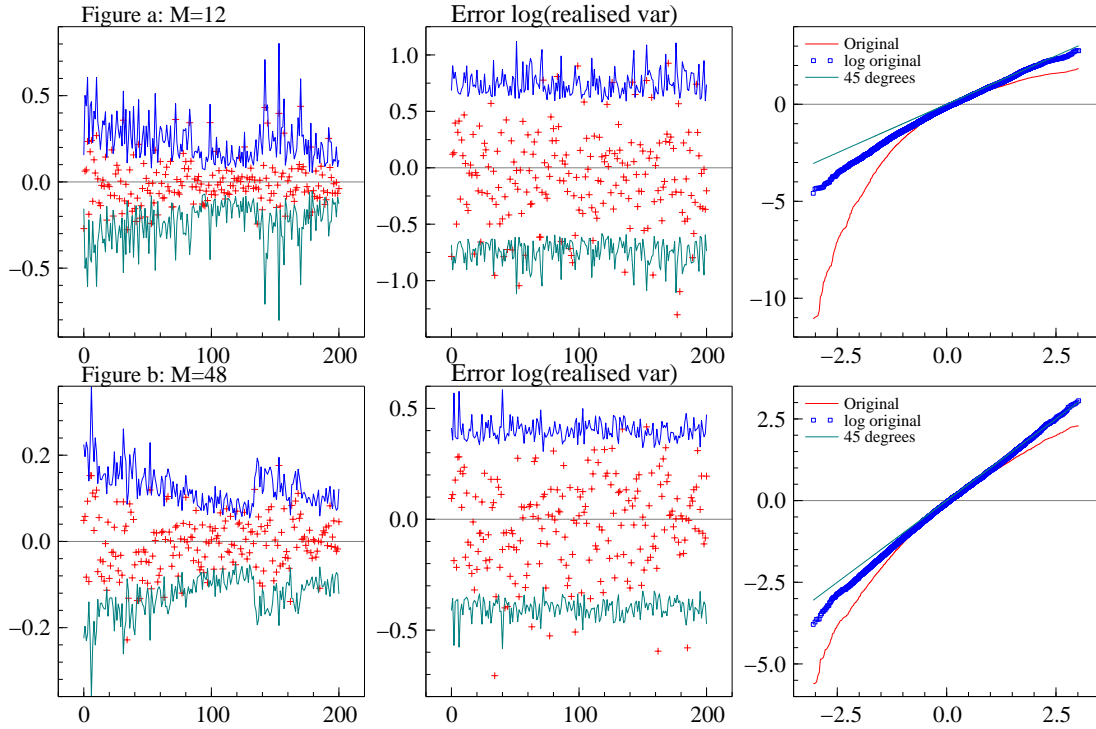


Figure 2: Actual $[y_M^*]_i - \sigma_i^2$ and twice asymptotic S.E.s. Middle graphs: $\log [y_M^*]_i - \log \sigma_i^2$ and twice asymptotic S.E.s. QQ plot of the standardised realised variance error (X-axis has the expected quantiles, Y-axis the observed). Code: `ssf_main.ox`.

The bottom graphs in Figure 2, which correspond to having $M = 48$, show similar effects. However, the standard error bands have sharply contracted, now ranging from 0.08 up to 0.3. The QQ plot is much better, although the asymptotic theory still only provides a rather rough guide to the finite sample behaviour.

The results for the log-based asymptotic theory (9) are much better. The top graphs shows that, even with $M = 12$, the standard error for $\log [y_M^*]_i - \log \sigma_n^2$ does not vary dramatically with n . A failing of the approximation is that there are quite a few errors which are extremely negative. This is picked up in the QQ plot which is much better than any of the ones we have so far seen, but fails in the left hand tail to a much greater degree than in the right hand tail. By the time M has reach 48 the asymptotics seems to give a rather better guide to the behaviour of the distribution. Further, the standard errors have again become much smaller, reducing from around 0.8 down to 0.35.

When we look at higher values of M these broad conclusions continue to hold, with the log based asymptotics substantially out performing the non-transform version. This can be seen in Figure 3, which looks at the cases where $M = 96$ and $M = 288$. Overall the Figure shows it is preferable to rely on the log-based theory.

M	Raw			Log		
	Bias	S.E.	Cove	Bias	S.E.	Cove
12	-.531	1.64	86.5	-.222	1.17	91.2
48	-.226	1.14	92.4	-.110	1.05	94.1
96	-.152	1.07	93.9	-.075	1.03	94.5
288	-.091	1.02	95.0	-.049	1.00	95.0

Table 1: *Bias and standard error of the realised variance errors using the raw asymptotics and the log-based asymptotics. Simulations use a Γ -OU process variance model. Cove denotes estimated finite sample coverage using the asymptotic theory setting the nominal level at 95.0. The Table deals with the no leverage ($\rho = 0$) and strong leverage ($\rho = -1$) cases. File: simple.ox.*

Table 1 gives an alternative view on these simulations. It records the bias from zero of the standardised realised variances (8) and log realised variances (9), together with their standard error. The standard error should be around one if the asymptotic theory is a good description of the behaviour of the statistics. Finally, the Table records the coverage rate of the statistic. This is the percentage of standardised statistics which are larger than two in absolute value.

The results again suggest the raw statistic has poorer behaviour in terms of bias, standard error and coverage compared to the log-version. By the time M reaches 48 the log-version of the statistic seems quite well approximated by the asymptotic theory.

2.2 Superposition

Similar results hold when we build more sophisticated models based on a superposition of OU type models. Such processes also have potential for modelling long-range dependence and self-similarity in variance. This is discussed in the OU case in Barndorff-Nielsen and Shephard (2001) and at more depth by Barndorff-Nielsen (2001) who formalises the use of superpositions as a way of modelling long-range dependence and approximate self-similarity. This follows earlier related work by Granger (1980), Cox (1991), Ding and Granger (1996), Engle and Lee (1999), Shephard (1996, pp. 36–37), Andersen and Bollerslev (1997), Barndorff-Nielsen, Jensen, and Sørensen (1998) and Comte and Renault (1998).

