Stochastic Volatility: Origins and Overview*

NEIL SHEPHARD
Oxford-Man Institute, University of Oxford, Oxford, UK;
Department of Economics, University of Oxford, Oxford, UK

TORBEN G. ANDERSEN
Kellogg School of Management, Northwestern University, Evanston, IL, USA;
CREATES, University of Aarhus, Aarhus, Denmark
NBER, Cambridge, MA, USA;

March 3, 2008

1 Introduction

Stochastic volatility (SV) models are used heavily within the fields of financial economics and mathematical finance to capture the impact of time-varying volatility on financial markets and decision making. The development of the subject has been highly multidisciplinary, with results drawn from financial economics, probability theory and econometrics blending to produce methods that aid our understanding of option pricing, efficient portfolio allocation and accurate risk assessment and management.

Time-varying volatility is endemic in financial markets. This has been known for a long time, with early comments including Mandelbrot (1963) and Fama (1965). 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 implies that higher rewards are required as an asset is exposed to more systematic risk. Of course, such risks may change through time in complicated ways, and it is natural to build stochastic models for the temporal evolution in volatility and codependence across assets. This allow us to explain, for example, empirically observed departures from Black-Scholes-Merton option prices and understand why we should expect to see occasional dramatic moves in financial markets. More generally, they bring the application of financial economics closer to the empirical reality of the world we live in, allowing us to make better decisions, inspire new theory and improve model building.

---

*Chapter prepared for the Handbook of Financial Time Series, Springer Verlag.

Autoregressive conditional heteroskedasticity (ARCH) processes are often described as SV, but we do not follow that nomenclature. The essential feature of ARCH models is that they explicitly model the conditional variance of returns given past returns observed by the econometrician. This one-step-ahead prediction approach to volatility modeling is very powerful, particularly in the field of risk management. It is convenient from an econometric viewpoint as it immediately delivers the likelihood function as the product of one-step-ahead predictive densities.

In the SV approach the predictive distribution of returns is specified indirectly, via the structure of the model, rather than explicitly. For a small number of SV models this predictive distribution can be calculated explicitly but, invariably, for empirically realistic representations it has to be computed numerically. This move away from direct one-step-ahead predictions has some advantages. In particular, in continuous time it is more convenient, and perhaps more natural, to model directly the volatility of asset prices as having its own stochastic process without worrying about the implied one-step-ahead distribution of returns recorded over an arbitrary time interval convenient for the econometrician, such as a day or a month. This does, however, raise some difficulties as the likelihood function for SV models is not directly available, much to the frustration of econometricians in the late 1980s and early 1990s.

Since the mid-1980s continuous-time SV has dominated the option pricing literature but early on econometricians struggled with the difficulties of estimating and testing these models. Only in the 1990s were novel simulation strategies developed to efficiently estimate SV models. These computationally intensive methods enable us, given enough coding and computing time, to efficiently estimate a broad range of fully parametric SV models. This has lead to refinements of the models, with many earlier tractable models being rejected from an empirical viewpoint. The resulting enriched SV literature has brought us much closer to the empirical realities we face in financial markets.

From the late 1990s SV models have taken center stage in the econometric analysis of volatility forecasting using high-frequency data based on realized volatility and related concepts. The reason is that the econometric analysis of realized volatility is tied to continuous-time processes, so SV is central. The close connection between SV and realized volatility has allowed financial econometricians to harness the enriched information set available through high-frequency data to improve, by an order of magnitude, the accuracy of their volatility forecasts over that traditionally offered by ARCH models based on daily observations. This has broadened the applications of SV into the important arena of risk assessment and asset allocation.

Below, we provide a selective overview of the SV literature. The exposition touches on models, inference, options pricing and realized volatility. The SV literature has grown organically, with a

variety of contributions playing important roles for particular branches of the literature, reflecting the highly multidisciplinary nature of the research.

2 The origin of SV models

The modern treatment of SV is typically cast in continuous time, but many older contributions employ discrete-time models. Specifically, the early econometric studies tended to favor discrete-time specifications, while financial mathematicians and financial economists often cast the problems in a diffusive setting when addressing portfolio choice and derivatives pricing. In response, econometricians have more recently developed practical inference tools for continuous-time SV models. We start with a description of some important early studies cast in a discrete-time setting and then cover the continuous-time formulations.

A central intuition in the SV literature is that asset returns are well approximated by a mixture distribution where the mixture reflects the level of activity or news arrivals. Clark (1973) originates this approach by specifying asset prices as subordinated stochastic processes directed by the increments to an underlying activity variable. Ignoring mean returns and letting the directing process being independent of the return innovations he stipulates,

$$Y_i = X_{\tau_i}, \quad i = 0, 1, 2, ..., \tag{1}$$

where Y_i denotes the logarithmic asset price at time i and $y_i = Y_i - Y_{i-1}$ the corresponding continuously compounded return over $[i-1, i]$, X_i is a normally distributed random variable with mean zero, variance $\sigma_X^2 \cdot i$, and independent increments, and τ_i is a real-valued process initiated at $\tau_0 = 0$ with non-negative and non-decreasing sample paths, i.e., it constitutes a time change. Clark focuses on the case where the increments to τ_i represent independent draws from a stationary distribution with finite variance, implying the subordinated return process also has independent increments with zero mean. More generally, as long as the time change process is independent of the price innovations, the asset returns are serially uncorrelated, albeit dependent, even if the time change increments are not stationary or independent. In fact, we have

$$y_i | (\tau_i - \tau_{i-1}) \sim N(0, \sigma_X^2 \cdot (\tau_i - \tau_{i-1})). \tag{2}$$

Thus, marginally, the asset returns follow a normal mixture, implying a symmetric but fat tailed distribution. The directing or mixing process, $\tau_t, t \geq 0$,, is naturally interpreted as an indicator of the intensity of price-relevant information flow over the interval $[0, t]$. Specifications of this type are generally referred to as Mixture of Distributions Hypotheses (MDH). They induce heteroskedastic return volatility and, if the time-change process is positively serially correlated, also volatility

clustering. Clark explores the i.i.d. time-change specification only and relates the time-change to trading volume. Many subsequent studies pursue the serially correlated volatility extension empirically and seek to identify observable market proxies for the latent time-change process. Complete specification of the joint dynamic distribution of return variation and related market variables allows for a more structural oriented approach to stochastic volatility modeling, see, e.g., Epps and Epps (1976), Tauchen and Pitts (1983), Andersen (1996), and Leisenfeld (2001).

