

Stochastic volatility

NEIL SHEPHARD

Nuffield College, University of Oxford OX1 1NF, UK

`neil.shephard@nuf.ox.ac.uk`

1 Overview

Stochastic volatility (SV) is the main concept used in the fields of financial economics and mathematical finance to deal with the endemic time-varying volatility and codependence found in financial markets. Such dependence has been known for a long time, early comments include Mandelbrot (1963) and Officer (1973). It was also clear to the founding fathers of modern continuous time finance that homogeneity was an unrealistic if convenient simplification, e.g. Black and Scholes (1972, p. 416) wrote “... there is evidence of non-stationarity in the variance. More work must be done to predict variances using the information available.” Heterogeneity has deep implications for the theory and practice of financial economics and econometrics. In particular, asset pricing theory is dominated by the idea that higher rewards may be expected when we face higher risks, but these risks change through time in complicated ways. Some of the changes in the level of risk can be modelled stochastically, where the level of volatility and degree of codependence between assets is allowed to change over time. Such models allow us to explain, for example, empirically observed departures from Black-Scholes-Merton prices for options and understand why we should expect to see occasional dramatic moves in financial markets.

The outline of this article is as follows. In section 2 I will trace the origins of SV and provide links with the basic models used today in the literature. In section 3 I will briefly discuss some of the innovations in the second generation of SV models. In section 4 I will briefly discuss the literature on conducting inference for SV models. In section 5 I will talk about the use of SV to price options. In section 6 I will consider the connection of SV with realised volatility. A extensive reviews of this literature is given in Shephard (2005).

2 The origin of SV models

The origins of SV are messy, I will give five accounts, which attribute the subject to different sets of people.

Clark (1973) introduced Bochner's (1949) time-changed Brownian motion (BM) into financial economics. He wrote down a model for the log-price M as

$$M_t = W_{\tau_t}, \quad t \geq 0, \quad (1)$$

where W is Brownian motion (BM), t is continuous time, τ is a time-change and $W \perp\!\!\!\perp \tau$, where $\perp\!\!\!\perp$ denotes independence. The definition of a time-change is a non-negative process with non-decreasing sample paths, although Clark also assumed τ has independent increments. Then $M_t | \tau_t \sim N(0, \tau_t)$. Further, so long (for each t) as $E\sqrt{\tau_t} < \infty$ then M is a martingale (written $M \in \mathcal{M}$) for this is necessary and sufficient to ensure that $E|M_t| < \infty$. More generally if (for each t) $\tau_t < \infty$ then M is a local martingale (written $M \in \mathcal{M}_{loc}$). Hence Clark was solely modelling the instantly risky component of the log of an asset price, written Y , which in modern semimartingale (written $Y \in \mathcal{SM}$) notation we would write as

$$Y = A + M.$$

The increments of A can be thought of as the instantly available reward component of the asset price, which compensates the investor for being exposed to the risky increments of M . The A process is assumed to be of finite variation (written $A \in \mathcal{FV}$).

To the best of my understanding the first published direct volatility clustering SV paper is that by Taylor (1982). His discrete time model of daily returns, computed as the difference of log-prices

$$y_i = Y_i - Y_{i-1}, \quad i = 1, 2, \dots,$$

where I have assumed that $t = 1$ represents one day to simplify the exposition. He modelled the risky part of returns, $m_i = M_i - M_{i-1}$ as a product process

$$m_i = \sigma_i \varepsilon_i. \quad (2)$$

Taylor assumed ε has a mean of zero and unit variance, while σ is some non-negative process, finishing the model by assuming $\varepsilon \perp\!\!\!\perp \sigma$. Taylor modelled ε as an autoregression and

$$\sigma_i = \exp(h_i/2),$$

where h is a non-zero mean Gaussian linear process. The leading example of this is the first order autoregression

$$h_{i+1} = \mu + \phi(h_i - \mu) + \eta_i, \quad \eta_i \sim NID(0, \sigma_\eta^2). \quad (3)$$

In the modern SV literature the model for ε is typically simplified to an i.i.d. process, for we deal with the predictability of asset prices through the A process rather than via M . This is

now often called the log-normal SV model in the case where ε is also assumed to be Gaussian. In general, M is always a local martingale.