Consider variance based on the sum of J independent OU processes

$$\sigma^2(t) = \sum_{i=1}^J \sigma^{2(i)}(t), \quad \text{where} \quad d\sigma^{2(i)}(t) = -\lambda_i \sigma^{2(i)}(t)dt + dz^{(i)}(\lambda_i t)$$

where the $\sigma^{2(i)}(t)$ process has the memory parameter λ_i . We parameterise the model, in the simplest possible way in the following gamma based example, so that

$$\sigma^{2(i)}(t) \sim \Gamma(w_i \nu, \alpha), \quad \text{where} \quad \{w_i \geq 0\} \quad \text{and} \quad \sum_{i=1}^J w_i = 1.$$

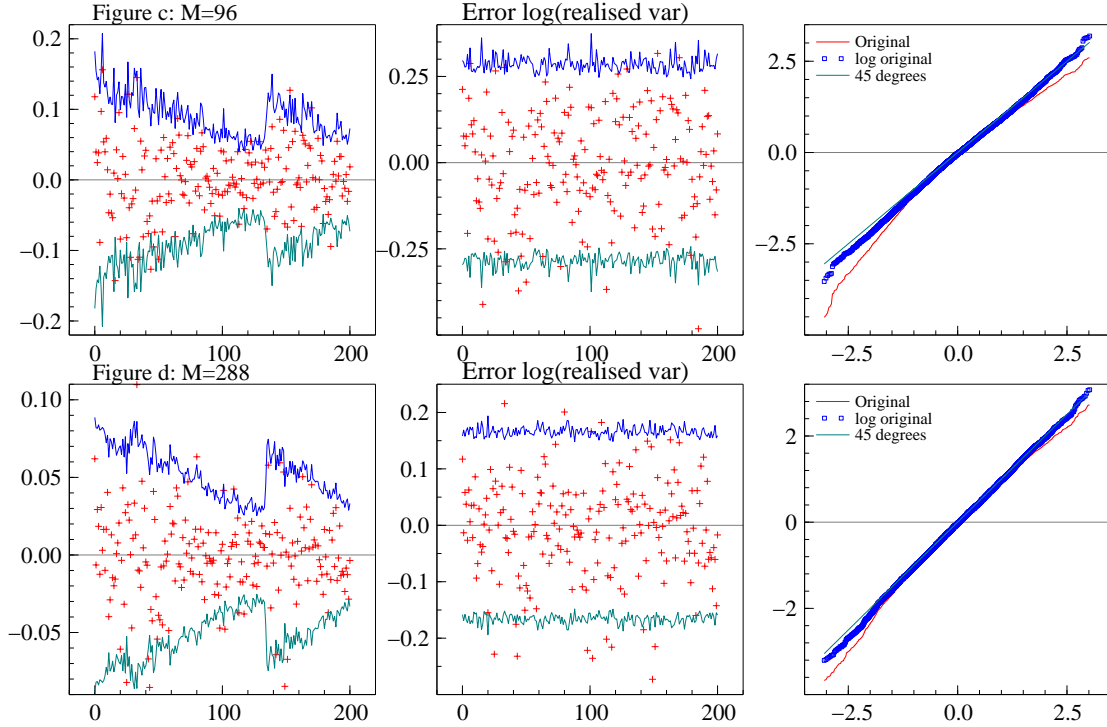


Figure 3: Actual $[y_M^*]_i - \sigma_i^2$ and twice asymptotic S.E.s. Middle graphs: $\log [y_M^*]_i - \log \sigma_i^2$ and twice asymptotic S.E.s. QQ plot of the standardised realised variance error for $M = 96$ and $M = 288$. (X-axis has the expected quantiles, Y-axis the observed). Code: `ssf_main.ox`.

Barndorff-Nielsen and Shephard (2002b) showed such extensions to the basic OU model were necessary in order to satisfactorily fit high frequency exchange rate data. Here we repeat the above analysis with $\nu = 4$ and $\alpha = 8$, but take $J = 2$, $w_1 = 0.8$ and $\lambda_1 = 4$, $\lambda_2 = 0.03$. This means that the second component in the variance has considerable memory, while the first component has very little indeed.

The normal QQ plots for the standardised asymptotic realised variance errors are given in Figure 4. They show the results for $M = 12$, $M = 48$, $M = 96$ and $M = 288$. We broadly repeat the results from the OU case, with the log version of the asymptotics being reasonably reliable even for moderate values of M , particularly in the right hand tail of the distribution. On the other hand the non-transformed version again requires a high value of M to yield satisfactory results.

2.3 Diffusion case

This section will repeat the experiments reported in the previous subsections but this time based on the Feller or Cox, Ingersoll, and Ross (1985) square root process for the volatility dynamics. In the context of SV models this is often called the Heston (1993) model. The experiment has