For future reference, it is convenient to cast the Clark formulation in equivalent continuous-time notation. To emphasize that the log-price process as specified as a martingale, we denote it M. We may then restate equation (1) in a manner which implies the identical distribution for discretely sampled data,

$$M_t = W_{\tau_t}, \quad t \geq 0, \tag{3}$$

where W is Brownian motion (BM) and W and τ are independent processes. Technically, as long as (for each t) $\mathrm{E}\sqrt{\tau_t} < \infty$, M is a martingale since this is necessary and sufficient to ensure that $\mathrm{E}|M_t| < \infty$.

Asset pricing theory asserts that securities exposed to systematic risk have expected positive excess returns relative to the risk-free interest rate. As a result, asset prices will not generally be martingales. Instead, assuming frictionless markets, a weak no-arbitrage condition implies that the asset price will be a special semimartingale, see, e.g., Back (1991). This leads to the more general formulation,

$$Y = Y_0 + A + M, \tag{4}$$

where the finite variation process, A, constitutes the expected mean return. If the asset represents a claim on the broader market portfolio, a simple and popular specification for A is $A_t = r_f t + \beta \tau_t$, with r_f denoting the risk-free rate and β representing a risk premium due to the undiversifiable variance risk. This means that the distributional MDH result in equation (2) generalizes to $Y_t | \tau_t \sim N(r_f t + \beta \tau_t, \tau_t)$.

Clark's main purpose was to advocate the MDH as an alternative to the empirically less attractive stable processes. Although his framework lends itself to the appropriate generalizations, he did not seek to accommodate the persistence in return volatility. In fact, only about a decade later do we find a published SV paper explicitly dealing with volatility clustering, namely Taylor (1982). Taylor models the risky part of returns as a product process,

$$m_i = M_i - M_{i-1} = \sigma_i \varepsilon_i. \tag{5}$$

4

ε is assumed to follow an autoregression with zero mean and unit variance, while σ is some non-negative process. He completes the model by assuming $\varepsilon \perp\!\!\!\perp \sigma$ and

$$\sigma_i = \exp(h_i/2), \tag{6}$$

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

$$h_{i+1} = \mu + \phi\left(h_i - \mu\right) + \eta_i, \tag{7}$$

where η is a zero mean, Gaussian white noise process. In the modern SV literature the model for ε is typically simplified to an i.i.d. process, as the predictability of asset prices is incorporated in the A process rather than in M. The resulting model is now often called the log-normal SV model if ε is also assumed to be Gaussian. Finally, we note that M is a martingale as long as $\mathrm{E}\left(\sigma_i\right) < \infty$, which is satisfied for all models considered above if h is stationary.[1]

A key feature of SV, not discussed by Taylor, is that it can accommodate an asymmetric return-volatility relation, often termed a statistical leverage effect in reference to Black (1976), even if it is widely recognized that the asymmetry is largely unrelated to any underlying financial leverage. The effect can be incorporated in discrete-time SV models by negatively correlating the Gaussian ε_i and η_i so that the direction of returns impact future movements in the volatility process, with price drops associated with subsequent increases in volatility. Leverage effects also generate skewness, via the dynamics of the model, in the distribution of $(M_{i+s} - M_i)|\sigma_i$ for $s \geq 2$, although $(M_{i+1} - M_i)|\sigma_i$ continues to be symmetric. This is a major impetus for the use of these models in pricing of equity index options for which skewness appear endemic.

We now move towards a brief account of some early contributions to the continuous-time SV literature. In that context, it is useful to link the above exposition to the corresponding continuous-time specifications. The counterpart to the (cumulative) product process for the martingale component in equation (5) is given by the stochastic integral representation,

$$M_t = \int_0^t \sigma_s \mathrm{d}W_s, \tag{8}$$

where the non-negative spot volatility σ is assumed to have càdlàg sample paths. Note that this allows for jumps in the volatility process. Moreover, SV models given by (8) have continuous

---

[1]Taylor's discussion of the product process was predated by a decade in the unpublished Rosenberg (1972). This remarkable paper appears to have been lost to the modern SV literature until recently, but is now available in Shephard (2005). Rosenberg introduces product processes, empirically demonstrating that time-varying volatility is partially predictable, and thus moving beyond Clark's analysis on this critical dimension. He also explores a variety of econometric methods for analyzing heteroskedasticity only reintroduced into the literature much later. Finally, he studies an SV model which in some respects is a close precursor of ARCH models even if he clearly does not recognize the practical significance of restrictions on his system that would lead to an ARCH representation.

sample paths even if σ does not. A necessary and sufficient condition for M to constitute a martingale is that $\mathrm{E}\sqrt{\int_0^t \sigma_s^2 \mathrm{d}s} < \infty$. The squared volatility process is often termed the spot variance. There is no necessity for σ and W to be independent, but when they are we obtain the important simplification that $M_t | \int_0^t \sigma_s^2 \mathrm{d}s \sim N\left(0, \int_0^t \sigma_s^2 \mathrm{d}s\right)$. This makes it evident that the structure is closely related to the MDH or time-change representation (3) of Clark. The directing process is labeled Integrated Variance, i.e., $IV_t = \int_0^t \sigma_s^2 \mathrm{d}s$, and arises naturally as a quantity of key interest in practical applications.