A key feature of SV, which is not discussed by Taylor, is that it can deal with leverage effects. Leverage effects are associated with the work of Black (1976) and Nelson (1991), and can be implemented in discrete time SV models by negatively correlating the Gaussian ε_i and η_i . This still implies that $M \in \mathcal{M}_{loc}$, but allows the direction of returns to influence future movements in the volatility process, with falls in prices associated with rises in subsequent volatility.

Taylor's discussion of the product process was predated by a decade in the unpublished Rosenberg (1972). Rosenberg introduces product processes, empirically demonstrating that time-varying volatility is partially forecastable and so breaks with the earlier work by Clark. He suggests an understanding of aggregational Gaussianity of returns over increasing time intervals and predates a variety of econometric methods for analysing heteroskedasticity.

In continuous time the product process is the standard SV model

$$M_t = \int_0^t \sigma_s dW_s, \quad (4)$$

where the non-negative spot volatility σ is assumed to have càdlàg sample paths (which means it can possess jumps). The squared volatility process is often called the spot variance.

The first use of continuous time SV models in financial economics was, to my knowledge, in Johnson (1979) who studied the pricing of options using time-changing volatility models in continuous time (see also Johnson and Shanno (1987) and Wiggins (1987)). The most well known paper in this area is Hull and White (1987). Each of these authors desired to generalise the Black and Scholes (1973) approach to option pricing models to deal with volatility clustering. In the Hull and White approach σ^2 follows the solution to the univariate SDE

$$d\sigma^2 = \alpha(\sigma^2)dt + \omega(\sigma^2)dB,$$

where B is a second Brownian motion and $\omega(\cdot)$ is a non-negative deterministic function.

The probability literature has demonstrated that SV models and their time-changed BM relatives are fundamental. This theoretical development will be the fifth strand of literature that I think of as representing the origins of modern stochastic volatility research. Suppose we simply assume that $M \in \mathcal{M}_{loc}^c$, a process with continuous local martingale sample paths. Then the celebrated Dambis-Dubins-Schwartz Theorem shows that M can be written as a time-changed Brownian motion. Further the time-change is the quadratic variation (QV) process

$$[M]_t = \text{p-lim}_{n \rightarrow \infty} \sum_{j=1}^n (M_{t_j} - M_{t_{j-1}}) (M_{t_j} - M_{t_{j-1}})', \quad (5)$$

for any sequence of partitions $t_0 = 0 < t_1 < \dots < t_n = t$ with $\sup_j \{t_j - t_{j-1}\} \rightarrow 0$ for $n \rightarrow \infty$. What is more, as M has continuous sample paths, so must $[M]$. Under the stronger condition that $[M]$ is absolutely continuous, then M can be written as a stochastic volatility process. This latter result, which is called the martingale representation theorem, is due to Doob (1953). Taken together this implies that time-changed BMs are canonical in continuous sample path price processes and SV models are special cases of this class. A consequence of the fact that for continuous sample path time-change BM, $[M] = \tau$ is that in the SV case

$$[M]_t = \int_0^t \sigma_s^2 ds.$$

The SV framework has an elegant multivariate generalisation. In particular, write a p -dimensional price process M as (4) but where σ is a matrix process whose elements are all càdlàg, W is a multivariate BM process. Further $[M]_t = \int_0^t \sigma_s \sigma_s' ds$.

3 Second generation model building

3.1 Univariate models

3.1.1 General observations

In initial diffusion-based models the volatility was Markovian with continuous sample paths. Research in the late 1990s and early 2000s has shown that more complicated volatility dynamics are needed to model either options data or high frequency return data. Leading extensions to the model are to allow jumps into the volatility SDE (e.g. Barndorff-Nielsen and Shephard (2001) and Eraker, Johannes, and Polson (2003)) or to model the volatility process as a function of a number of separate stochastic processes or factors (e.g. Chernov, Gallant, Ghysels, and Tauchen (2003), Barndorff-Nielsen and Shephard (2001)).