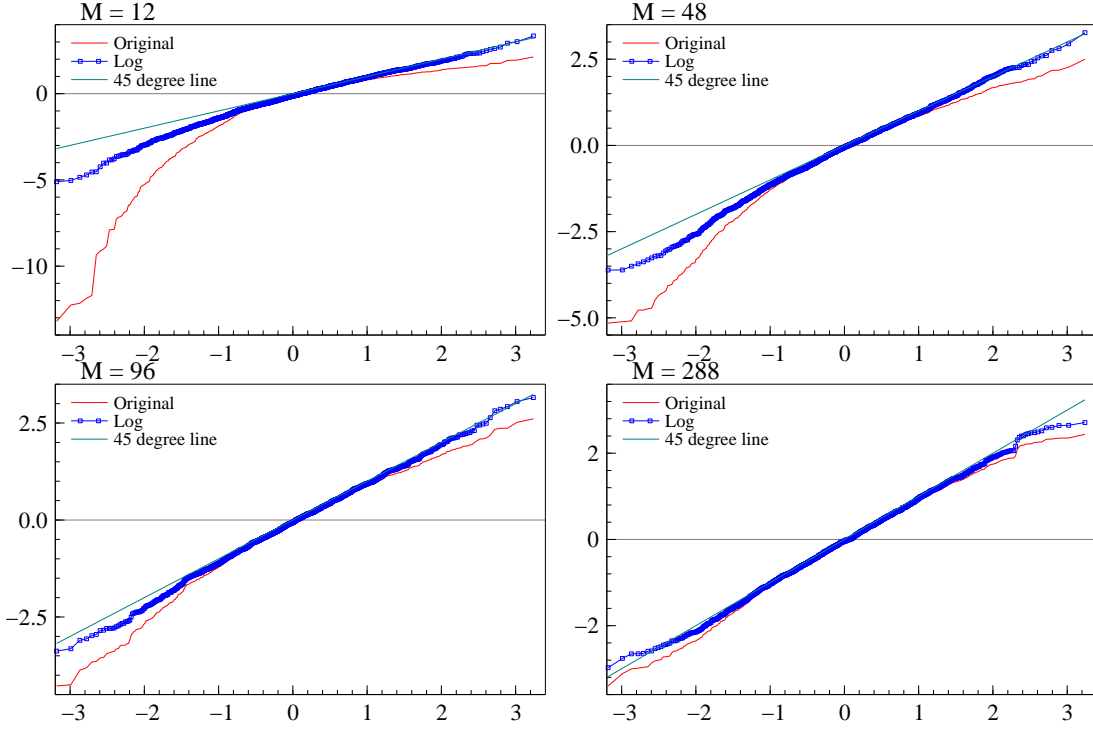


Figure 4: *Superposition case. QQ plot for standardised $\log [y_M^*]_i - \log \sigma_i^2$. Based on $M = 12, 48, 96$ and $M = 288$. (X-axis has the expected quantiles, Y-axis the observed). Code: `ssf_main.ox`.*

two aims:

- To demonstrate that the results we indicated above are not sensitive to the type of volatility processes used in building the model.
- To explore the effect of leverage terms (that is correlation between the returns and future volatility movements) on the performance of our theory. This is interesting as it is outside the assumptions that Barndorff-Nielsen and Shephard (2002b), Barndorff-Nielsen and Shephard (2003) and Barndorff-Nielsen and Shephard (2002a) have been able to prove the asymptotics for realised variance.

We write the dynamics as

$$dy^*(t) = \sigma(t)dw(t)$$

and

$$d\sigma^2(t) = -\lambda \{ \sigma^2(t) - \xi \} dt + \omega \sigma(t) db(\lambda t), \quad \xi \geq \omega^2/2 \quad (11)$$

where $b(t)$ is a standard Brownian motion process. To allow for the possibility of leverage we will assume

$$\text{Cor} \{ b(\lambda t), w(t) \} = \rho t \sqrt{\lambda}.$$

The correlation parameter ρ indexes the leverage effect in the model and would be expected to be negative for equity data (e.g. Black (1976) and Nelson (1991)). The square root process has a marginal distribution

$$\sigma^2(t) \sim \Gamma(2\omega^{-2}\xi, 2\omega^{-2}) = \Gamma(\nu, a), \quad \nu \geq 1,$$

with a mean of $\xi = \nu/a$ and a variance of $\omega^2 = \nu/a^2$. Throughout this section we again take $\hbar = 1$ and $\nu = 4$ and $a = 8$.




Figure 5: *Simulation from a diffusion based SV model based on the square root volatility processes with low and high amounts of leverage. QQ plots are drawn. Top line has $M=72$, bottom $M=288$. Left hand graphs show results for no leverage, right hand graphs high leverage. Drawn is the results for the raw and log based theory. Code is available at: heston.ox*

The results, based on 10,000 replications, in the no leverage case are given in the left hand part of Figure 5. It shows that moving to the diffusion based volatility model does not really change any of the conclusions from the previous subsection — the asymptotics still provide a useful guide to the finite sample behaviour of these statistics. Although not surprising, since this case is covered by our theory, this is a reassuring result.

These results are reinforced by Table 2 which shows the mean and standard error of the normalised statistics (8) and (9). In the Table the former is called the raw statistics, the latter the log version. The Table shows that there is a negative bias in the raw statistic, which corresponds to the realised variance being too small and at the same time the corresponding

denominator being too small. This bias is an order of magnitude smaller on the log-version of the statistic. In both cases the standard error of the normalised statistic is roughly one.

The Table also gives some results on the coverage performance of the asymptotic theory. This records the percentage of times the realised variance minus the actual variance is larger, in absolute value, than twice the feasible asymptotic standard error. Thus, if the asymptotic theory was exact then we would expect the coverage percentage to be 95. The results suggest that this is not a poor approximation for moderately large values of M .

The results with strong leverage are given in the right hand side of Figure 5². They are very much in line with those we reported for the non-leverage case and suggest that our analysis may be robust to this effect. The right hand side of Table 2 confirms these observations. Proving this conjecture turns out to be challenging mathematically even in the univariate case and is the subject of on-going research. In a stimulating piece of work Meddahi (2002) has shown that the effect of leverage on the unconditional mean square error of the realised covariation error is asymptotically negligible in a wide class of diffusion based volatility models. This again points to the more general result of our asymptotics working in the leverage case.