An early application of continuous-time SV models was the unpublished work by Johnson (1979) who studied option pricing using time-changing volatility. While this project evolved into Johnson and Shanno (1987), a more well known paper in the use of continuous-time SV models for option pricing is Hull and White (1987) who allow the spot volatility process to follow a general diffusion. In their approach the spot variation process is given as the solution to a univariate stochastic differential equation,

$$\mathrm{d}\sigma^2 = \alpha(\sigma^2)\mathrm{d}t + \omega(\sigma^2)\mathrm{d}B, \tag{9}$$

where B is a second Brownian motion and $\omega(.)$ is a non-negative deterministic function. By potentially correlating the increments of W and B, Hull and White provide the first coherent leverage model in financial economics. They compute option prices by numerical means for the special case,

$$\mathrm{d}\sigma^2 = \alpha\sigma^2\mathrm{d}t + \omega\sigma^2\mathrm{d}B. \tag{10}$$

This formulation is closely related to the so-called GARCH diffusion which arises as the diffusion limit of a sequence of GARCH(1,1) models, see Nelson (1990), and has been used for volatility forecasting. Another related representation is the square-root process which belongs to the affine model class and allows for analytically tractable pricing of derivatives, as discussed in more detail later.

Wiggins (1987) also starts from the general univariate diffusion (9) but then focuses on the special case where log volatility follows a Gaussian Ornstein-Uhlenbeck (OU) process,

$$\mathrm{d}\log\sigma^2 = \alpha(\mu - \log\sigma^2)\mathrm{d}t + \omega\mathrm{d}B, \quad \alpha > 0. \tag{11}$$

The log-normal SV model of Taylor (1982) can be thought of as an Euler discretization to this continuous-time model over a unit time period. Ito's formula implies that this log-normal OU model can be written as

$$\mathrm{d}\sigma^2 = \left\{\theta - \alpha\log\sigma^2\right\}\sigma^2\mathrm{d}t + \omega\sigma^2\mathrm{d}B. \tag{12}$$

It is evident that it resembles the previous models in important respects although it is also distinctly different in the drift specification.

The initial diffusion-based SV models specify volatility to be Markovian with continuous sample paths. This is a constraint on the general SV structure (8) which requires neither of these assumptions. Research in the late 1990s and early 2000s has shown that more complex 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)).

A final noteworthy observation is that SV models and time-changed Brownian motions provide fundamental representations for continuous-time martingales. If M is a process with continuous martingale sample paths then the celebrated Dambis-Dubins-Schwartz Theorem, e.g., Rogers and Williams (1996, p. 64), ensures that M can be written as a time-changed BM with the time-change being the quadratic variation (QV) process,

$$[M]_t = \text{p} - \lim \sum_{j=1}^{n} \left(M_{t_j} - M_{t_{j-1}} \right)^2,\tag{13}$$

for any sequence of partitions $t_0 = 0 < t_1 < ... < t_n = t$ with $\sup_j \{t_j - t_{j-1}\} \to 0$ for $n \to \infty$. What is more, as M has continuous sample paths, so must $[M]$. Under the stronger condition that $[M]$ is absolutely continuous, M can be written as a stochastic volatility process. This latter result, known as 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 arise as special cases. In the SV case we thus have,

$$[M]_t = \int_0^t \sigma_s^2 \mathrm{d}s.\tag{14}$$

Hence, the increments to the quadratic variation process are identical to the corresponding integrated return variance generated by the SV model.

3 Second generation model building

3.1 Univariate models

3.1.1 Jumps

All the work discussed previously assumes that the asset price process is continuous. Yet, theory asserts that discrete changes in price should occur when significant new information is revealed. In fact, equity indices, Treasury bonds and foreign exchange rates all do appear to jump at the

moment significant macroeconomic or monetary policy news are announced. Likewise, individual stock prices often react abruptly to significant company-specific news like earnings reports, see, e.g. Andersen, Bollerslev, Diebold, and Vega (2007) and Johannes and Dubinsky (2006). As long as these jumps are unknown in terms of timing and/or magnitude this remains consistent with the no-arbitrage semimartingale setting subject only to weak regularity conditions. The cumulative sum of squared price jumps contribute to the return quadratic variation, thus generating distinct diffusive (integrated variance) and jump components in volatility.

Moreover, empirical work using standard SV models, extended by adding jumps to the price process, document significant improvements in model fit, e.g., Andersen, Benzoni, and Lund (2002) and Eraker, Johannes, and Polson (2003). This follows, of course, earlier theoretical work by Merton (1976) on adding jumps to the Black-Scholes diffusion. Bates (1996) was particularly important for the option pricing literature as he documents the need to include jumps in addition to SV for derivatives pricing, at least when volatility is Markovian.

Another restrictive feature of the early literature was the absence of jumps in the diffusive volatility process. Such jumps are considered by Eraker, Johannes, and Polson (2003) who deem this extension critical for adequate model fit. A very different approach for SV models was put forth by Barndorff-Nielsen and Shephard (2001) who build volatility models from pure jump processes. In particular, in their simplest model, σ^2 represent the solution to the SDE

$$\mathrm{d}\sigma_t^2 = -\lambda\sigma_t^2\mathrm{d}t + \mathrm{d}z_{\lambda t}, \quad \lambda > 0, \tag{15}$$

where z is a subordinator with independent, stationary and non-negative increments. The unusual timing convention for $z_{\lambda t}$ ensures that the stationary distribution of σ^2 does not depend on λ. These non-Gaussian OU processes are analytically tractable as they belong to the affine model class discussed below.

Geman, Madan, and Yor (2002) provide a new perspective within the general setting by defining the martingale component of prices as a time-change Lévy process, generalizing Clark's time-change of Brownian motion. Empirical evidence in Barndorff-Nielsen and Shephard (2006) suggest these rather simple models may potentially perform well in practice. Note, if one builds the time-change of the pure jump Lévy process from of an integrated non-Gaussian OU process then the resulting process will not have any Brownian components in the continuous-time price process.