3.1.2 Long memory

In the SV literature considerable progress has been made on working with both discrete and continuous time long memory SV. This involves specifying a long-memory model for σ in discrete or continuous time.

Breidt, Crato, and de Lima (1998) and Harvey (1998) looked at discrete time models where the log of the volatility was modelled as a fractionally integrated process. In continuous time there is work on modelling the log of volatility as fractionally integrated Brownian motion by Comte and Renault (1998). More recent work, which is econometrically easier to deal with, is the square root model driven by fractionally integrated BM introduced in an influential paper by Comte, Coutin, and Renault (2003) and the infinite superposition of non-negative OU processes introduced by Barndorff-Nielsen (2001).

3.1.3 Jumps

In detailed empirical work a number of researchers have supplemented standard SV models by adding jumps to the price process or to the volatility dynamics. Bates (1996) was particularly important as it showed the need to include jumps in addition to SV, at least when volatility is Markovian. Eraker, Johannes, and Polson (2003) deals with the efficient inference of these types of models. A radical departure in SV models was put forward by Barndorff-Nielsen and Shephard (2001) who suggested building volatility models out of pure jump processes called non-Gaussian OU processes. Closed form option pricing based on this structure is studied briefly in Barndorff-Nielsen and Shephard (2001) and in detail by Nicolato and Venardos (2003). All these non-Gaussian OU processes are special cases of the affine class advocated by Duffie, Pan, and Singleton (2000) and Duffie, Filipovic, and Schachermayer (2003).

3.2 Multivariate models

Diebold and Nerlove (1989) introduced volatility clustering into traditional factor models, which are used in many areas of asset pricing. In continuous time their type of model has the interpretation

$$M_t = \sum_{j=1}^J \int \beta_{(j)s} dF_{(j)s} + G_t,$$

where the factors $F_{(1)}, F_{(2)}, \dots, F_{(J)}$ are independent univariate SV models and G is correlated multivariate BM. Some of the related papers on the econometrics of this topic include King, Sentana, and Wadhvani (1994) and Fiorentini, Sentana, and Shephard (2004), who all fit this kind of model. These papers assume that the factor loading vectors are constant through time.

A more limited multivariate discrete time model was put forward by Harvey, Ruiz, and Shephard (1994) who allowed $M_t = C \int_0^t \sigma_s dW_s$, where σ is a diagonal matrix process and C is a fixed matrix of constants with a unit leading diagonal. This means that the risky part of prices is simply a rotation of a p -dimensional vector of independent univariate SV processes.

4 Inference based on return data

4.1 Moment based inference

The task is to carry out inference on $\theta = (\theta_1, \dots, \theta_K)'$, the parameters of the SV model based on a sequence of returns $y = (y_1, \dots, y_T)'$. Taylor (1982) and Melino and Turnbull (1990) calibrated their models using the method of moments. Systematic studies, using a GMM approach, of which moments to heavily weight in SV models was given in Andersen and Sørensen (1996), Genon-Catalot, Jeantheau, and Larédo (2000), Sørensen (2000) and Hoffmann (2002).

A difficulty with using moment based estimators for continuous time SV models is that it is not straightforward to compute the moments y . In the case of no leverage, general results for the second order properties of y and their squares were given in Barndorff-Nielsen and Shephard (2001). Some quite general results under leverage are also given in Meddahi (2001).

In the discrete time log-normal SV models the approach advocated by Harvey, Ruiz, and Shephard (1994) has been influential. Their approach was to remove the predictable part of the returns, so we think of $Y = M$ again, and work with $\log y_i^2 = h_i + \log \varepsilon_i^2$. If the volatility has short memory then this form of the model can be handled using the Kalman filter, while long memory models are often dealt with in the frequency domain. Either way this delivers a Gaussian quasi-likelihood which can be used to estimate the parameters of the model. The linearised model is non-Gaussian due to the long left hand tail of $\log \varepsilon_i^2$ which generates outliers when ε_i is small.

4.2 Simulation based inference

In the 1990s a number of econometricians started to use simulation based inference to tackle SV models. To discuss these methods it will be convenient to focus on the simplest discrete time log-normal SV model given by (2) and (3).

