

Modelling and measuring volatility*

OLE E. BARNDORFF-NIELSEN

*The T.N. Thiele Centre for Mathematics in Natural Science,
Department of Mathematical Sciences,
University of Aarhus, Ny Munkegade, DK-8000 Aarhus C, Denmark
& CREATES, University of Aarhus
oebn@imf.au.dk*

NEIL SHEPHARD

*Oxford-Man Institute, University of Oxford,
Blue Boar Court, Alfred Road, Oxford OX1 4EH, UK
& Department of Economics, University of Oxford
neil.shephard@economics.ox.ac.uk*

November 21, 2008

Abstract

Keywords: Lévy process, realised volatility, realised kernel, stochastic volatility
JEL Classifications: C01, C14, C32

1 Introduction

Financial volatility is loosely speaking the standard deviation of asset returns. It has been known for a long time that volatility changes through time. Both ex-post and ex-ante volatility measures are in common use. The most well known ex-post measure is realised volatility, while ex-ante measures include those generated by ARCH type models and option based numbers such as implied volatility and the VIX. Reviews of this literature include, amongst others, Andersen, Bollerslev, and Diebold (2008), Barndorff-Nielsen and Shephard (2007), Taylor (2005) and Engle (2008).

In this article we provide a brief review of part of the literature on this topic, focusing on high frequency ex-post measures of volatility and models of volatility driven by Lévy processes. Throughout our discussion we think about a vector of efficient prices

$$Y_t = Y_0 + \int_0^t \mu_u du + \int_0^t \sigma_u dW_u, \quad t \geq 0, \quad (1)$$

where W_t is Brownian motion and μ_t, σ_t are spot drift and volatility, all adapted to $\mathcal{F}_t$ — the continuous time record of the past Y_t . For simplicity of exposition we will ignore jumps in the price process.

*Prepared for the “Encyclopedia of Quantitative Finance” (editor Rama Cont), John Wiley & Sons.

2 Ex-post measures

Ito's lemma implies

$$\begin{aligned}\text{Var} \left\{ \left(Y_t - Y_0 - \int_0^t \mu_u du \right) | \mathcal{F}_0 \right\} &= \text{E} \left\{ \int_0^t \sigma_u \sigma_u' du | \mathcal{F}_0 \right\} \\ &= \text{E} \{ [Y, Y]_t | \mathcal{F}_0 \},\end{aligned}$$

linking together ex-ante variance forecasts with ex-ante forecasts of the quadratic variation $[Y, Y]$ of Y . But we can use high frequency data to estimate the $[Y, Y]_t$ process, allowing us to project this quantity into the future without an explicit model for the spot volatility matrix σ_t . This method of constructing forecasts of volatility through QV was developed by Andersen, Bollerslev, Diebold, and Labys (2001).

There is an active literature on estimating the $[Y, Y]_t$ process. Most of it looks at estimating discrete increments of it over specified times periods, such as a day, $IV_i = [Y, Y]_i - [Y, Y]_{i-1}$. The most familiar estimator is the realised variance, for the i -th day, which is

$$RV_i = \sum_{i-1 < t_j \leq i} (Y_{t_j} - Y_{t_{j-1}}) (Y_{t_j} - Y_{t_{j-1}})', \quad (2)$$

where the t_j are times at which data is available — which could represent the times of trades, quotes (or a subset of such trades and quotes) or fixed intervals of time such as every 2 minutes. In the univariate case the square root of RV is the realised volatility (Rvol). Many option contracts are directly written on Rvol and RV, but usually this is where volatility is measured over the month and daily returns are used in (2).