3.1.2 Long memory

In the first generation of SV models the volatility process was given by a simple SDE driven by a BM. This implies that spot volatility is a Markov process. There is considerable empirical evidence that, whether volatility is measured using high-frequency data over a few years or using daily data

8

recorded over decades, the dependence in the volatility structure decays at a rapid rate for shorter lags, but then at a much slower hyperbolic rate at longer lags. Moreover, consistent with the hypothesis that long memory is operative in the volatility process, the estimates for the degree of fractional integration appear remarkably stable irrespective of the sampling frequencies of the underlying returns or the sample period, see Andersen and Bollerslev (1997). As an alternative, it is possible to approximate the long memory feature well by specifying the (log) volatility process via a sum of first-order autoregressive components, leading to multi-factor SV models as pursued by, e.g., Chernov, Gallant, Ghysels, and Tauchen (2003).

The literature has been successful in directly accommodating the longer run volatility dependencies through both discrete-time and continuous-time long memory SV models. In principle, this is straightforward as it only requires specifying a long-memory model for σ. Breidt, Crato, and de Lima (1998) and Harvey (1998) study discrete-time models where log volatility is modeled as a fractionally integrated process. They show this can be handled econometrically by quasi-likelihood estimators which are computationally simple, although not fully efficient. In continuous time Comte and Renault (1998) model log volatility as a fractionally integrated BM. More recent work includes the infinite superposition of non-negative OU processes introduced by Barndorff-Nielsen (2001). The two latter models have the potential advantage that they potentially can be used for options pricing without excessive computational effort.

3.2 Multivariate models

Diebold and Nerlove (1989) cast a multivariate SV model within the factor structure used in many areas of asset pricing. Restated in continuous time, their model for the $(N \times 1)$ vector of martingale components of the log asset price vector takes the form,

$$M = \sum_{j=1}^{J} \left(\beta_{(j)} F_{(j)} \right) + G, \tag{16}$$

where the factors $F_{(1)}, F_{(2)}, ..., F_{(J)}$ are independent univariate SV models, $J < N$, and G is a correlated $(N \times 1)$ BM, and the $(N \times 1)$ vector of factor loadings, $\beta_{(j)}$, remains constant through time. This structure has the advantage that the martingale component of time-invariant portfolios assembled from such assets will inherit this basic factor structure. Related papers on the econometrics of this model structure and their empirical performance include King, Sentana, and Wadhwani (1994) and Fiorentini, Sentana, and Shephard (2004).

A more limited multivariate discrete-time model was put forth by Harvey, Ruiz, and Shephard (1994) who suggest having the martingale components be given as a direct rotation of a p-dimensional vector of univariate SV processes. Another early contribution was a multivariate

9

extension of Jacquier, Polson, and Rossi (1994) which evolved into Jacquier, Polson, and Rossi (1999). In recent years, the area has seen a dramatic increase in activity as is evident from the chapter on Multivariate SV in this Handbook by Chib, Omori and Asai.

4 Inference based on return data

4.1 Moment based inference

A long standing difficulty for applications based on SV models was that the models were hard to estimate efficiently in comparison with their ARCH cousins due to the latency of the volatility state variable. In ARCH models, by construction, the likelihood (or quasi-likelihood) function is readily available. In SV models this is not the case which early on inspired two separate approaches. First, there is a literature on computationally intensive methods which approximate the efficiency of likelihood-based inference arbitrarily well, but at the cost of using specialized and time-consuming techniques. Second, a large number of papers have built relatively simple, inefficient estimators based on easily computable moments of the model. We briefly review the second literature before focusing on the former. We will look at the simplification high frequency data brings to these questions in Section 6.

The task is to carry out inference based on a sequence of returns $y = (y_1, ..., y_T)'$ from which we will attempt to learn about $\theta = (\theta_1, ..., \theta_K)'$, the parameters of the SV model. The early SV paper by Taylor (1982) calibrated the discrete-time model using the method of moments. Melino and Turnbull (1990) improve the inference by relying on a larger set of moment conditions and combining them more efficiently as they exploit the generalized method of moments (GMM) procedure. The quality of the (finite sample) GMM inference is quite sensitive to both the choice of the number of moments to include and the exact choice of moments among the natural candidates. Andersen and Sørensen (1996) provide practical guidelines for the GMM implementation and illustrate the potentially sizeable efficiency gains in the context of the discrete-time lognormal SV model. One practical drawback is that a second inference step is needed to conduct inference regarding the realizations of the latent volatility process. A feasible approach is to use a linear Kalman filter approximation to the system, given the first stage point estimates for the parameters, and extract the volatility series from the filter. However, this is highly inefficient and the combination of a two-step approach and a relatively crude approximation renders it hard to assess the precision of the inference for volatility.

Harvey, Ruiz, and Shephard (1994) apply the natural idea of using the Kalman filter for joint quasi-likelihood estimation of the model parameters and the time-varying volatility for the lognormal SV model defined via (5) and (7). This method produces filtered as well as smoothed

estimates of the underlying volatility process. The main drawback is that the method is quite inefficient as the linearized system is highly non-Gaussian.

For continuous-time SV models, it is generally much harder to derive the requisite closed form solutions for the return moments. Nonetheless, Meddahi (2001) provides a general approach for generating moment conditions for the full range of models that fall within the so-called Eigenfunction SV class. A thorough account of the extensive literature on moment-based SV model inference, including simulation-based techniques, is given in the chapter by Eric Renault.

4.2 Simulation-based inference

Within the last two decades, a number of scholars have started to develop and apply simulation-based inference devices to tackle SV models. Concurrently two approaches were brought forward. The first was the application of Markov chain Monte Carlo (MCMC) techniques. The second was the development of indirect inference or the so-called efficient method of moments. To discuss these methods it is convenient to focus on the simplest discrete-time log-normal SV model given by (5) and (7).