MCMC allows us to simulate from $\theta, h|y$, where $h = (h_1, \dots, h_T)'$. Discarding the h draws yields samples from $\theta|y$. Summarising yields fully efficient parametric inference. In an influential paper Jacquier, Polson, and Rossi (1994) implemented a MCMC algorithm for this problem. A subsequent paper by Kim, Shephard, and Chib (1998) gives quite an extensive discussion of various MCMC algorithms. This is a subtle issue and makes a very large difference to the computational efficiency of the methods (e.g. Jacquier, Polson, and Rossi (2003) and Yu (2005)).

Kim, Shephard, and Chib (1998) introduced the first filter using a so-called particle filter. As well as being of substantial scientific interest for decision making, the advantage of having a filtering method is that it allows us to compute marginal likelihoods for model comparison and one-step ahead predictions for model testing.

Although MCMC based papers are mostly couched in discrete time, a key advantage of the general approach is that it can be adapted to deal with continuous time models by the idea of augmentation. This was fully worked out in Elerian, Chib, and Shephard (2001), Eraker (2001) and Roberts and Stramer (2001).

A more novel non-likelihood approach was introduced by Smith (1993) and later developed by Gourieroux, Monfort, and Renault (1993) and Gallant and Tauchen (1996) into what is now called indirect inference or the efficient method of moments. Here I will briefly give a stylised version of this approach.

Suppose there is an auxiliary model for the returns (e.g. GARCH) whose density, $g(y; \psi)$, is easy to compute and, for simplicity of exposition, has $\dim(\psi) = \dim(\theta)$. Then compute its MLE, which we write as $\hat{\psi}$. We assume this is a regular problem so that $\partial \log g(y; \hat{\psi}) / \partial \psi = 0$ recalling that y is the observed return vector. Simulate a very long process from the SV model using parameters θ , which we denote by y^+ and evaluate the score not using the data but this simulation. This produces

$$\left. \frac{\partial \log g(y^+; \psi)}{\partial \psi} \right|_{\psi = \hat{\psi}}, \quad y^+ \sim f(y; \theta).$$

Then move θ around until the score is again zero, but now under the simulation. Write the point where this happens as $\tilde{\theta}$. It is called the indirect inference estimator.

5 Options

5.1 Models

SV models provide a basis for realistic modelling of option prices. We recall the central role played by Johnson and Shanno (1987) and Wiggins (1987). The most well known paper in this area is by Hull and White (1987), who looked at a diffusion volatility model with leverage effects. They assumed that volatility risk was unrewarded and priced their options either by approximation or by simulation. Hull and White (1987) indicated that SV models could produce smiles and skews in option prices, which are frequently observed in market data. The skew is particularly important in practice and Renault and Touzi (1996) prove that can be achieved in SV models via leverage effects.

The first analytic option pricing formulae were developed by Stein and Stein (1991) and Heston (1993). The only other closed form solution I know of is the one based on the Barndorff-Nielsen and Shephard (2001) class of non-Gaussian OU SV models. Nicolato and Venardos (2003) provide a detailed study of such option pricing solutions. See also the textbook exposition in Cont and Tankov (2004, Ch. 15). Slightly harder computationally to deal with is the more general affine class of models highlighted by Duffie, Filipovic, and Schachermayer (2003).

5.2 Econometrics of SV option pricing

In theory, option prices themselves should provide rich information for estimating and testing volatility models. I will discuss the econometrics of options in the context of the stochastic discount factor (SDF) approach, which has a long history in financial economics and is emphasised in, for example, Cochrane (2001) and Garcia, Ghysels, and Renault (2005). For simplicity I will assume interest rates are constant. We start with the standard Black-Scholes (BS) problem,

which will take a little time to recall, before being able to rapidly deal with the SV extension. We model

$$d \log Y = (r + p - \sigma^2/2) dt + \sigma dW, \quad d \log \widetilde{M} = h dt + b dW,$$

where \widetilde{M} is the SDF process, r the riskless short rate, and σ , h , b and p , the risk premium, are assumed constant for the moment.