M	No leverage						Strong leverage					
	Raw			Log			Raw			Log		
	Bias	S.E.	Cove	Bias	S.E.	Cove	Bias	S.E.	Cove	Bias	S.E.	Cove
18	-.428	1.40	88.4	-.203	1.12	92.0	-.411	1.38	88.8	-.190	1.11	92.5
36	-.275	1.18	91.8	-.137	1.06	94.0	-.262	1.18	91.9	-.125	1.06	93.7
72	-.196	1.09	92.9	-.104	1.04	94.1	-.182	1.08	93.5	-.093	1.03	94.4
288	-.099	1.03	94.5	-.055	1.02	94.8	-.083	1.03	94.4	-.039	1.02	94.6
576	-.067	1.02	94.8	-.036	1.01	95.2	-.069	1.01	95.2	-.039	1.00	95.2
1152	-.034	1.00	95.3	-.013	1.00	95.3	-.061	1.00	95.1	-.040	1.00	95.1
2304	-.024	1.00	95.5	-.009	1.00	95.5	-.036	1.00	95.5	-.021	1.00	95.5

Table 2: *Bias and standard error of the realised variance errors using the raw asymptotics and the log-based asymptotics. Simulations use a CIR variance model. Cove denotes estimated finite sample coverage using the asymptotic theory setting the nominal level at 95.0. The Table deals with the no leverage ($\rho = 0$) and strong leverage ($\rho = -1$) cases. File: simple.ox.*

2.4 Alternative estimators of quarticity

An important feature of the above results is the use of empirical quarticity as $M \rightarrow \infty$

$$\frac{M}{h} \sum_{j=1}^M y_{j,i}^4 \xrightarrow{p} 3 \int_{(i-1)h}^{ih} \sigma^4(u) du,$$

which allows us to produce a feasible limit theory for realised variance. A difficulty with estimating quarticity using the fourth powers of returns is that it is not very robust. In some

²In the simulation, common random numbers are used in the leverage and non-leverage cases.

recent work Barndorff-Nielsen and Shephard (2002a) managed to avoid this by showing that as $M \rightarrow \infty$ so for $s > 0$ the time series estimator

$$\frac{M}{h} \sum_{j=1}^{M-s} y_{j,i}^2 y_{j+s,i}^2 \xrightarrow{p} \int_{(i-1)h}^{ih} \sigma^4(u) du$$

This can be used to provide an alternative to the standard feasible limit theory (8). In particular it implies that the standardised realised variance

$$\frac{[y_M^*]_i - \sigma_i^2}{\sqrt{2S^{-1} \sum_{s=1}^S \sum_{j=1}^{M-s} y_{j,i}^2 y_{j+s,i}^2}} \xrightarrow{L} N(0, 1). \quad (12)$$

Here we study its finite sample behaviour, as well as the log-transformed version

$$\frac{\log[y_M^*]_i - \log \sigma_i^2}{\sqrt{\frac{2}{[y_M^*]_i^2} S^{-1} \sum_{s=1}^S \sum_{j=1}^{M-s} y_{j,i}^2 y_{j+s,i}^2}} \xrightarrow{L} N(0, 1). \quad (13)$$

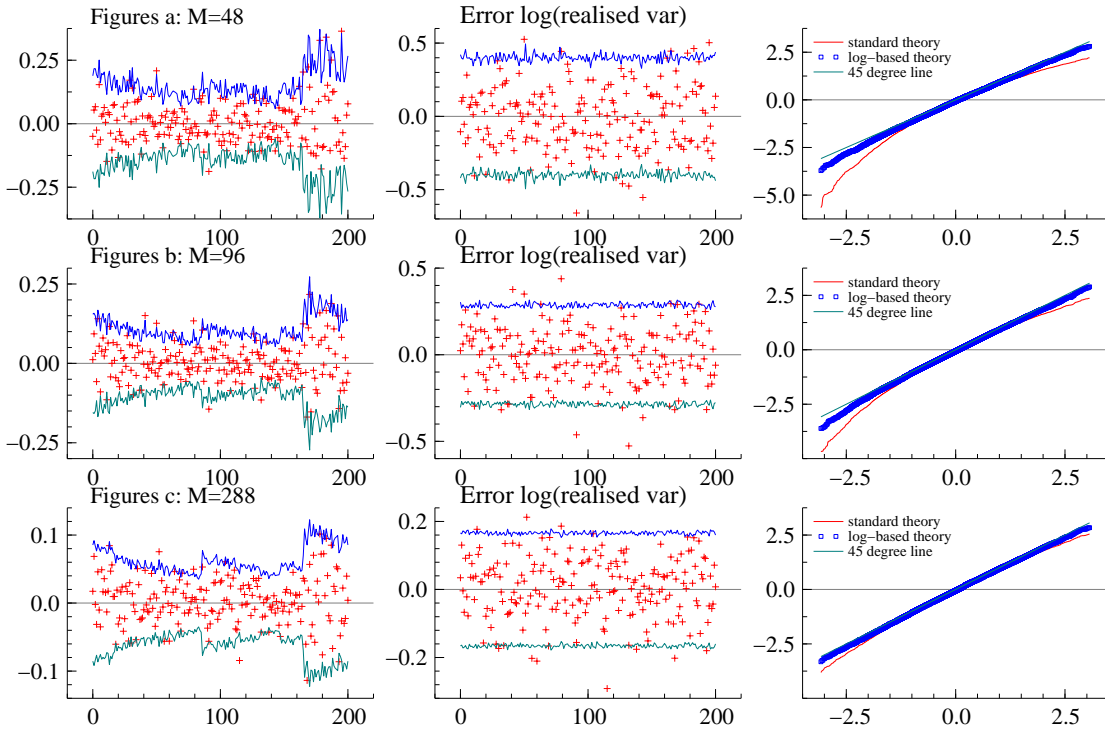


Figure 6: *Finite sample performance based on an alternative estimator of quarticity. Here we use an average of 4 lags. Actual $[y_M^*]_i - \sigma_i^2$ and twice asymptotic S.E. Middle graphs: $\log [y_M^*]_i - \log \sigma_i^2$ and twice asymptotic S.E.s. QQ plot of the standardised realised variance error for $M = 48, 96$ and $M = 288$. (X-axis has the expected quantiles, Y-axis the observed). Code: `simple.oa`.*

Throughout this subsection we return to using simulations to investigate the finite sample behaviour of these asymptotic results (12) and (13). These will again will be based upon the

type of OU based variance process described in the introduction. In particular we will start work with a process where σ^2 has a $\Gamma(4, 8)$ stationary distribution, $\lambda = -\log(0.99)$ and $\bar{h} = 1$. The time series versions of these results will be based on using (12) and (13) with $S = 4$.