MCMC allows us to simulate from high dimensional posterior densities, such as the smoothing variables $h|y, \theta$, where $h = (h_1, ..., h_T)'$ are the discrete time unobserved log-volatilities. Shephard (1993) notes that SV models are a special case of a Markov random field so MCMC can be used for simulation of $h|y, \theta$. Hence, the simulation output inside an EM algorithm can be used to approximate the maximum likelihood estimator of θ. However, the procedure converges slowly. Jacquier, Polson, and Rossi (1994) demonstrate that a more elegant inference may be developed by becoming Bayesian and using the MCMC algorithm to simulate from $h, \theta|y$. Once the ability to compute many simulations from this $T + K$ dimensional random variable (there are K parameters), one can discard the h variables and simply record the many draws from $\theta|y$. Summarizing these draws allows for fully efficient parametric inference in a relatively sleek way. Later, Kim, Shephard, and Chib (1998) provide an extensive discussion of alternative methods for implementing the MCMC algorithm. This is a subtle issue and can make a large difference to the computational efficiency of the methods.

Kim, Shephard, and Chib (1998) also introduce a genuine filtering method for recursively sampling from

$$h_1, ..., h_i|y_1, ..., y_{i-1}, \theta, \quad i = 1, 2, ..., T. \tag{17}$$

These draws enable estimation, by simulation, of $\mathrm{E}\left(\sigma_i^2|y_1, ..., y_{t-1}, \theta\right)$ as well as the corresponding density and the density of $y_i|y_1, ..., y_{t-1}, \theta$ using the so-called particle filter, see, e.g., Gordon, Salmond, and Smith (1993) and Pitt and Shephard (1999). These quantities are useful inputs for

financial decision making as they are derived conditional only on current information. Moreover, they allow for computation of marginal likelihoods for model comparison and for one-step-ahead predictions for specification testing.[2]. Although these MCMC based papers are couched in discrete time, it is also noted that the general approach can be adapted to handle models operating with data generated at higher frequencies through data augmentation. This strategy was implemented for diffusion estimation by Jones (1998), Eraker (2001), Elerian, Chib, and Shephard (2001), and Roberts and Stramer (2001).

The MCMC approach works effectively under quite general circumstances, although it is dependent on the ability to generate appropriate and efficient proposal densities for the potentially complex conditional densities that arise during the recursive sampling procedure. An alternative is to develop a method that maximizes a simulation based estimate of the likelihood function. This may require some case-by-case development but it has been implemented for a class of important discrete-time models by Danielsson and Richard (1993) using the Accelerated Gaussian Importance Sampler. The procedure was further improved through improved simulation strategies by Fridman and Harris (1998) and Leisenfeld and Richard (2003). A formal approach for simulated maximum likelihood estimation of diffusions is developed by Pedersen (1995) and simultaneously, with a more practical orientation, by Santa-Clara (1995). Later refinements and applications for SV diffusion models include Elerian, Chib, and Shephard (2001), Brandt and Santa-Clara (2002), Durham and Gallant (2002), and Durham (2003).

Another successful approach for diffusion estimation was developed via a novel extension to the Simulated Method of Moments of Duffie and Singleton (1993). Gourieroux, Monfort, and Renault (1993) and Gallant and Tauchen (1996) propose to fit the moments of a discrete-time auxiliary model via simulations from the underlying continuous-time model of interest, thus developing the approach into what is now termed Indirect Inference or the Efficient Method of Moments (EMM). The latter approach may be intuitively explained as follows. First, an auxiliary model is chosen to have a tractable likelihood function but with a generous parameterization that should ensure a good fit to all significant features of the time series at hand. For financial data this typically involves an ARMA-GARCH specification along with a dynamic and richly parameterized (semi-nonparametric or SNP) representation of the density function for the return innovation distribution. The auxiliary model is estimated by (quasi-) likelihood from the discretely observed data. This provides a set of score moment functions which, ideally, encode important information regarding the probabilistic structure of the actual data sample. Next, a very long sample is simulated from the continuous-time model. The underlying continuous-time parameters are varied in order to

---

[2]A detailed account of the Particle Filter is given by Johannes and Polson in this Handbook

produce the best possible fit to the quasi-score moment functions evaluated on the simulated data. If the underlying continuous-time model is correctly specified it should be able to reproduce the main features of the auxiliary score function extracted from the actual data. It can be shown, under appropriate regularity, that the method provides asymptotically efficient inference for the continuous-time parameter vector. A useful side-product is an extensive set of model diagnostics and an explicit metric for measuring the extent of failure of models which do not adequately fit the quasi-score moment function. Gallant, Hsieh, and Tauchen (1997) provide an in-depth discussion and illustration of the use of these methods in practice. Moreover, the task of forecasting volatility conditional on the past observed data (akin to filtering in MCMC) or extracting volatility given the full data series (akin to smoothing in MCMC) may be undertaken in the EMM setting through the reprojection method developed and illustrated in Gallant and Tauchen (1998).

An early use of Indirect Inference for SV diffusion estimation is Engle and Lee (1996) while EMM has been extensively applied with early work exploring short rate volatility (Andersen and Lund (1997)), option pricing under SV (Chernov and Ghysels (2000)), affine and quadratic term structure models (Dai and Singleton (2000), Ahn, Dittmar, and Gallant (2002)), SV jump-diffusions for equity returns (Andersen, Benzoni, and Lund (2002)) and term structure models with regime-shifts (Bansal and Zhou (2002)).

An alternative approach to estimation of spot volatility in continuous time is given by Foster and Nelson (1996). They develop an asymptotic distribution theory for a local variance estimator, computed from the lagged data,

$$\widehat{\sigma_t^2} = h^{-1} \sum_{j=1}^{M} \left(Y_{t-hj/M} - Y_{t-h(j-1)/M} \right)^2. \tag{18}$$

They study the behavior of the estimator as $M \to \infty$ and $h \downarrow 0$ under a set of regularity conditions, ruling out, e.g., jumps in price or volatility. This "double asymptotics" yields a Gaussian limit theory as long as $h \downarrow 0$ and $M \to \infty$ at the correct, connected rates. This is related to the realized volatility approach detailed in a separate section below although, importantly, the latter focuses on the integrated volatility rather than the spot volatility and thus avoids some of the implementation issues associated with the double limit theory.