The econometric theory associated with RV uses an in-fill asymptotic analysis, imagining $t_j - t_{j-1} \downarrow 0$. This is somewhat problematic as (1) ignores market microstructure effects which bite in the limit. However, for the moment ignore this. Then $RV_i \xrightarrow{p} IV_i$ by the definition of QV. The central limit theory extension of this, in the univariate case, is

$$\widehat{V}_i^{-1/2} \sqrt{n} (RV_i - IV_i) \xrightarrow{d} N(0, 1),$$

and was developed by Barndorff-Nielsen and Shephard (2002) (who suggested a variety of $\widehat{V}_i$ quantities). See also Jacod (1994) and Mykland and Zhang (2006). Figure 1 illustrates this (together with results from another estimator discussed below), showing 95% confidence intervals using 20 minute returns for General Electric, based on trades made on the stock exchange in November 2004.

Zhou (1998) was the first to formally study the properties of RV type statistics in the presence of market microstructure noise — which causes returns to be autocorrelated in time. Here we think of noise as $U = X - Y$ where X is the observed trade or quote process and Y continues to be the underlying efficient price. Modern research has developed three methods for being robust to noise:




Figure 1: Daily estimates of ex-post variability: realised variance and realised kernel, together with 95% confidence intervals.

preaveraging (Jacod, Li, Mykland, Podolskij, and Vetter (2007)), multiscale (Zhang, Mykland, and Aït-Sahalia (2005) and Zhang (2006)) and realised kernels (Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008a) and Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008b)).

We will focus on realised kernels, which generalise realised variance to

$$K(X) = \sum_{h=-H}^H k\left(\frac{h}{H+1}\right) \gamma_h, \quad \gamma_h = \sum_{j=|h|+1}^n r_j r_{j-|h|}$$

where $\{r_j\}$ are high frequency data and k is a Parzen weight function

$$k(x) = \begin{cases} 1 - 6x^2 + 6x^3, & 0 \leq x \leq 1/2, \\ 2(1 - x)^3, & 1/2 \leq x \leq 1, \\ 0, & x > 1. \end{cases}$$

This is like a HAC estimator, which is popular in econometrics, but there is no scaling by sample size. It is (i) robust to irregularly spaced data, (ii) endogenous noise, (iii) autocorrelation in the noise. A key feature of realised kernels is the choice of the bandwidth H . The value which approximately minimises the asymptotic MSE of the estimator is $H = 0.97\xi^{4/5}n^{3/5}$, $\xi^2 = \text{Var}(X - Y)/IV$. Figure 1 illustrates the use of this method, which shows the improvement from using the high frequency data. More details on how to implement this method are given in Barndorff-Nielsen, Hansen, Lunde, and Shephard (2009). The multivariate implementation is discussed in Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008b).

These high frequency data based volatility measures can be combined to make volatility forecasts in a variety of ways. Andersen, Bollerslev, Diebold, and Labys (2001) Andersen, Bollerslev,

Diebold, and Labys (2003) and Andersen, Bollerslev, Diebold, and Ebens (2001) have pioneered the use of simple linear time series methods. An alternative is to use lagged RV measures as an explanatory variable in the volatility dynamics of an ARCH model (e.g. Engle and Gallo (2006)). The latter approach is followed in the multivariate case by Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008b).

3 Pure jump based volatility models

There is now quite a lot of evidence that there are jumps in the volatility process (Todorov and Tauchen (2008)) and that jumps play a key role in generating statistical leverage effects (Barndorff-Nielsen, Kinnebrock, and Shephard (2009)). A paper which is a distinct break from the traditional diffusion based stochastic volatility models associated with Hull and White (1987), is Barndorff-Nielsen and Shephard (2001). In their simplest model they wrote the univariate spot variance as an OU type process

$$d\sigma_t^2 = -\lambda\sigma_t^2 dt + dZ_{\lambda t}, \quad \lambda > 0,$$

where Z is a subordinator. Recall a subordinator is a Lévy process with non-negative increments, which means σ_t^2 has no Brownian component at all. In the literature models of these type are called BNS models. They have the feature that if $\text{Var}(Z_1) < \infty$ then $\text{Corr}(\sigma_t^2, \sigma_{t+s}^2) = \exp(-\lambda s)$.