Figure 6 gives the results for the time series based method. Thus this Figure is directly comparable with the results given in Figures 2 and 3. Throughout the results indicate that the new time series based estimator may very slightly improve the finite sample behaviour of the statistics, however the difference is very marginal compared to shifting from the raw asymptotic approximation to the one based on the logarithmic transformation.

To see how these results vary with the model, we have rerun these calculations varying the persistence parameter λ and the marginal distribution of $\Gamma(4, 8)$. Table 3 reinforces the same points as made by the Figure using the bias and the standard errors of the standardised feasible approximations given in (12) and (13). These are shown with the corresponding results (8) and (9) based on the fourth moments reported in the previous sub-Sections. These again show a slight improvement using the time series method compared to those obtained by using the fourth moment.

Model	M	Time series estimator				4th moment estimator			
		Raw		Log		Raw		Log	
		Bias	S.E.	Bias	S.E.	Bias	S.E.	Bias	S.E.
$e^\lambda = e^{-0.99}, \Gamma(4, 8)$	48	-.216	1.11	-.104	1.03	-.224	1.13	-.111	1.04
	96	-.155	1.05	-.0809	1.01	-.158	1.05	-.0832	1.01
	288	-.103	1.02	-.0611	1.00	-.104	1.02	-.0614	1.01
$e^\lambda = e^{-0.9}, \Gamma(4, 8)$	48	-.228	1.13	-.112	1.04	-.238	1.15	-.121	1.06
	96	-.163	1.07	-.0857	1.02	-.165	1.07	-.0881	1.02
	288	-.0961	1.01	-.0541	0.998	-.0970	1.01	-.0548	1.00
$e^\lambda = e^{-0.99}, \Gamma(2, 4)$	48	-.215	1.11	-.101	1.04	-.226	1.14	-.110	1.05
	96	-.150	1.06	-.0734	1.02	-.152	1.07	-.0755	1.03
	288	-.0902	1.01	-.0481	1.00	-.0913	1.02	-.0491	1.00

Table 3: *Bias and standard error of the realised variance errors. First group based on an $\Gamma(4, 8)$ -OU process with $\lambda = -\log(0.99)$. Second group changes to an $\Gamma(4, 8)$ -OU process with $\lambda = -\log(0.9)$. Third to $\Gamma(2, 4)$ -OU process with $\lambda = -\log(0.99)$ File: simple.ox.*

3 Some theoretical observations

3.1 Asymptotic distribution of $\log [y_M^*]_i$

The variance of $\log [y_M^*]_i - \log \sigma_i^2$ seem to only mildly fluctuate with i , with values around $2/M$.

The conditional variance in the mixed normal asymptotic theory is

$$\frac{2}{M} \frac{\hbar^{-1} \int_{(i-1)\hbar}^{i\hbar} \sigma^4(u) du}{\left(\hbar^{-1} \int_{(i-1)\hbar}^{i\hbar} \sigma^2(u) du \right)^2} \geq \frac{2}{M},$$

by Jensen's inequality. If σ is continuous the lower bound is obtained when $\hbar \rightarrow 0$ for

$$\frac{\hbar^{-1} \int_{(i-1)\hbar}^{i\hbar} \sigma^4(u) du}{\left(\hbar^{-1} \int_{(i-1)\hbar}^{i\hbar} \sigma^2(u) du \right)^2} \xrightarrow{a.s.} 1.$$

On the other hand, for $\hbar \rightarrow \infty$ and assuming the spot variance is ergodic then

$$\frac{\hbar^{-1} \int_{(i-1)\hbar}^{i\hbar} \sigma^4(u) du}{\left(\hbar^{-1} \int_{(i-1)\hbar}^{i\hbar} \sigma^2(u) du \right)^2} \xrightarrow{a.s.} \frac{\mathbb{E}\sigma^4(t)}{\{\mathbb{E}\sigma^2(t)\}^2}.$$

In the above examples we had $\sigma^2(t) \sim \Gamma(\nu, \alpha)$ which implies

$$\frac{\mathbb{E}\sigma^4(t)}{\{\mathbb{E}\sigma^2(t)\}^2} = \frac{\nu/\alpha^2 + (\nu/\alpha)^2}{(\nu/\alpha)^2} = 1 + \frac{1}{\nu}.$$

Throughout ν took the value 4, which implies again the lower bound of $2/M$ is a good rough approximation. Of course in examples where the fourth moment of the variance process does not exist then the lower bound will be wildly off.

3.2 Relationship between integrals and sums

In the asymptotics for $[y_M^*]_i$ we replace

$$2\hbar\sigma_i^{[4]} \quad \text{with} \quad \frac{2}{3}M \sum_{j=1}^M y_{j,i}^4,$$

while in the $\log [y_M^*]_i$ case we estimate

$$2\hbar\sigma_i^{[4]} / (\sigma_i^2)^2 \quad \text{by} \quad \frac{2}{3 [y_M^*]_i^2} M \sum_{j=1}^M y_{j,i}^4.$$

Although these sums are consistent estimators of the required integrals, in practice they could be rather noisy. Here we report evidence on this issue using the same gamma based superposition model employed in the previous section.