We will price all contingent payoffs $g(Y_T)$ as $C_t = E\left(\frac{\widetilde{M}_T}{\widetilde{M}_t} g(Y_T) | \mathcal{F}_t\right)$, the expected discounted value of the claim where $T > t$. For this model to make financial sense we require that $\widetilde{M}_t Y_t$ and $\widetilde{M}_t \exp(tr)$ are local martingales, which is enough to mean that adding other independent BMs to the $\log \widetilde{M}$ process makes no difference to C or Y , the observables. These two constraints imply, respectively, $p + b\sigma = 0$ and $h = -r - b^2/2$. This means that (C^{BS}, Y) is driven by a single W .

When we move to the standard SV model we can remove this degeneracy. The functional form for the SV Y process is unchanged, but we now allow

$$d \log \widetilde{M} = h dt + a dB + b dW, \quad d\sigma^2 = \alpha dt + \omega dB,$$

where we assume that $B \perp\!\!\!\perp W$ to simplify the exposition. The SV structure will mean that p will have to change through time in response to the moving σ^2 . B is again redundant in the SDF (but not in the volatility) so the usual SDF conditions again imply $h = -r - \frac{1}{2}a^2$ and $p + b\sigma = 0$. This implies that the move to the SV case has little impact, except that the sample path of $\sigma^2 \perp\!\!\!\perp W$. So the generalised BS (GBS) price is

$$\begin{aligned} C_t^{GBS}(\sigma_t^2) &= E\left(\frac{\widetilde{M}_T}{\widetilde{M}_t} g(Y_T) | \mathcal{F}_t\right) \\ &= E\left\{C_t^{BS}\left(\frac{1}{T-t} \int_t^T \sigma_u^2 du\right) | \sigma_t^2, Y_t\right\}. \end{aligned}$$

Now C^{GBS} is a function of both Y_t and σ_t^2 , which means that (C^{GBS}, Y) is not degenerate. From an econometric viewpoint this is an important step, meaning inference on options is just the problem of making inference on a complicated bivariate diffusion process. When we allow leverage back into the model, the analysis becomes slightly more complicated algebraically.

In some recent work econometricians have been trying to use data from underlying assets and option markets to jointly model the dynamics of (C^{GBS}, Y) . The advantage of this joint estimation is that we can pool information across data types and estimate all relevant effects which influence Y , σ^2 and \widetilde{M} . Relevant papers include Chernov and Ghysels (2000), Pastorello, Patilea, and Renault (2003), Das and Sundaram (1999) and Bates (2000).

6 Realised volatility

The advent of very informative high frequency data has prompted econometricians to study estimators of the increments of the quadratic variation (QV) process and then to use this estimate to project QV into the future in order to predict future levels of volatility. The literature on this starts with independent, concurrent papers by Andersen and Bollerslev (1998), Barndorff-Nielsen and Shephard (2001) and Comte and Renault (1998). Some of this work echoes earlier important contributions from, for example, Rosenberg (1972) and Merton (1980).

A simple estimator of $[Y]$ is the *realised QV process*

$$[Y_\delta]_t = \sum_{j=1}^{\lfloor t/\delta \rfloor} (Y_{\delta j} - Y_{\delta(j-1)}) (Y_{\delta j} - Y_{\delta(j-1)})',$$

thus as $\delta \downarrow 0$ so $[Y_\delta]_t \xrightarrow{p} [Y]_t$. If $A \in \mathcal{FV}^c$, then $[Y] = [M]$, while if we additionally assume that M is SV then $[Y_\delta]_t \xrightarrow{p} \int_0^t \sigma_s \sigma_s' ds$.

In practice it makes sense to look at the increments of the QV process. Suppose we are interested in analysing daily return data, but in addition have higher frequency data measured at the time interval δ . The i -th daily realised QV is defined as

$$V(Y_\delta)_i = \sum_{j=1}^{\lfloor 1/\delta \rfloor} (Y_{i+\delta j} - Y_{i+\delta(j-1)}) (Y_{i+\delta j} - Y_{i+\delta(j-1)})' \xrightarrow{p} V(Y)_i = [Y]_i - [Y]_{i-1},$$

the i -th daily QV. The diagonal elements of $V(Y_\delta)_i$ are called realised variances and their square roots are called realised volatilities.