A main advantage of this OU process is that

$$\sigma_t^2 = e^{-\lambda t} \sigma_0^2 + \int_0^t e^{-\lambda(t-s)} dZ_{\lambda s}$$

and

$$\int_0^t \sigma_u^2 du = \lambda^{-1}(1 - e^{-\lambda t})\sigma_0^2 + \lambda^{-1} \int_0^t \{1 - e^{-\lambda(t-s)}\} dZ_{\lambda s}.$$

This makes them easy to analyse and simulate from.

More sophisticated dynamics can be achieved by adding together independent OU processes, each with a different decay parameter λ . Such superposition processes can be extended to the case of an infinite number of components, as analysed by Barndorff-Nielsen (2001).

A different route is to generate ARMA type dynamics. Such processes have been advocated by Brockwell (2001), Brockwell and Marquardt (2005) and Todorov and Tauchen (2006).

The OU process has been extended to the multivariate case by Barndorff-Nielsen and Stelzer (2007) and Pigorsch and Stelzer (2007). They set up an OU process for the positive semi-definite covariance matrix $\Sigma_t = \sigma_t \sigma_t'$ and analysed the properties of the resulting implied multivariate return process.

These BNS models have been used extensively in applications. Derivative pricing based on this type of model was first discussed by Nicolato and Venardos (2003). Subsequent research includes

Benth (2004). Benth, Groth, and Kufakunesu (2007) study the pricing of the variance and volatility swaps based on the BNS model. Minimal martingale measures for BNS models are analysed by Benth and Meyer-Brandis (2006) and Hubalek and Sgarra (2007). Benth, Groth, and Wallin (2007) study the calculation of Greeks for BNS type models.

Benth, Karlsen, and Reikvam (2003) and Delong and Kluppelberg (2007) have used BNS models to provide analytic solutions to the portfolio allocation problem.

These models and variants based off them have been used as models of the prices traded in energy markets where extreme movements are quite common. References include Benth, Kallsen, and Meyer-Brandis (2007) and Kluppelberg, Meyer-Brandis, and Schmidt (2008).

References

- Andersen, T. G., T. Bollerslev, and F. X. Diebold (2008). Parametric and nonparametric measurement of volatility. In Y. Aït-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. Correction published in 2003, volume 98, page 501.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys (2003). Modeling and forecasting realized volatility. *Econometrica* 71, 579–625.
- Barndorff-Nielsen, O. E. (2001). Superposition of Ornstein-Uhlenbeck type processes. *Theory of Probability and its Applications* 46, 175–194.
- Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2008a). Designing realised kernels to measure the ex-post variation of equity prices in the presence of noise. *Econometrica* 76, 1481–1536.
- Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2008b). Multivariate realised kernels: consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading. Unpublished paper: Oxford-Man Institute, University of Oxford.
- Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2009). Realised kernels in practice: trades and quotes. *Econometrics Journal*. Forthcoming.
- Barndorff-Nielsen, O. E., S. Kinnebrock, and N. Shephard (2009). Measuring downside risk: realised semivariance. In T. Bollerslev, J. Russell, and M. Watson (Eds.), *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*. Oxford University Press. Forthcoming.
- 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 (2002). 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 (2007). Variation, jumps and high frequency data in financial econometrics. In R. Blundell, T. Persson, and W. K. Newey (Eds.), *Advances in Economics and Econometrics. Theory and Applications, Ninth World Congress*, Econometric Society Monographs, pp. 328–372. Cambridge University Press.
- Barndorff-Nielsen, O. E. and R. Stelzer (2007). Positive-definite matrix processes. *Probability and Mathematical Statistics* 27, 3–43.
- Benth, F. E. (2004). *Option Theory with Stochastic Analysis*. Springer.
- Benth, F. E., M. Groth, and R. Kufakunesu (2007). Valuing volatility and variance swaps for a non-gaussian ornstein-uhlenbeck stochastic volatility model. *Applied Mathematical Finance* 14, 347–363.