Figure 7 shows a plot of $2\hbar\sigma_i^{[4]}$ and its estimator $2M/3 \sum_{j=1}^M y_{j,i}^4$ against i for a variety of values of M . The key observation is that the estimator is very noisy when M is moderate, with

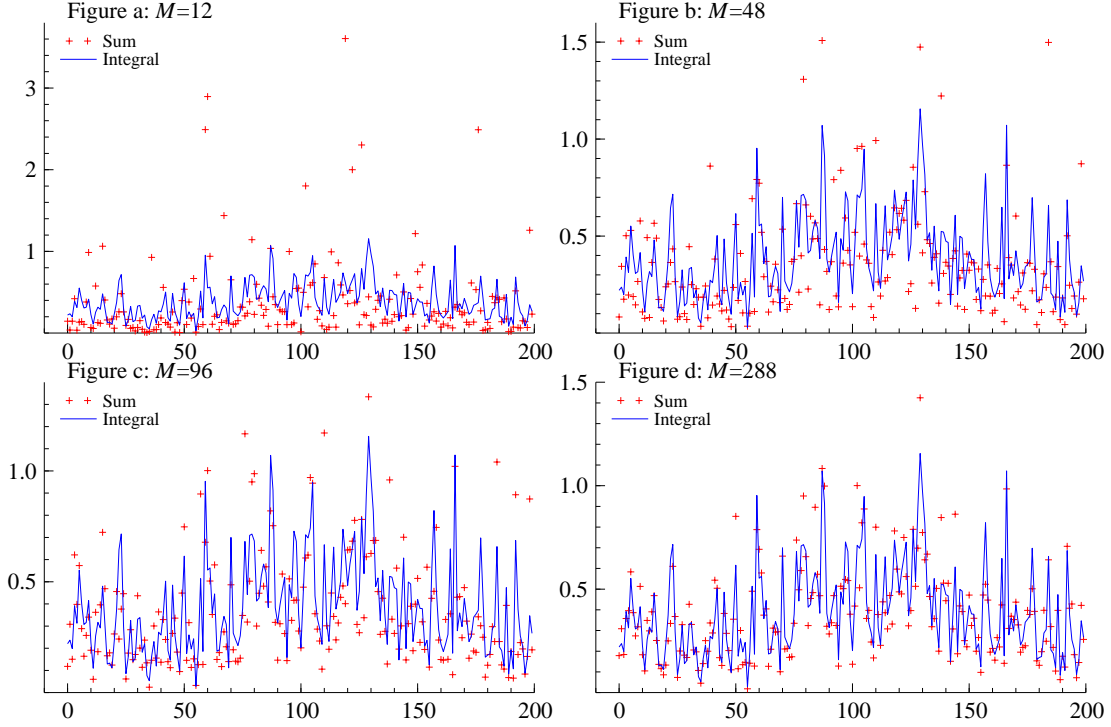


Figure 7: $[y_M^*]_i$ case. Plot of the integrals $2\hbar\sigma_i^{[4]}$ and their consistent estimators $\frac{2M}{3}\sum_{j=1}^M y_{j,i}^4$. Based on $M = 12, 48, 96$ and $M = 288$. Code: `ssf_main.ox`.

quite large values of M needed in order to accurately estimate $2\hbar\sigma_i^{[4]}$. This is a major cause of the poor nature of the finite sample behaviour in the QQ plots we recorded in Figures 1, 2 and 4 above.

Figure 8 gives the corresponding results for the asymptotics for $\log [y_M^*]_i$. It shows that $2\hbar\sigma_i^{[4]}/(\sigma_i^2)^2$ does not vary very much with n and hovers just about 2. Thus by dividing by $(\sigma_i^2)^2$ we have approximately stabilised $\sigma_i^{[4]}$. The plot of $\frac{2M}{3[y_M^*]_i^2}\sum_{j=1}^M y_{j,i}^4$ shows that this estimator is still quite variable, however it is a very substantial improvement over the case of $2M/3\sum_{j=1}^M y_{j,n}^4$. An interesting feature is that the estimator sometimes goes below 2, which we have seen is the lower bound for the ratio of integrals.

3.3 Finite sample corrections

The observation that the feasible bound sometimes goes below 2 suggests imposing it directly in the standardisation formula. In particular we work with

$$\frac{\log [y_M^*]_i - \log \sigma_i^2}{s_i} \xrightarrow{\mathcal{L}} N(0, 1), \quad \text{where} \quad s_i^2 = \max \left\{ \frac{2}{3 [y_M^*]_i^2} \sum_{j=1}^M y_{j,i}^4, \frac{2}{M} \right\}.$$