Andersen, Bollerslev, Diebold, and Labys (2001) have shown that to forecast the volatility of future asset returns, then a key input should be predictions of future daily QV. Recall, from Ito's formula, that if $Y \in \mathcal{SM}^c$ and $M \in \mathcal{M}$, then writing \mathcal{F}_t as the filtration generated by the continuous history of Y up to time t then

$$\mathbb{E}(y_i y_i' | \mathcal{F}_{i-1}) \simeq \mathbb{E}(V(Y)_i | \mathcal{F}_{i-1}).$$

A review of some of this material is given by Barndorff-Nielsen and Shephard (2005).

A difficulty with this line of argument is that the QV theory only tells us that $V(Y_\delta)_i \xrightarrow{p} V(Y)_i$, it gives no impression of the size of $V(Y_\delta)_i - V(Y)_i$. Jacod (1994) and Barndorff-Nielsen and Shephard (2002) have strengthened the consistency result to provide a univariate central limit theory

$$\frac{\delta^{-1/2} ([Y_\delta]_t - [Y]_t)}{\sqrt{2 \int_0^t \sigma_s^4 ds}} \xrightarrow{d} N(0, 1),$$

while giving a method for consistently estimating the integrated quarticity $\int_0^t \sigma_s^4 ds$ using high frequency data. This analysis was generalised to the multivariate case by Barndorff-Nielsen and

Shephard (2004a). This type of analysis greatly simplifies parametric estimation of SV models for we can now have estimates of the volatility quantities SV models directly parameterise. Barndorff-Nielsen and Shephard (2002), Bollerslev and Zhou (2002) and Phillips and Yu (2005) study this topic from different perspectives.

Recently there has been interest in studying the impact of market microstructure effects on the estimates of realised covariation. This causes the estimator of the QV to become biased. Leading papers on this topic are Zhou (1996), Fang (1996), Bandi and Russell (2003), Hansen and Lunde (2006) and Zhang, Mykland, and Ait-Sahalia (2005). Further, one can estimate the QV of the continuous component of prices in the presence of jumps using the so-called realised bipower variation process. This was introduced by Barndorff-Nielsen and Shephard (2004b) and Barndorff-Nielsen and Shephard (2006).

7 Acknowledgments

My research is supported by the UK's ESRC through the grant "High frequency financial econometrics based upon power variation."

References

- 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, 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. and B. Sørensen (1996). GMM estimation of a stochastic volatility model: a Monte Carlo study. *Journal of Business and Economic Statistics* 14, 328–352.
- Bandi, F. M. and J. R. Russell (2003). Microstructure noise, realized volatility, and optimal sampling. Unpublished paper, Graduate School of Business, University of Chicago.
- Barndorff-Nielsen, O. E. (2001). Superposition of Ornstein-Uhlenbeck type processes. *Theory of Probability and its Applications* 45, 175–194.
- 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 (2004a). Econometric analysis of realised covariation: high frequency covariance, regression and correlation in financial economics. *Econometrica* 72, 885–925.
- Barndorff-Nielsen, O. E. and N. Shephard (2004b). Power and bipower variation with stochastic volatility and jumps (with discussion). *Journal of Financial Econometrics* 2, 1–48.