- Benth, F. E., M. Groth, and O. Wallin (2007). Derivative-free Greeks for the Barndorff-Nielsen and Shephard stochastic volatility model. E-print, No. 3 March 2007, University of Oslo.
- Benth, F. E., J. Kallsen, and T. Meyer-Brandis (2007). A non-Gaussian Ornstein-Uhlenbeck process for electricity spot price modelling and derivative pricing. *Applied Mathematical Finance* 14, 153–169.
- Benth, F. E., K. H. Karlsen, and K. Reikvam (2003). Merton’s portfolio optimization problem in a Black and Scholes market with non-Gaussian stochastic volatility of Ornstein-Uhlenbeck type. *Mathematical Finance* 13, 215–244.
- Benth, F. E. and T. Meyer-Brandis (2006). The density process of the minimal entropy martingale measure in a stochastic volatility model with jumps. *Finance and Stochastics* 9, 563–575.
- Brockwell, P. J. (2001). Lévy driven CARMA processes. *Annals of the Institute of Statistical Mathematics* 52, 113–124.
- Brockwell, P. J. and T. Marquardt (2005). Lévy-driven and fractionally integrated ARMA processes with continuous time parameter. *Statistica Sinica* 15, 477–494.
- Delong, L. and C. Kluppelberg (2007). Optimal investment and consumption in a black-scholes market with stochastic coefficients driven by a non-gaussian ornstein-uhlenbeck process. Preprint, Technical University of Munich.
- Engle, R. F. (2008). *Anticipating Correlations*. Princeton University Press. Forthcoming.
- Engle, R. F. and J. P. Gallo (2006). A multiple indicator model for volatility using intra daily data. *Journal of Econometrics* 131, 3–27.
- Hubalek, F. and C. Sgarra (2007). On the Esscher transforms and the other equivalent martingale measures for the Barndorff-Nielsen and Shephard stochastic volatility models with jumps. Unpublished paper: Vienna University of Technology.
- Hull, J. and A. White (1987). The pricing of options on assets with stochastic volatilities. *Journal of Finance* 42, 281–300.
- Jacod, J. (1994). Limit of random measures associated with the increments of a Brownian semimartingale. Preprint number 120, Laboratoire de Probabilités, Université Pierre et Marie Curie, Paris.
- Jacod, J., Y. Li, P. A. Mykland, M. Podolskij, and M. Vetter (2007). Microstructure noise in the continuous case: the pre-averaging approach. Unpublished paper: Department of Statistics, University of Chicago.
- Kluppelberg, C., T. Meyer-Brandis, and A. Schmidt (2008). Electricity spot price modelling with a view towards extreme spike risk. Unpublished paper: Centre of Mathematical Sciences, Munich University of Technology.
- Mykland, P. A. and L. Zhang (2006). ANOVA for diffusions and Ito processes. *Annals of Statistics* 34, 1931–1963.
- Nicolato, E. and E. Venardos (2003). Option pricing in stochastic volatility models of the Ornstein-Uhlenbeck type. *Mathematical Finance* 13, 445–466.
- Pigorsch, C. and R. Stelzer (2007). A multivariate generalisation of the Ornstein-Uhlenbeck stochastic volatility model. Unpublished paper.
- Taylor, S. J. (2005). *Asset Price Dynamics, Volatility, and Prediction*. Princeton, New Jersey: Princeton University Press.
- Todorov, V. and G. Tauchen (2006). Simulation methods for Lévy-driven CARMA stochastic volatility models. *Journal of Business and Economic Statistics* 24, 450–469.
- Todorov, V. and G. Tauchen (2008). Volatility jumps. Unpublished paper: Department of Economics, Duke University.
- Zhang, L. (2006). Efficient estimation of stochastic volatility using noisy observations: a multi-scale approach. *Bernoulli* 12, 1019–1043.
- Zhang, L., P. A. Mykland, and Y. Aït-Sahalia (2005). A tale of two time scales: determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association* 100, 1394–1411.
- Zhou, B. (1998). Parametric and nonparametric volatility measurement. In C. L. Dunis and B. Zhou (Eds.), *Nonlinear Modelling of High Frequency Financial Time Series*, Chapter 6, pp. 109–123. New York: John Wiley Sons Ltd.