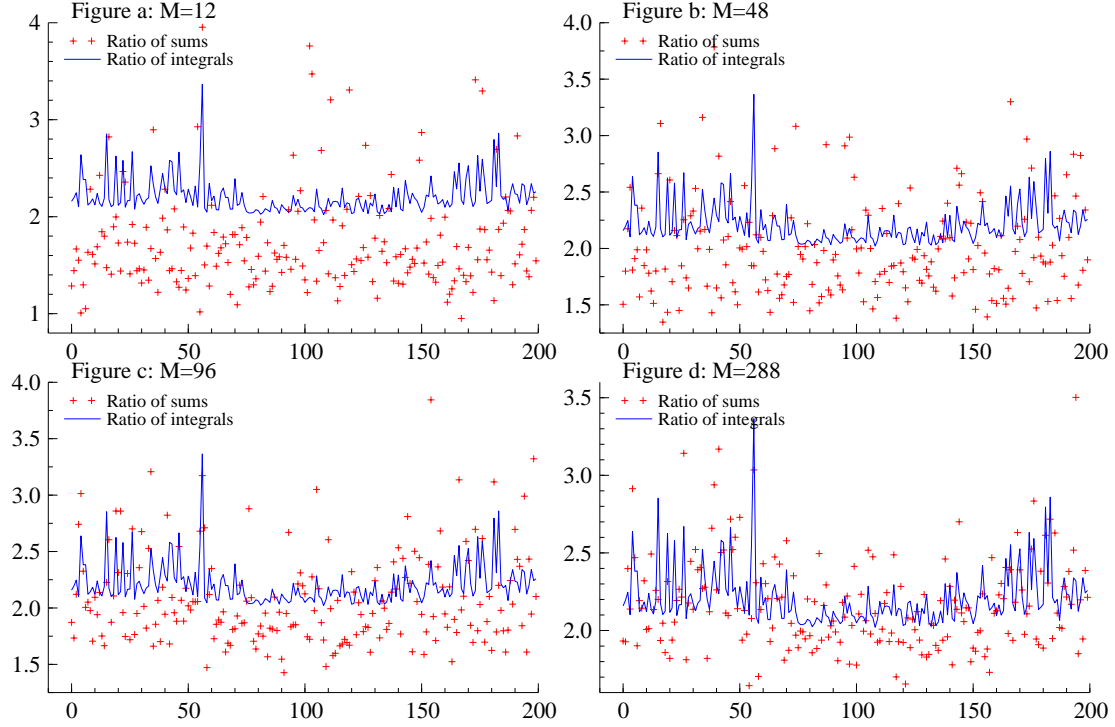


Figure 8: $\log [y_M^*]_i$ case. Plot of the ratio of integrals $2\hbar\sigma_i^{[4]} / (\sigma_i^2)^2$ and their consistent estimators $\frac{2}{3[y_M^*]_i^2} M \sum_{j=1}^M y_{j,i}^4$. Based on $M = 12, 48, 96$ and $M = 288$. Code: `ssf_main.ox`.

It also makes sense to make a finite sample mean correction to this for $[y_M^*]_i$ is an unbiased estimator of σ_i^2 when $\alpha(t) = 0$, but this implies $\log [y_M^*]_i$ will only biasedly estimate $\log \sigma_i^2$. This implies, using approximate log-normality, working with

$$\frac{\log [y_M^*]_i - \log \sigma_i^2 + \frac{1}{2}s_i^2}{s_i} \xrightarrow{\mathcal{L}} N(0,1) \quad \text{where} \quad s_n^2 = \max \left\{ \frac{2}{3 [y_M^*]_i^2} \sum_{j=1}^M y_{j,i}^4, \frac{2}{M} \right\}$$

The corresponding QQ plot given in Figure 9 suggests this improves the finite sample behaviour of the method. Separate calculations indicate there is very little difference between the performance of this statistic and the infeasible version

$$\frac{\log [y_M^*]_i - \log \sigma_i^2 + \frac{1}{2}2\hbar M^{-1}\sigma_i^{[4]} / (\sigma_i^2)^2}{\sqrt{2\hbar M^{-1}\sigma_i^{[4]} / (\sigma_i^2)^2}}.$$

3.4 Alternative transformations

Throughout this paper we have studied the effect of using a log transformation on the realised variance. We have seen this is very beneficial. We now ask if this is the best transformation to use?




Figure 9: $\log [y_M^*]_i$ case. QQ plot for standardised $\log [y_M^*]_i - \log \sigma_i^2$ (denoted *log original*) and the finite sample version $\frac{\log [y_M^*]_i - \log \sigma_i^2 + \frac{1}{2} s_i^2}{s_i}$ (denoted *log with Max* in the Figure). Based on $M = 12, 48, 96$ and $M = 288$.

Consider a natural exponential family

$$p(x; \theta) = a(\theta)b(x)e^{\theta x},$$

where $p(x; \theta)$ is a probability density function and θ are some parameters. Let $\tau(\mu)$ be the variance function of x expressed as a function of the mean μ of x . For each $\kappa \in (0, 1)$ introduce the transformation

$$\eta_\kappa(\mu) = \int \tau(\mu)^{-\kappa} d\mu,$$

i.e. $\eta_\kappa(\mu)$ is the indefinite integral of $\tau(\mu)^{-\kappa}$. Then (subject to the usual type of smoothness conditions) we have (cf. Barndorff-Nielsen (1978, pp.177-179)) the following.

- For $\kappa = \frac{1}{3}$ the skewness of $\eta_\kappa(x)$ is approximately 0.
- For $\kappa = \frac{1}{2}$ the variance of $\eta_\kappa(x)$ is approximately constant.
- For $\kappa = \frac{1}{2}$ the spread of the log likelihood function, expressed in terms of $\eta_\kappa(\mu)$, is approximately constant³.

³The precise statement is as follows. Let $l(\eta)$ denote the log likelihood expressed as a function of $\eta = \eta_{1/2}(\mu)$. Then $l''(\hat{\eta}) = 0$.

- For $\kappa = \frac{2}{3}$ the log likelihood function, expressed in terms of $\eta_k(\mu)$, is approximately ‘normal’⁴.

Example Barndorff-Nielsen and Shephard (2002b) showed that the marginal distribution of realised variance is well approximated by the inverse Gaussian $IG(\delta, \gamma)$ family where $\delta > 0, \gamma > 0$. This has the density

$$f(x) = \frac{\delta}{\sqrt{2\pi}} e^{\delta\gamma x^{-3/2}} \exp\left\{-\frac{1}{2}(\delta^2 x^{-1} + \gamma^2 x)\right\}, \quad x > 0.$$

Further work on this approximation is given in Bollerslev and Forsberg (2002), while Andersen, Bollerslev, Diebold, and Labys (2001) looked at a log-normal approximation. Here δ^2 is a scale parameter and $\tau(\mu) = \delta^{-2}\mu^3$. Hence $\log x$ is approximately normal in the sense of having an approximately symmetric distribution. And $x^{-1/2}$ has approximately constant variance. For fixed value of the mean δ/γ of x the approximation is better the larger the value of δ . \square

The above discussion indicates that the log transformation has particular attractive features in the context of realised variances.

4 Conclusion

In this note we have looked at the finite sample performance of our asymptotic approximation to the distribution of realised variance. The evidence suggests that the finite sample log version of the result:

$$\frac{\log [y_M^*]_i - \log \sigma_i^2 + \frac{1}{2}s_i^2}{s_i} \xrightarrow{\mathcal{L}} N(0, 1) \quad \text{where} \quad s_i^2 = \max\left\{\frac{2}{3[y_M^*]_i^2} \sum_{j=1}^M y_{j,i}^4, \frac{2}{M}\right\},$$

is reasonably reliable for moderate values of M , say 12 or above.

5 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 second author using the Ox language of Doornik (2001). We would like to thank Torben Andersen, Tim Bollerslev and Nour Meddahi for helpful conversations on this topic.