- Barndorff-Nielsen, O. E. and N. Shephard (2005). Variation, jumps, market frictions and high frequency data in financial econometrics. Unpublished discussion paper: Nuffield College, Oxford.
- Barndorff-Nielsen, O. E. and N. Shephard (2006). Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics* 4. Forthcoming.
- Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *Review of Financial Studies* 9, 69–107.
- Bates, D. S. (2000). Post-'97 crash fears in the S-&P 500 futures option market. *Journal of Econometrics* 94, 181–238.
- Black, F. (1976). Studies of stock price volatility changes. *Proceedings of the Business and Economic Statistics Section, American Statistical Association*, 177–181.
- Black, F. and M. Scholes (1972). The valuation of option contracts and a test of market efficiency. *Journal of Finance* 27, 399–418.
- Black, F. and M. Scholes (1973). The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637–654.
- Bochner, S. (1949). Diffusion equation and stochastic processes. *Proceedings of the National Academy of Science of the United States of America* 85, 369–370.
- Bollerslev, T. and H. Zhou (2002). Estimating stochastic volatility diffusion using conditional moments of integrated volatility. *Journal of Econometrics* 109, 33–65.
- Breidt, F. J., N. Crato, and P. de Lima (1998). On the detection and estimation of long memory in stochastic volatility. *Journal of Econometrics* 83, 325–348.
- Chernov, M., A. R. Gallant, E. Ghysels, and G. Tauchen (2003). Alternative models of stock price dynamics. *Journal of Econometrics* 116, 225–257.
- Chernov, M. and E. Ghysels (2000). A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. *Journal of Financial Economics* 56, 407–458.
- Clark, P. K. (1973). A subordinated stochastic process model with fixed variance for speculative prices. *Econometrica* 41, 135–156.
- Cochrane, J. H. (2001). *Asset Pricing*. Princeton: Princeton University Press.
- Comte, F., L. Coutin, and E. Renault (2003). Affine fractional stochastic volatility models. Unpublished paper: University of Montreal.
- Comte, F. and E. Renault (1998). Long memory in continuous-time stochastic volatility models. *Mathematical Finance* 8, 291–323.
- Cont, R. and P. Tankov (2004). *Financial Modelling with Jump Processes*. London: Chapman and Hall.
- Das, S. R. and R. Sundaram (1999). Of smiles and smirks: A term structure perspective. *Journal of Financial and Quantitative Analysis* 34, 211–240.
- Diebold, F. X. and M. Nerlove (1989). The dynamics of exchange rate volatility: a multivariate latent factor ARCH model. *Journal of Applied Econometrics* 4, 1–21.
- Doob, J. L. (1953). *Stochastic Processes*. New York: John Wiley and Sons.
- Duffie, D., D. Filipovic, and W. Schachermayer (2003). Affine processes and applications in finance. *Annals of Applied Probability* 13, 984–1053.

- Duffie, D., J. Pan, and K. Singleton (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68, 1343–1376.
- Elerian, O., S. Chib, and N. Shephard (2001). Likelihood inference for discretely observed non-linear diffusions. *Econometrica* 69, 959–993.
- Eraker, B. (2001). Markov chain Monte Carlo analysis of diffusion models with application to finance. *Journal of Business and Economic Statistics* 19, 177–191.
- Eraker, B., M. Johannes, and N. G. Polson (2003). The impact of jumps in returns and volatility. *Journal of Finance* 53, 1269–1300.
- Fang, Y. (1996). Volatility modeling and estimation of high-frequency data with Gaussian noise. Unpublished Ph.D. thesis, Sloan School of Management, MIT.
- Fiorentini, G., E. Sentana, and N. Shephard (2004). Likelihood-based estimation of latent generalised ARCH structures. *Econometrica* 72, 1481–1517.
- Gallant, A. R. and G. Tauchen (1996). Which moments to match. *Econometric Theory* 12, 657–81.
- Garcia, R., E. Ghysels, and E. Renault (2005). The econometrics of option pricing. In Y. Ait-Sahalia and L. P. Hansen (Eds.), *Handbook of Financial Econometrics*. North Holland. Forthcoming.
- Genon-Catalot, V., T. Jeantheau, and C. Larédo (2000). Stochastic volatility as hidden Markov models and statistical applications. *Bernoulli* 6, 1051–1079.
- Gourieroux, C., A. Monfort, and E. Renault (1993). Indirect inference. *Journal of Applied Econometrics* 8, S85–S118.
- Hansen, P. R. and A. Lunde (2006). Realized variance and market microstructure noise (with discussion). *Journal of Business and Economic Statistics* 24. Forthcoming.
- Harvey, A. C. (1998). Long memory in stochastic volatility. In J. Knight and S. Satchell (Eds.), *Forecasting Volatility in Financial Markets*, pp. 307–320. Oxford: Butterworth-Heinemann.
- Harvey, A. C., E. Ruiz, and N. Shephard (1994). Multivariate stochastic variance models. *Review of Economic Studies* 61, 247–264.
- 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.
- Hoffmann, M. (2002). Rate of convergence for parametric estimation in stochastic volatility models. *Stochastic Processes and their Application* 97, 147–170.
- 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.
- Jacquier, E., N. G. Polson, and P. E. Rossi (1994). Bayesian analysis of stochastic volatility models (with discussion). *Journal of Business and Economic Statistics* 12, 371–417.
- Jacquier, E., N. G. Polson, and P. E. Rossi (2003). Stochastic volatility models: Univariate and multivariate extensions. *Journal of Econometrics*. Forthcoming.
- Johnson, H. (1979). Option pricing when the variance rate is changing. Working paper, University of California, Los Angeles.