5 Options

5.1 Models

As discussed previously, the main impetus behind the early SV diffusion models was the desire to obtain a realistic basis for option pricing. A particularly influential contribution was Hull and White (1987) who studied a diffusion with leverage effects. Assuming volatility risk is fully diversifiable,

they price options either by approximation or by simulation. The results suggest that SV models are capable of producing smiles and skews in option implied volatilities as often observed in market data. Renault (1997) studies these features systematically and confirms that smiles and smirks emerge naturally from SV models via leverage effects.

The first analytic SV option pricing formula is by Stein and Stein (1991) who model σ as a Gaussian OU process. European option prices may then be computed using a single Fourier inverse which, in this literature, is deemed "closed form." A conceptual issue with the Gaussian OU model is that it allows for a negative volatility process. Heston (1993) overcomes this by employing a version of the so-called square root volatility process. Bates (1996) extends the framework further to allow for jumps in the underlying price and shows that these are critical for generating a reasonable fit to option prices simultaneously across the strike and time-to-maturity spectrum. Another closed-form option pricing solution is given by Nicolato and Venardos (2003) who rely on the non-Gaussian OU SV models of Barndorff-Nielsen and Shephard (2001).

All models above belong to the affine class advocated by Duffie, Pan, and Singleton (2000). These models are used extensively because they provide analytically tractable solutions for pricing a wide range of derivative securities. The general case involves solving a set of ordinary differential equations inside a numerical Fourier inverse but this may be done quickly on modern computers. These developments have spurred more ambitious inference procedures for which the parameters of affine SV models for both the underlying asset and the risk-neutral dynamics governing market pricing are estimated jointly from data on options and the underlying. Chernov and Ghysels (2000) estimate the affine SV diffusions for the actual and risk-neutral measures simultaneously using EMM. Pan (2002) exploits at-the-money options while allowing for an affine SV jump-diffusion representation under the actual and risk-neutral measure. Her inference is conducted via GMM, exploiting the closed-form expressions for the joint conditional moment-generating function of stock returns and volatility developed in Duffie, Pan, and Singleton (2000); see also Singleton (2001). Eraker (2004) expands the model specification, using MCMC based inference, to include a wider cross-section of option strikes and allowing for jumps in the volatility process as well. Finally, it is possible to develop option pricing on time-change Lévy processes, see, e.g., Carr and Wu (2004) who develop the derivatives pricing in a setting inspired by Geman, Madan, and Yor (2002).

6 Realized volatility

A couple of relatively recent developments have moved SV models towards the center of volatility research. This process is related to the rapid increase in research under the general heading of realized volatility.

One major change is the advent of commonly available and very informative high-frequency data, such as minute-by-minute return data or entire records of quote and/or transaction price data for particular financial instruments. The first widely disseminated data of this type were foreign exchange quotes gathered by Olsen & Associates, discussed in detail in the seminal work of Dacorogna, Gencay, Müller, Olsen, and Pictet (2001). Later scholars started using tick-by-tick data from the main equity and futures exchanges in the U.S. and Europe. This naturally moved the perspective away from fixed time intervals, such as a day, and into the realm where, at least in theory, one thinks of inference regarding the price process over different horizons based on ever changing information sets. This type of analysis is, of course, ideally suited to a continuous-time setting as any finite-horizon distribution then, in principle, may be obtained through time aggregation. Moreover, this automatically ensures modeling coherence across different sampling frequencies. Hence, almost by construction, volatility clustering in continuous time points us towards SV models.

A related development is the rapidly accumulating theoretical and empirical research on how to exploit this high-frequency data to estimate 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. This literature deals with various aspect of so-called realized variation, also often more generically referred to as realized volatility. This section briefly introduces some of the main ideas, leaning on contributions from Andersen, Bollerslev, Diebold, and Labys (2001) and Barndorff-Nielsen and Shephard (2002). A more detailed account is given in the chapter by Andersen and Benzoni in this handbook.

In realized variation theory, high-frequency data are used to estimate the QV process. We let δ denote a time period between high-frequency observations and define the *realized QV process* as,

$$[Y_\delta]_t = \sum_{j=1}^{\lfloor t/\delta \rfloor} \{Y(\delta j) - Y(\delta (j-1))\}^2. \tag{19}$$

Then, by the definition of the QV process, as $\delta \downarrow 0$ so

$$[Y_\delta]_t \overset{p}{\to} [Y]_t, \tag{20}$$

which the probability literature has shown to be well behaved if Y is a semimartingale. If the expected return process has continuous sample paths, then $[Y] = [M]$, and if additionally M is a SV process then $[Y_\delta]_t \overset{p}{\to} \int_0^t \sigma_s^2 \mathrm{d}s$.

In practice, it is preferable to measure increments of the quadratic variation process over one full trading day (or week). This measure is often referred to as the daily realized variance while its square root then is denoted the daily realized volatility, following the terminology of the financial mathematics literature. This should not be confused with the more generic terminology that

refers to all transformations of realized quadratic variation measures as realized volatility. The main reason for aggregating the realized variation measures to a daily frequency is the presence of pronounced and systematic intraday patterns in return volatility. These stem from highly regular, but dramatic, shifts in the quote and transactions intensity across the trading day as well as the release of macroeconomic and financial news according to specific time tables. Often, new information creates short-run dynamics akin to a price discovery process with an immediate price jump followed by a brief burst in volatility, see, e.g., Andersen and Bollerslev (1998). As a result, the intraday volatility process displays rather extreme variation and contains various components with decidedly low volatility persistence. Consequently, the direct modeling of the ultra high-frequency volatility process is both complex and cumbersome. Yet, once the return variation process is aggregated into a time series of daily increments, the strong inter-daily dependence in return volatility is brought out very clearly as the systematic intraday variation, to a large extent, is annihilated by aggregation across the trading day. In fact, the evidence for inter-daily volatility persistence is particularly transparent from realized volatility series compared to the traditional volatility measures inferred from daily return data.