⁴The precise statement is as follows. Let $l(\eta)$ denote the log likelihood expressed as a function of $\eta = \eta_{2/3}(\mu)$. Then $l^{(3)}(\hat{\eta}) = 0$.

References

- Andersen, T. G. and T. Bollerslev (1997). Heterogeneous information arrivals and return volatility dynamics: Uncovering the long-run in high frequency returns. *Journal of Finance* 52, 975–1005.
- Andersen, T. G. and T. Bollerslev (1998). Answering the skeptics: yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39, 885–905.
- Andersen, T. G., T. Bollerslev, and F. X. Diebold (2003). Parametric and nonparametric measurement of volatility. In Y. Ait-Sahalia and L. P. Hansen (Eds.), *Handbook of Financial Econometrics*. Amsterdam: North Holland. Forthcoming.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and H. Ebens (2001). The distribution of realized stock return volatility. *Journal of Financial Economics* 61, 43–76.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys (2001). The distribution of exchange rate volatility. *Journal of the American Statistical Association* 96, 42–55.
- Back, K. (1991). Asset pricing for general processes. *Journal of Mathematical Economics* 20, 371–395.
- 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. (1978). *Information and Exponential Families in Statistical Theory*. Chichester: John Wiley and Sons.
- Barndorff-Nielsen, O. E. (2001). Superposition of Ornstein-Uhlenbeck type processes. *Theory of Probability and its Applications* 45, 175–194.
- Barndorff-Nielsen, O. E., J. L. Jensen, and M. Sørensen (1998). Some stationary processes in discrete and continuous time. *Advances in Applied Probability* 30, 989–1007.
- Barndorff-Nielsen, O. E. and N. Shephard (2001). Non-Gaussian Ornstein–Uhlenbeck-based models and some of their uses in financial economics (with discussion). *Journal of the Royal Statistical Society, Series B* 63, 167–241.
- Barndorff-Nielsen, O. E. and N. Shephard (2002a). Econometric analysis of realised covariation: high frequency covariance, regression and correlation in financial economics. Unpublished paper: Nuffield College, Oxford, Economics working paper 2002-W13.
- Barndorff-Nielsen, O. E. and N. Shephard (2002b). Econometric analysis of realised volatility

- and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society, Series B* 64, 253–280.
- Barndorff-Nielsen, O. E. and N. Shephard (2003). Realised power variation and stochastic volatility. *Bernoulli* 9, 243–265. Correction is available at www.levyprocess.org.
- Bartlett, M. S. and D. G. Kendall (1946). Statistical analysis of variance-heterogeneity and the logarithmic transformation. *Journal of the Royal Statistical Society Supplement* 8, 128–138.
- Black, F. (1976). Studies of stock price volatility changes. *Proceedings of the Business and Economic Statistics Section, American Statistical Association*, 177–181.
- Bollerslev, T. and L. Forsberg (2002). Bridging the gap between the distribution of realized (ECU) volatility and ARCH modeling (of the Euro): The GARCH normal inverse Gaussian model. *Journal of Applied Econometrics* 17, 535–548.
- Chriss, N. and W. Morokoff (1999). Volatility and variance swaps. *Risk* 12, 55–59.
- Christensen, B. J. and N. R. Prabhala (1998). The relation between implied and realized volatility. *Journal of Financial Economics* 37, 125–150.
- 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.
- Cox, J. C., J. E. Ingersoll, and S. A. Ross (1985). A theory of the term structure of interest rates. *Econometrica* 53, 385–407.
- Demeterfi, K., E. Derman, M. Kamal, and J. Zou (1999). A guide to volatility and variance swaps. *Journal of Derivatives* 6, 9–32.
- 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. (2001). *Ox: Object Oriented Matrix Programming, 3.0*. London: Timberlake Consultants Press.
- 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.

- Florens-Zmirou, D. (1993). On estimating the diffusion coefficient from discrete observations. *Journal of Applied Probability* 30, 790–804.
- Foster, D. P. and D. B. Nelson (1996). Continuous record asymptotics for rolling sample variance estimators. *Econometrica* 64, 139–174.
- Genon-Catalot, V., C. Laredo, and D. Picard (1992). Non-parametric estimation of the diffusion coefficient by wavelet methods. *Scandinavian Journal of Statistics* 19, 317–335.
- 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.
- Granger, C. W. J. (1980). Long memory relationships and the aggregation of dynamic models. *Journal of Econometrics* 14, 227–238.
- Hansen, B. E. (1995). Regression with non-stationary volatility. *Econometrica* 63, 1113–1132.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility, with applications to bond and currency options. *Review of Financial Studies* 6, 327–343.
- Howison, S. D., A. Rafailidis, and H. O. Rasmussen (2000). A note on the pricing and hedging of volatility derivatives. Unpublished paper: Mathematical Institute, University of Oxford.
- Jacod, J. and A. N. Shiryaev (1987). *Limit Theorems for Stochastic Processes*. Springer-Verlag: Berlin.
- Meddahi, N. (2002). A theoretical comparison between integrated and realized volatilities. *Journal of Applied Econometrics* 17, 479–508.
- Meddahi, N. and E. Renault (2002). Temporal aggregation of volatility models. *Journal of Econometrics*. Forthcoming.
- Nelson, D. B. (1991). Conditional heteroskedasticity in asset pricing: a new approach. *Econometrica* 59, 347–370.
- Pagan, A. R. and A. Ullah (1999). *Nonparametric Econometrics*. Cambridge: Cambridge University Press.
- Poterba, J. and L. Summers (1986). The persistence of volatility and stock market fluctuations. *American Economic Review* 76, 1124–1141.
- Protter, P. (1990). *Stochastic Integration and Differential Equations: A New Approach*. New York: Springer-Verlag.
- Rothenberg, T. J. (1984). Approximating the distributions of econometric estimators and test statistics. In Z. Griliches and M. D. Intriligator (Eds.), *Handbook of Econometrics*, Volume 2, pp. 881–935. Amsterdam: Elsevier.

- 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.
- 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.