- Johnson, H. and D. Shanno (1987). Option pricing when the variance is changing. *Journal of Financial and Quantitative Analysis* 22, 143–151.
- Kim, S., N. Shephard, and S. Chib (1998). Stochastic volatility: likelihood inference and comparison with ARCH models. *Review of Economic Studies* 65, 361–393.
- King, M., E. Sentana, and S. Wadhvani (1994). Volatility and links between national stock markets. *Econometrica* 62, 901–933.
- Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business* 36, 394–419.
- Meddahi, N. (2001). An eigenfunction approach for volatility modeling. Unpublished paper: University of Montreal.
- Melino, A. and S. M. Turnbull (1990). Pricing foreign currency options with stochastic volatility. *Journal of Econometrics* 45, 239–265.
- Merton, R. C. (1980). On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics* 8, 323–361.
- Nelson, D. B. (1991). Conditional heteroskedasticity in asset pricing: a new approach. *Econometrica* 59, 347–370.
- Nicolato, E. and E. Venardos (2003). Option pricing in stochastic volatility models of the Ornstein-Uhlenbeck type. *Mathematical Finance* 13, 445–466.
- Officer, R. R. (1973). The variability of the market factor of the New York stock exchange. *Journal of Business* 46, 434–453.
- Pastorello, S., V. Patilea, and E. Renault (2003). Iterative and recursive estimation in structural non-adaptive models. *Journal of Business and Economic Statistics* 21, 449–509.
- Phillips, P. C. B. and J. Yu (2005). A two-stage realized volatility approach to the estimation for diffusion processes from discrete observations. Unpublished paper: Cowles Foundation for Research in Economics, Yale University.
- Renault, E. and N. Touzi (1996). Option hedging and implied volatilities in a stochastic volatility model. *Mathematical Finance* 6, 279–302.
- Roberts, G. O. and O. Stramer (2001). On inference for nonlinear diffusion models using the Hastings-Metropolis algorithms. *Biometrika* 88, 603–621.
- Rosenberg, B. (1972). The behaviour of random variables with nonstationary variance and the distribution of security prices. Working paper 11, Graduate School of Business Administration, University of California, Berkeley. Reprinted in Shephard (2005).
- Shephard, N. (2005). *Stochastic Volatility: Selected Readings*. Oxford: Oxford University Press.
- Smith, A. A. (1993). Estimating nonlinear time series models using simulated vector autoregressions. *Journal of Applied Econometrics* 8, S63–S84.
- Sørensen, M. (2000). Prediction based estimating equations. *The Econometrics Journal* 3, 123–147.
- Stein, E. M. and J. Stein (1991). Stock price distributions with stochastic volatility: an analytic approach. *Review of Financial Studies* 4, 727–752.
- Taylor, S. J. (1982). Financial returns modelled by the product of two stochastic processes — a study of daily sugar prices 1961–79. In O. D. Anderson (Ed.), *Time Series Analysis: Theory and Practice*, 1, pp. 203–226. Amsterdam: North-Holland.

- Wiggins, J. B. (1987). Option values under stochastic volatilities. *Journal of Financial Economics* 19, 351–372.
- Yu, J. (2005). On leverage in a stochastic volatility model. *Journal of Econometrics* 127, 165–178.
- Zhang, L., P. 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*. Forthcoming.
- Zhou, B. (1996). High-frequency data and volatility in foreign-exchange rates. *Journal of Business and Economic Statistics* 14, 45–52.