Andersen, Bollerslev, Diebold, and Labys (2001) show that a key input for forecasting the volatility of future asset returns should be predictions of the future daily quadratic return variation. Recall from Ito's formula that, if Y is a continuous sample path semimartingale then

$$Y_t^2 = [Y]_t + 2\int_0^t Y_s \mathrm{d}Y_s = [Y]_t + 2\int_0^t Y_s \mathrm{d}A_s + 2\int_0^t Y_s \mathrm{d}M_s. \tag{21}$$

Letting $\mathcal{F}_t$ denote the filtration generated by the continuous history of Y_t up to time t and exploiting that M is a martingale, we have

$$\mathrm{E}(Y_t^2|\mathcal{F}_0) = \mathrm{E}\left([Y]_t|\mathcal{F}_0\right) + 2\mathrm{E}\left(\int_0^t Y_s \mathrm{d}A_s|\mathcal{F}_0\right). \tag{22}$$

In practice, over small intervals of time, the second term is small, so that

$$\mathrm{E}(Y_t^2|\mathcal{F}_0) \simeq \mathrm{E}\left([Y]_t|\mathcal{F}_0\right). \tag{23}$$

This implies that forecasting future squared daily returns can be done effectively through forecasts for future realized QV increments. A natural procedure estimates a time series model directly from the past observable realized daily return variation and uses it to generate predictions for future realized variances, as implemented through an ARFIMA model for realized log volatility in Andersen, Bollerslev, Diebold, and Labys (2003). The incorporation of long memory through fractional integration proves particularly important for forecast performance while only a few autoregressive lags are needed to accommodate shorter run dependencies. Hence, long lags of appropriately

weighted (hyperbolic decaying) realized log volatilities prove successful in forecasting future volatility.

A potential concern with this approach is that the QV theory only tells us that $[Y_\delta] \overset{p}{\to} [Y]$, but does not convey information regarding the likely size of the measurement error, $[Y_\delta]_t - [Y]_t$. Jacod (1994) and Barndorff-Nielsen and Shephard (2002) strengthen the consistency result to provide a central limit theory for the univariate version of this object. They show that the measurement errors are asymptotically uncorrelated and

$$\frac{\delta^{-1/2}\left([Y_\delta]_t - [Y]_t\right)}{\sqrt{2\int_0^t \sigma_s^4 \mathrm{d}s}} \overset{d}{\to} N(0,1). \tag{24}$$

The latter also develop a method for consistently estimating the integrated quarticity, $\int_0^t \sigma_s^4 \mathrm{d}s$, from high-frequency data, thus enabling feasible inference on the basis of the above result. This analysis may help simplify parametric estimation as we obtain estimates of the key volatility quantities that SV models directly parameterize. In terms of volatility forecasting, the use of long lags of weighted realized volatilities tends to effectively diversify away the impact of measurement errors so that the predictive performance is less adversely impacted than one may suspect, see Andersen, Bollerslev, and Meddahi (2006).

In the very recent past there have been various elaborations to this literature. We briefly mention two. First, there has been interest in studying the impact of market microstructure effects on the estimates of realized variance. This causes the estimator of the QV to become biased. Leading papers on this topic are Hansen and Lunde (2006), Zhang, Mykland, and Aït-Sahalia (2005), Bandi and Russell (2006) and Barndorff-Nielsen, Hansen, Lunde, and Shephard (2006). Second, one can estimate the QV of the continuous component of prices in the presence of jumps using the so-called realized bipower variation process. This was introduced by Barndorff-Nielsen and Shephard (2004).

References

Ahn, D.-H., R. F. Dittmar, and A. R. Gallant (2002). Quadratic term structure models: Theory and evidence. *Review of Financial Studies 15*, 243–288.

Andersen, T. G. (1996). Return volatility and trading volume: an information flow interpretation of stochastic volatility. *Journal of Finance 51*, 169–204.

Andersen, T. G., L. Benzoni, and J. Lund (2002). An empirical investigation of continuous-time equity return models. *Journal of Finance 57*, 1239–1284.

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). Deutsche mark-dollar volatility: intraday activity patterns, macroeconomic announcements, and longer run dependencies. *Journal of Finance 53*, 219–265.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and 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.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and C. Vega (2007). Real-time price discovery in global stock, bond and foreign exchange markets. *Journal of International Economics 73*, 251–277.

Andersen, T. G., T. Bollerslev, and N. Meddahi (2006). Market microstructure noise and realized volatility forecasting. Unpublished paper: Department of Economics, Duke University.

Andersen, T. G. and J. Lund (1997). Estimating continuous-time stochastic volatility models of the short term interest rate. *Journal of Econometrics 2*, 343–77.

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.

Back, K. (1991). Asset pricing for general processes. *Journal of Mathematical Economics 20*, 371–395.

Bandi, F. M. and J. R. Russell (2006). Seperating microstructure noise from volatility. *Journal of Financial Economics 79*, 655–692.

Bansal, R. and H. Zhou (2002). Term structure of interest rates with regime shifts. *Journal of Finance 57*, 1997–2043.

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 (2006). Designing realised kernels to measure the ex-post variation of equity prices in the presence of noise. Unpublished paper: Nuffield College, Oxford.

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 (2004). 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 (2006). Impact of jumps on returns and realised variances: econometric analysis of time-deformed Lévy processes. *Journal of Econometrics 131*, 217–252.

Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options. *Review of Financial Studies 9*, 69–107.

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.

Brandt, M. W. and P. Santa-Clara (2002). Simulated likelihood estimation of diffusions with an application to exchange rates dynamics in incomplete markets. *Journal of Financial Economics 63*, 161–210.

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.

Carr, P. and L. Wu (2004). Time-changed Lévy processes and option pricing. *Journal of Financial Economics 71*, 113–141.

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.

Comte, F. and E. Renault (1998). Long memory in continuous-time stochastic volatility models. *Mathematical Finance 8*, 291–323.

Dacorogna, M. M., R. Gencay, U. A. Müller, R. B. Olsen, and O. V. Pictet (2001). *An Introduction to High-Frequency Finance*. San Diego: Academic Press.

Dai, Q. and K. J. Singleton (2000). Specification analysis of affine term structure models. *Journal of Finance 55*, 1943–1978.

Danielsson, J. and J. F. Richard (1993). Accelerated Gaussian importance sampler with application to dynamic latent variable models. *Journal of Applied Econometrics 8*, S153–S174.

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., J. Pan, and K. J. Singleton (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica 68*, 1343–1376.

Duffie, D. and K. J. Singleton (1993). Simulated moments estimation of markov models of asset prices. *Econometrica 61*, 929–952.

Durham, G. (2003). Likelihood-based specification analysis of continuous-time models of the short-term interest rate. *Journal of Financial Economics 70*, 463–487.

Durham, G. and A. R. Gallant (2002). Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes (with discussion). *Journal of Business and Economic Statistics 20*, 297–338.

Elerian, O., S. Chib, and N. Shephard (2001). Likelihood inference for discretely observed non-linear diffusions. *Econometrica 69*, 959–993.

Engle, R. F. and G. G. J. Lee (1996). Estimating diffusion models of stochastic volatility. In P. E. Rossi (Ed.), *Modelling Stock Market Volatility — Bridging the GAP to Continuous Time*, pp. 333–384. Academic Press.

Epps, T. W. and M. L. Epps (1976). The stochastic dependence of security price changes and transaction volumes: implications for the mixture-of-distributions hypothesis. *Econometrica 44*, 305–321.

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. (2004). Do stock prices and volatility jump? reconciling evidence from spot and option prices. *Journal of Finance 59*, 1367–1403.

Eraker, B., M. Johannes, and N. G. Polson (2003). The impact of jumps in returns and volatility. *Journal of Finance 53*, 1269–1300.

Fama, E. F. (1965). The behaviour of stock market prices. *Journal of Business 38*, 34–105.

Fiorentini, G., E. Sentana, and N. Shephard (2004). Likelihood-based estimation of latent generalised ARCH structures. *Econometrica 72*, 1481–1517.

Foster, D. P. and D. B. Nelson (1996). Continuous record asymptotics for rolling sample variance estimators. *Econometrica 64*, 139–174.

Fridman, M. and L. Harris (1998). A maximum likelihood approach for non-Gaussian stochastic volatility models. *Journal of Business and Economic Statistics 16*, 284–291.

Gallant, A. R., D. Hsieh, and G. Tauchen (1997). Estimation of stochastic volatility models with diagnostics. *Journal of Econometrics 81*, 159–192.

Gallant, A. R. and G. Tauchen (1996). Which moments to match. *Econometric Theory 12*, 657–81.

Gallant, A. R. and G. Tauchen (1998). Reprojection partially observed systems with applications to interest rate diffusions. *Journal of the American Statistical Association 93*, 10–24.

Geman, H., D. B. Madan, and M. Yor (2002). Stochastic volatility, jumps and hidden time changes. *Finance and Stochastics 6*, 63–90.

Gordon, N. J., D. J. Salmond, and A. F. M. Smith (1993). A novel approach to nonlinear and non-Gaussian Bayesian state estimation. *IEE-Proceedings F 140*, 107–113.

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*, 127–218.

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-Heineman.

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.

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 Probabilitiés, 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 (1999). Stochastic volatility: Univariate and multivariate extensions. Unpublished paper: GSB, University of Chicago and HEC, Montreal.

Johannes, M. and A. Dubinsky (2006). Earnings announcements and equity options. Unpublished paper: Columbia University.

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.

Jones, C. S. (1998). Bayesian estimation of continuous-time finance models. Working Paper, Simon Graduate School of Business, University of Rochester.

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. Wadhwani (1994). Volatility and links between national stock markets. *Econometrica 62*, 901–933.

Leisenfeld, R. (2001). A generalized bivariate mixture model for stock price volatility and trading volume. *Journal of Econometrics 104*, 141–178.

Leisenfeld, R. and J.-F. Richard (2003). Univariate and multivairate stochastic volatility models: Estimation and diagnostics. *Journal of Empirical Finance 10*, 505–531.

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. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics 3*, 125–144.

Nelson, D. B. (1990). ARCH models as diffusion approximations. *Journal of Econometrics 45*, 7–38.

Nicolato, E. and E. Venardos (2003). Option pricing in stochastic volatility models of the Ornstein-Uhlenbeck type. *Mathematical Finance 13*, 445–466.

Pan, J. (2002). The jump-risk premia implicit in option prices: evidence from an integrated time-series study. *Journal of Financial Economics 63*, 3–50.

Pedersen, A. R. (1995). A new approach to maximum likelihood estimation for stochastic differential equations on discrete observations. *Scandinavian Journal of Statistics 27*, 55–71.

Pitt, M. K. and N. Shephard (1999). Filtering via simulation: auxiliary particle filter. *Journal of the American Statistical Association 94*, 590–599.

Renault, E. (1997). Econometric models of option pricing errors. In D. M. Kreps and K. F. Wallis (Eds.), *Advances in Economics and Econometrics: Theory and Applications*, pp. 223–78. Cambridge: Cambridge University Press.

Roberts, G. O. and O. Stramer (2001). On inference for nonlinear diffusion models using the Hastings-Metropolis algorithms. *Biometrika 88*, 603–621.

Rogers, L. C. G. and D. Williams (1996). *Diffusions, Markov Processes and Martingales. Volume 2, Ito Calculus* (2 ed.). Chichester: Wiley.

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).

Santa-Clara, P. (1995). Simulated likelihood estimation of diffusions with an application to the short term interest rate. PhD Dissertation, INSEAD.

Shephard, N. (1993). Fitting non-linear time series models, with applications to stochastic variance models. *Journal of Applied Econometrics 8*, S135–52.

Shephard, N. (2005). *Stochastic Volatility: Selected Readings*. Oxford: Oxford University Press.

Singleton, K. J. (2001). Estimation of affine asset pricing models using the empirical characteristic function. *Journal of Econometrics 102*, 111–141.

Stein, E. M. and J. Stein (1991). Stock price distributions with stochastic volatility: an analytic approach. *Review of Financial Studies 4*, 727–752.

Tauchen, G. and M. Pitts (1983). The price variability-volume relationship on speculative markets. *Econometrica 51*, 485–505.

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.

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.